University of Reading

Remote Health Monitoring – A Systems Approach to Using IoT Technologies

Thesis submitted for the degree of Doctor of Philosophy

Department of Biomedical Engineering

School of Biological Sciences

Ian Keith Poyner

February 2023

Declaration of Authorship

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Ian Keith Poyner

Acknowledgements

My PhD journey has been long with several diversions and interruptions, and I am grateful for all the kindness and support of the people throughout my life who have made this indulgence possible for me.

At the forefront of this endeavour is Professor R. Simon Sherratt, the most amicable, supportive and knowledgeable tutor that I could have hope for. Every interaction with Simon has left me enthused and, hopefully, a little wiser. My appreciation also goes to other members of the Department of Biomedical Engineering, particularly Professor Faustina Hwang who has so joyfully guided me through the processes of a PhD.

I look back at the 9-year-old boy and if key teachers had not guided me, my life would have been very different. They may never know it, but Mrs Holden, Mr Bellhouse, Mr Huggins and many others have instilled self-control, the language of mathematics and a love of physics. My gratitude also to the colleagues in the wide variety of organisations that I have worked with who have taught me ‘how’ to do engineering in a practical sense.

Above all, my love to Avril and James. James is a more intuitive engineer than I will ever be. Avril, you are my love and moral compass. Now onwards to the next great adventure in our lives.

Abstract

Remote Health Monitoring (RHM) has benefitted greatly from powerful smartphones and the very high data rate mobile networks. However, RHM benefits may be unobtainable for many users living with health challenges or in remote areas not served by telecommunications companies. IoT (Internet of Things) systems may redress some of these inequalities and extend RHM to a much wider community of users.

This thesis takes a systems-engineering approach to consider the service as a whole, and identifies the regulatory, business and user needs, especially reliability and privacy of personal health information. Relevant frameworks and technical requirements are assessed for a constrained device, including energy efficiency and security.

IoT networks, such as LoRaWAN, provide options for low cost, low power data transfer which are secure and do not depend upon network operators, especially when transmitted via satellites. Additionally, Machine Learning (ML) on constrained embedded devices is now practical, further reducing the need to transmit data for off-board processing. However, challenges remain for providing reliable and adaptable services to users whose health, and potentially life, relies on RHM services. Regulators are providing guidelines, but it is probable that legislation may in future enforce this guidance.

A TI CC2652 board was used to practically measure the relative energy consumption of transmitting packets of data via Bluetooth Low Energy (BLE) compared to on-board processing. A BLE message with a data payload of MTU = 251 bytes consumes approximately 660 – 676 nJ, which will also be dependent upon transmitted signal strength. This equates approximately to the CPU processing 11,380 – 11,655 for-loops. This provides a metric by which specific on-board processing and machine learning strategies can be assessed as to their energy efficiencies compared to offloading the raw data for processing. Advancements in ML for edge devices, such as TinyML and TensorFlow Lite for Microcontrollers, may enable very specific models to be run on the device within this energy budget. For comparison, this is approximately 50 times lower than the energy consumption of a BLE triple advertisement by the SPHERE SPW-1 wearable which consumes between 37 µJ (at -20dBm) and 60 µJ (at 4 dBm).

Contents

List of Figures

List of Tables

Chapter 1 Introduction

1.1 Introduction

The objective of this research is to investigate the suitability of Internet of Things (IoT) technologies to support remote health monitoring (RHM), particularly for users outside of a hospital, care-home or smart-home. The potential advantages include greater usability for all users, especially those living with cognitive challenges or operating in remote environments, better informed decision-making for clinicians, carers and supervisors and overall service affordability, particularly in the Least Developed Countries (LDC). Two key concerns addressed are communications where users do not have access to broadband or mobile networks, and the security and privacy issues of transmitting Personal Health Information (PHI).

eHealth is seen as central to the future of health services worldwide; it is not limited to consumer electronics devices for simple health parameters measurement. Demographics and increases in long term chronic disease will require fundamental changes in the way the world considers healthcare. The current status is unsustainable with ever increasing costs. It is necessary to reduce resources on the healthcare system. eHealth is recognised as the application of technology that can help to monitor people's health to reduce these costs in two ways: attendance and prevention (Bellido-Outeirino. et al., 2009).

This research investigates Remote Health Monitoring (RHM), that is monitoring a person’s condition outside a medical institution or normal place of work. Unlike a smart-home, the users are supported even when they are away from their normal residence. A key feature of remote monitoring is that a user can have confidence in living a more normal life outside of institutional care, or working in remote environments, knowing that someone is there to watch over them (Bellido-Outeirino. et al., 2009).

Energy efficiency is a key requirement in order to prolong the life of a battery. It influences many of the design considerations. Energy consumption for CPU activity and transmitting data to the cloud are measured with the results providing a metric against which trade-offs between edge processing and cloud processing can be assessed.

1.2 Motivation

The World Health Organization (WHO) estimates that more than one billion people are in need of one or more assistive products and this number is projected to increase to beyond two billion by 2050. However, only one in ten people in need currently have access to assistive technology. This can result in exclusion, isolation and being locked into poverty. This research addresses how people living or working in rural or remote areas may benefit at least partially from some of the advances in remote healthcare enabled by RHM (remote health monitoring).

Healthcare technologies are enabling people living with conditions to remain independent with the knowledge that carers can be alerted if there is a problem. This is applicable across all demographics, such as a teenager coping with diabetes; a pregnant woman being supported by a remote midwife; or a person living independently with early onset dementia. Supporters can review progress and provide guidance or encouragement to remote users without one of the parties having to make a long, and potentially painful, journey to a community health hub. During outbreaks of infections data collected over wide areas and remote communities could help manage the spread and to prioritise relief.

In other contexts, RHM may benefit citizens and workers in a wider context, including:

- Emergency responders dealing with crises and major incidents where the telecommunications infrastructure has been disabled, such as recovery from earthquakes and tsunamis (Centelles et al., 2019).
- Remote workers operating in areas not served by traditional networks.
- People in remote areas/wilderness, such as youth groups undertaking outdoor challenges.

Examples of health conditions where RHM can support the user are as follows:

- Cardiovascular Disease (CVD), the main cause of death in the world representing 30% of all global deaths, can often be prevented with proper health care.
- Diabetes, with the number of people suffering expected to rise to 380 million by 2025, can be mitigated with frequent monitoring to enable proper dosing, reducing the risk of fainting and in later life blindness, loss of circulation and other complications.

- Other examples of diseases that would benefit from continuous or prolonged monitoring include hypertension, asthma, Alzheimer's disease, Parkinson's disease, renal failure, post-operative monitoring, stress-monitoring and prevention of sudden infant death syndrome (Latre et al., 2011).

AAL (Active and Assisted Living) solutions, including adopting smart homes technologies, have grown rapidly. However, simply adopting a smart-home may not work for some users due to a number of issues, including loss of support when the user leaves the home, affordability, linguistic or cognitive barriers and technical constraints, such as poor mobile communications or intermittent power supplies. IoT devices, coupled to LPWAN (low-power wide area network) communications may enable many more communities to enjoy at least some of the fundamental benefits of RHM.

In summary, this research assesses the potential for IoT technology to support the following remote health monitoring scenarios (Islam et al., 2015):

- Acute/critical care.
- Monitoring a range of vital signs (blood pressure, body temperature, oxygen saturation and ECG).
- Pattern of life for vulnerable persons, typically where another person or organisation has a duty of care.
- Disaster relief and emergency responders at major events.
- Monitoring or supervising people in remote areas, including lone workers and youth expeditions.

1.3 Publications

Based on the research done, this report draws upon two conference papers that have been published and a further two in draft:

> Poyner, I. K. and Sherratt, R. S. (2019) *Improving access to healthcare in rural communities - IoT as part of the solution.* In: 3rd IET International Conference on Technologies for Active and Assisted Living (TechAAL 2019), 25 Mar 2019, London, UK. doi: https://doi.org/10.1049/cp.2019.0104,(Poyner and Sherratt, 2019).

Poyner, I. and Sherratt, S. (2018) *Privacy and security of consumer IoT devices for the pervasive monitoring of vulnerable people.* In: Living in the Internet of Things: Cybersecurity of the IoT, 28 - 29 March 2018, Savoy Place, London, doi: 10.1049/cp.2018.0043 (Poyner and Sherratt).

Poyner, I. K. and Sherratt, R. S. (2019) *Healthcare in rural communities and developing countries – IoT solutions to improve access.* In: Living in the Internet of Things (PETRAS 2019), 1-2 May 2019, London, UK. (In Press)

Poyner, I. K. and Sherratt, R. S. (2019) *Satellite constellations – opportunities for IoT-enabled healthcare.* In: Living in the Internet of Things (PETRAS 2019), 1-2 May 2019, London, UK. (In Press)

The author has also contributed to industry best practice, including:

IoT Security Assurance Framework, Release 3.0, November 2021 (IoT Security Foundation, Edited by Hall, Poyner, Phillips, Marshall and Markall).

IoTSF Framework to NSIT IR8259A Informative Reference, release 1.0.0, 10 November 2020, Point of Contact Ian Poyner.

Some of these works are embedded in Chapter B

1.4 Contributions and Organisation of Thesis

Chapter 2 considers the role of patient-generated health data and who are the key stakeholders using the data and what are the workflows in which this information is used. A systems engineering approach is taken by firstly identifying the user, business and regulatory requirements for RHM services. The contribution of this work is to provide a systematic approach to identifying the requirements of an RHM service, rather than researching a specific technology in isolation.

Chapter 3 derives technical requirements following from the user requirements identified in the previous chapter. A key consideration is the security and privacy of personal health information and the frameworks that support this are discussed. A range of low-power IoT communications are examined including both licensed 3GPP and 5G solutions and unlicensed (but regulated) systems, generally operating in the ISM (Instrument, Scientific and Medical) bands. The chapter is concluded by looking at the emerging field of embedded machine learning on constrained IoT devices.

Chapter 4 identifies the SAPPHIRE wearable as a candidate RHM device to evaluate how performing on-board processing of data compares to the energy requirements of transmitting the data, via Bluetooth Low Energy (BLE) in this case, to the cloud or core servers for processing. The approach follows earlier work investigating the SPHERE wearable.

Chapter 5 is the practical measurement of the energy consumption of a device based on the SAPPHIRE wearable. CPU cycles processing a 'for-loop' simulate on-board edge processing or an embedded ML model. The contribution of this work is to provide a metric by which edge processing/ML can be assessed in terms of energy efficiency compared to offloading the data via BLE.

Chapter 6 presents the conclusions to this thesis. Chapter 7 then provides suggestions for further work arising from this research.

References are included in Chapter 8.

Appendix A contains the details of the results from the experimental work described in Chapter 5. Spreadsheets containing 2,500 data points from each measurement are embedded into the results tables, together with example images of the waveforms generated by the device and BLE messages.

Appendix B has embedded copies of papers published during this research and also two publications in which the author has contributed during the period of this PhD research.

Chapter 2 Summary of Literature Review

This Chapter firstly distinguishes between medical monitoring and fitness/wellness tracking and how Principal Generated Health Data (PGHD – referred to patient generated health data in the literature) can support Remote Health Monitoring (RHM) services.

Use cases are developed that are wider in scope than many of the remote health projects in the literature. For example, supervision of remote workers and monitoring of teams of young people on expeditions are considered.

The stakeholders across the whole RHM service are then identified, together with some of their potential concerns. This is followed by a discussion on workflows and the need for PGHD to integrated into existing systems, at least initially to facilitate widespread adoption of RHM. Existing solutions, such as AAL (Active/Ambient and Assisted Living), are reviewed to determine their appropriateness and any shortfalls.

An extensive examination of the user requirements is derived, based upon the 'PESTLE' (political, economic, sociocultural, technological, legal and environmental) framework. This explores the requirements from across all the stakeholders in a technology-agnostic approach, to draw out the key themes that are particularly relevant to RHM services.

Three technical approaches: Edge Processing, Low Power Wide Area Networks (LPWAN) and on-device cryptography, are then assessed as to how they could support solutions for the key user requirement themes identified above.

2.1 Medical and Health Remote Monitoring

The use of technology to contribute to an individual's and a community's overall health has grown rapidly. It is often termed eHealth (electronic health), although this has been used with many definitions and a wide range of stakeholders and technologies (Oh et al., 2005). A subset of eHealth is mHealth (mobile health), which again has no standardised definition, but the Global Observatory for eHealth (GOe) defines "*as medical and public health practice supported by mobile devices, such as mobile phones, patient monitoring devices, personal digital assistants (PDAs), and other wireless devices*" (WHO Global Observatory for eHealth, 2011). Often mHealth is associated directly with the use of a

mobile phone, for example "*the use of mobile communication technologies to promote health by supporting healthcare practices (e.g. health data collection, delivery of healthcare information, or patient observation and provision of care*)" (Aranda-Jan et al., 2014) or the "*application of mobile technologies to the health sector*" (Karageorgos et al., 2019).

To avoid confusion with mHealth, and to broaden the research beyond only medical and public health, the term **Remote Health Monitoring (RHM)** services will be adopted. RHM is not only associated with the technology, but addresses all the other components of providing a service including people (stakeholders), processes (workflows), information flows and facilities, sometimes referred to as PPITF or PPT. In this study, RHM focusses on when non-smartphone technologies may offer advantages.

There has been a large increase in the number of health monitoring applications brought to market over the last 15 years. The mass production of IoT components, such as accelerometers and temperature sensors used in a wide variety of devices, have made a wide range of health devices affordable. In addition, the mass adoption of smartphones has been a significant catalyst supporting fitness applications as they provide a convenient platform at a relatively low cost for hosting apps and sensors and communicating with cloud services for monitoring and analysis. Several authors (Akmandor and Jha, 2018, Karageorgos et al., 2019, Rashidi and Mihailidis, 2013) provide examples of smart health-care applications including:

- measuring physiological signals from the human body
- diagnosing diseases and enabling proactive prevention strategies
- monitoring postoperative conditions
- injecting pharmaceutical compounds into the body
- adverse drug reaction
- providing rehabilitation
- wandering prevention tools
- citizen education and behaviour change communication
- provider training, education and work planning
- supply chain management

However, many devices that are labelled or marketed as 'medical devices' are merely health monitoring aids, such as fitness and wellness apps on mobile phones. For example, of the many user-centred (as opposed to those used in a clinical environment) devices and apps that monitor a user's heart rhythm, only the KardiaMobile has been approved for the ambulatory detection of atrial fibrillation by the UK's NICE and US FDA (other devices such as the Carnation Ambulatory Monitor (NICE MIB276) and Zio XT (NICE MTG52) are intended as a service that monitors the heart for 14 days with the data assessed by cardiac technicians or clinicians (NICE MIB152)). In Great Britain, the Medical Devices Regulations 2002 (SI 2002 No 618, as amended) (UK MDR 2002) (Secretary of State for Health, 2002) defines a medical device as an instrument, apparatus, appliance, material or other article, whether used alone or in combination, together with any software necessary for its proper application, which—

(a) is intended by the manufacturer to be used for human beings for the purpose of-
   (i) diagnosis, prevention, monitoring, treatment or alleviation of disease,
   (ii) diagnosis, monitoring, treatment, alleviation of or compensation for an injury or handicap,
   (iii) investigation, replacement or modification of the anatomy or of a physiological process, or
   (iv) control of conception; and
(b) does not achieve its principal intended action in or on the human body by pharmacological, immunological or metabolic means, even if it is assisted in its function by such means.

Any device that is used to diagnose or provide therapy, administer medication or treatment must comply with the appropriate regulatory approvals. In Great Britain this is the UK MDR 2002, in the EU and Northern Ireland the EU Medical Devices Regulation (2017/745) applies and in the US the Food and Drug Administration (FDA) Center for Devices and Radiological Health (CDRH) is responsible for regulating medical devices through its 510(k) process. These regulatory approval processes have different levels of classification corresponding to the risks to the patient and user, and there are exemptions for the lowest risk situations. In addition, producers of medical devices may be required to implement a quality system, typically based on ISO/EN/IEEE 11073 Personal Health Device, ISO 13845:2016 Medical Devices – Quality Management Systems, the Medical

Device and Health IT Joint Security Plan (Healthcare and Public Health Sector Coordinating Council, 2019) and may also need to gain CE approval.

Related to the above, the terms health app and medical app are often used interchangeably but do not necessarily mean the same thing. Health apps are software programs on mobile devices that process health-related data for their users and can be used by a health-conscious person to maintain, improve, or manage the health of an individual or the community. Medical apps may share the same technological functions and devices, with health professionals, patients, and family caregivers being the main user groups. Medical apps are intended for clinical and medical purposes and can be legally regulated as mobile medical devices (Maaß et al., 2022).

For this research, the focus is on low-risk IoT devices that generally involve a clinician making the actual medical diagnosis or a user controlling the intervention, rather than the device autonomously determining and applying medication or therapy. Note also that for this research, only general medical devices (categorised under Part II of the UK MDR 2002) are considered because certification of active implantable medical devices (Part III) is much more demanding as they represent the highest risk category. In vitro diagnostic devices (Part IV) are also not appropriate to this research. The use of IoT devices in hospital, clinical or ambulance environments is also not in scope of this research.

It is critical that the distinction between a medical device and a health or fitness device should be understood at the outset of developing RHM service.

2.1.1 PGHD - Principal/Patient Generated Health Data

PGHD (principal/patient generated health data) – *"are health-related data—including health history, symptoms, biometric data, treatment history, lifestyle choices, and other information—created, recorded, gathered, or inferred by or from patients or their designees (i.e., care partners or those who assist them) to help address a health concern. PGHD are distinct from data generated in clinical settings and through encounters with providers in two important ways. First, patients, not providers, are primarily responsible for capturing or recording these data. Second, patients direct the sharing or distributing of these data to health care providers and other stakeholders. In these ways, PGHD*

complement provider-directed capture and flow of health-related data across the health care system" (Shapiro et al., 2012). This definition is broadened in this thesis to include physiological parameters of remote personal and emergency responders being monitored by their supervisors. Also, the term principal is used instead of patient to encompass the wider variety of users and that many wearers of an RHM do not see themselves as being a patient.

2.2 Stakeholders and Workflows

2.2.1 Stakeholders

Most studies on mHealth/RHM have focused on the technical development and adoption of mobile applications, but other aspects have received less attention (Petersen et al., 2015). When considering a service and its systems, and even more so when a complex system-of-systems is included, it is important to understand all the stakeholders directly impacted by the system of interest (SOI) and also to identify the actors in the wider system of interest (WSOI) that may impact the SOI. The stakeholders will vary widely upon jurisdictions, the conditions being monitored, the technologies used and other considerations.

Petersen et al (Petersen et al., 2015) identified stakeholders in mHealth systems. The terminology used has been adapted for this thesis to reflect the broader use cases beyond clinical cases and to make the language more inclusive as many users will not identify themselves as patients (Petersen's original terminology in brackets).

Table 2-1: Stakeholder List

Stakeholder	Description/examples
Principals (patients)	The main subjects of the system, whose physiological or pattern of life parameters are generating the data. Principal includes more than a person living with a medical condition, and could include: • At risk pregnant woman being remotely monitored by a midwife.

	• Assisted living. • Emergency responders dealing with major incidents where normal telecommunications infrastructure has been damaged. • Remote workers. • People in remote areas/wilderness, such as youth groups undertaking outdoor challenges.
Family and Caregivers	Families and others responsible for care, including carers employed to provide a service. They may be using RHM to be alerted to any sudden health concern or to monitor longer term patterns.
Clinicians	Can use RHM systems to access PGHD and, in some cases, to prescribe medications electronically. Their concerns may include costs, security and ease of use (integration with existing workflows). Clinicians are unlikely to have the time to review and respond to PGHD in near real-time, but liability for 'missing' any change in symptoms may vary with jurisdiction and use-case. Clinicians may also have concerns regarding the trustworthiness (reliability, accuracy) of a system for basing interventions.
Health care facilities	Hospitals, surgery centres, long-term care facilities, community group homes, home health agencies.
Sponsors and health insurers (payors and purchasers including health insurers).	The sponsor may include the principal or family, social services, youth group supervisors (schools), employers (for emergency responders and remote workers), as well as those who wish to improve wellness and reduce absenteeism amongst their workforce).
Regulators and standards bodies (policy actors)	Legislators, regulators and industry standards bodies, such as those who govern the use of data, licensing of medical devices and aligning of incentives.

	The wider health system may also influence how PGHD is integrated and exploited within EHR (electronic health records) and exchanged between clinical teams.
Researchers	May use data for clinical trials, comparative effectiveness research and other areas
Service providers/integrators (vendors, suppliers, app developers and consultants)	Wide variety of organisations and business models. May be contracted by the sponsors, be subject to SLA (service level agreements) and bill the sponsor. May employ/engage care-providers. Often will operate the back-end systems that receive PGHD from the device and presents it to the caregiver, family, or clinician. Likely to be a 'data controller' or 'data processor' of protected health information (PHI) under HIPAA or sensitive information under GDPR.
Device suppliers	Involved in the design, development, deployment and through-life support of hardware, firmware and applications.
Network operators	Network operators include traditional telecom companies providing landline services, mobile phone operators, IoT network operators (as a service company or ad-hoc private volunteers), satellite operators, Emergency Services Network and other specialised providers. A network operator may offer varying levels of quality of service (QoS) under a service level agreement (SLA) with the service provider or sponsor. In many other cases, users may rely simply on the standard offerings of the network operators with no guarantees of availability, security, latency or signal quality. A network operator may cover several regions or have roaming agreements that enable users to traverse other territories.

Out of scope are hospitals and emergency services (ambulances and paramedics), except for situations where teams are responding to major natural disasters and the infrastructure

has been compromised. This is because they usually expect to operate to a very high level of reliability and often already deploy and operate advanced monitoring and communications systems.

Principal	
Family and caregivers	
Clinicians	
Health care facilities	
Sponsors and health insurers	
Regulators and Standards Bodies	
Researchers	
Service providers/integrators	
Device suppliers	
Network operators	

Out of scope	

Figure 2-1: Stakeholders Aide Memoire

2.2.2 Workflows

In healthcare, clinical pathways or care pathways are implemented to standardise the delivery of care. A care pathway is a series of processes that reduces variability and improves clinical outcomes for the patient, with three typical pathways being (Harikrishnan, 2017):

1. Data pathway- activity and habits are monitored, data aggregated and analysed.
2. Care pathway- care and clinical processes designed to optimise the outcome for the patient (user).
3. Payment pathway – care providers being paid for services to the patient (user).

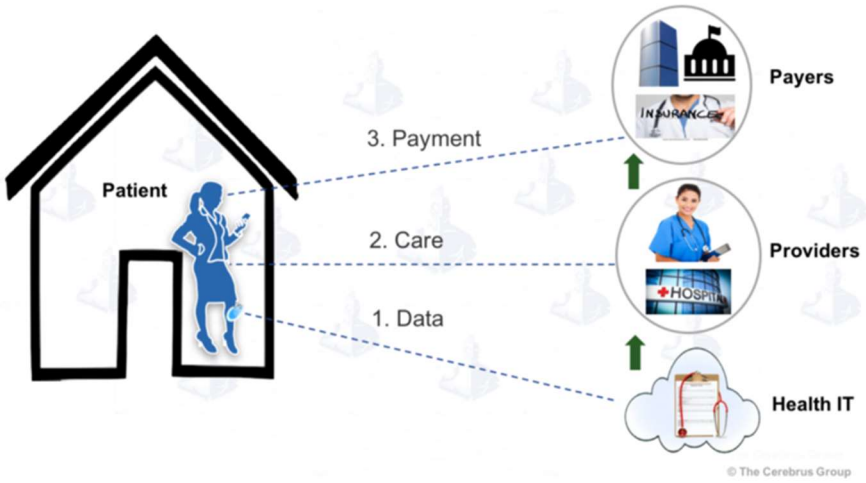

Figure 2-2: Three healthcare pathways (from Harikrishnan https://www.iotforall.com/health-monitoring-using-iot)

An IoT-based system may need to address all three pathways in order to be adopted; this is a major distinction from IoT devices aimed at consumer personal-fitness/health monitoring.

Remote health monitoring using patient-gathered health data currently remains limited due to the clinicians' concern with the reliability and accuracy of the gathered data to support decision-making (Alaboud et al., 2022). The paper identified four main themes:

1. Data generation and collection – understanding which measurements are important may require medical judgement. Standardisation collection and management of PGHD would be useful to inform medical decisions effectively.
2. Data integration and accessibility – challenge of integrating PGHD into the EHR, subject to privacy and data storage concerns, without creating additional workload to clinicians or needing to open a separate application. Ideally, the system would trigger an alert for the clinician for any data outside of trend, to avoid the clinician having to review data daily from hundreds of patients.
3. Data presentation – a data summary or dashboard.
4. Data interpretation and utilisation.

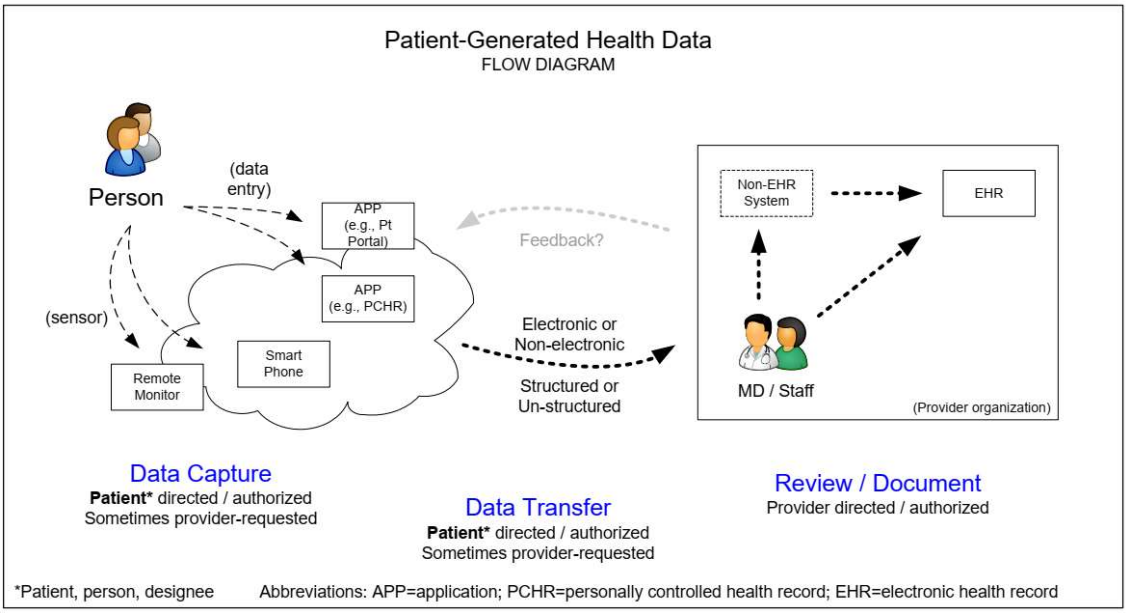

Figure 2-3 Patient-Generated Health Data (Shapiro, 2012)

The data may be structured or unstructured, numeric, text, waveform etc. (Shapiro et al., 2012), although it will be transmitted electronically. The data needs to be transferred and then reviewed and actioned. In some cases, the data may be discarded or incorporated into a medical health record, depending upon its provenance, quality and relevance.

2.2.3 Appropriateness of Smartphone AAL solutions

Patients should receive care whenever they need it and in many forms, not just face-to-face visits (Kohn et al.) and that such visits can be inefficient, particularly for patients such as the elderly, those with serious health complications, needing continuous monitoring or who are unable to move (Bhatti et al., Kohn et al., 2001). It is claimed to be economically and socially advantageous to reduce the burden of disease treatment by enhancing prevention and early detection while allowing people to stay at home for as long as possible (Ko et al., 2010). Assistance is not limited to physiological measurements, but can include identifying wandering behaviour, reminders for taking medication and prompts for routine tasks to help people with dementia carry out their daily activities (Rashidi and Mihailidis, 2013).

Remote monitoring systems also have value in other situations, such as in emergencies and after disasters (e.g.: earthquakes), to coordinate care, physiological monitoring and rescue for victims and to monitor the emergency responders themselves (Centelles et al., 2019), and enabling large-scale field studies of behaviours and chronic diseases (Ko et al., 2010).

Smartphone devices have enabled many health monitoring services (mHealth) as they have become more pervasive, user-accepted and powerful than ever with different types of low-cost sensors (e.g., accelerometer, gyroscope, camera, magnetometer, pedometer, goniometer, actometer, biometric and pressure). Several Active Assisted Living (AAL) projects use a smartphone as the core computing and communications unit (for examples see http://www.aal-europe.eu/), such as managing Parkinson's disease (Pasluosta et al., 2015). However, the usefulness of smartphones is uncertain due to not being cost-effective for some patients, whilst other individuals may be visually impaired, unable to use their hands effectively, or even unable to use the technology at all (Ghamari et al., 2016).

Smartphones also typically rely on licenced bandwidth for communication, and unless mandated, it is generally the mobile network operators who determine whether it is cost effective to provide coverage in remote areas. This can leave some isolated communities without mobile data coverage; this is discussed further in section 3.3.5.

However, some IoT networks operate in the unlicenced radio bands, and the basic radio base-station infrastructure could be installed for relatively modest amounts (a base-station can cost less than £1,000 to purchase). Some IoT networks have already been installed for other purposes and may have the potential to be used for RHM. For example, the Cook Islands have installed an IoT network as part of the Smart Island project for energy metering, road traffic management, street lighting operation and environmental controls mainly due to a 5G network not being practical due to the tropical and mountainous terrain and the required data rates being low (Hayes, 2022). Vegetation and large trees can impede radio wave propagation more than large buildings because a tree does not reflect radio waves, so although 83% of the islanders use mobile services and 38% have a fixed landline, there will still be areas where people could be without coverage. A similar scenario may arise in sparsely populated rural areas in other countries, or a temporary network may be installed to monitor workers such as logging or pipeline installation and maintenance.

2.3 User (Service) Requirements

This section draws upon the following conference paper, included in Appendix B:

> Poyner, I. K. and Sherratt, R. S. (2019) *Healthcare in rural communities and developing countries – IoT solutions to improve access.* In: Living in the Internet of Things (PETRAS 2019), 1-2 May 2019, London, UK. DOI: 10.1049/cp.2019.0104 (Poyner and Sherratt, 2019)

Extracts from the paper are highlighted in grey.

"Some technical developers are contentedly tinkering with Arduino or Raspberry Pi modules, clip cables, IDLE Python, plug-ins and demonstrating cool applications. But IoT system engineering requires also consumer focused and formal analyses of the application scope, business integration, synergies, scale economies, suitable resources, roadmaps, interoperability, standardization and deep investigation of the market; in order to achieve the high potential of IoT ecosystems." (Fernandez and Pallis, 2014)

"Health systems in LMICs usually do not provide appropriate support for the development and sustainability of m-health interventions. Instead, their development is usually driven by NGOs and private enterprises. Many services are not built for large-scale implementation, but rather for small pilot studies."(Quaglio et al., 2016)

The two quotes above highlight that to gain an enduring and effective uptake of an RHM system by the target population, there are many challenges beyond the technological. Instead, during the initial concept and design of a system/service the potential concerns of each stakeholder need to be identified and assessed (user requirements) and the required levels of quality of service (QoS) need to be understood. However, QoS have only been mentioned rarely in industry and research (Fabbricatore et al., 2011). Having assigned weightings to the various potential concerns, technical trade-offs can then be performed with an awareness of how they may impact across the stakeholders and value-chain (to derive the system requirements). The factors described in this section are necessarily generalised, and the influence that each factor will exert will vary dependent upon the target use cases and technologies employed.

The concerns identified from the literature covers a broad range of viewpoints, including physicians in Germany and the US, a review of RHM projects launched in LMICs, business-schools and IoT-centred assessments. Some concerns recur as a theme across many stakeholders, but it is important that each stakeholder may have a different perception or place different emphasis upon any given attribute. For example, for the theme of confidentiality, a principal may wish to protect the privacy of their health conditions, a clinician may need to control access to data and a service provider may be more focussed on data security and compliance with regulations.

In this section, the stakeholder requirements are assessed following a modified PESTLE (political, economic, sociocultural, technology, legal, environmental) analysis. The priority of concerns will vary depending upon the criticality of how any RHM system is being used and what level of reliability/quality of service is expected; for example, a device monitoring a life-threatening condition may need to provide very high levels of accuracy and coverage, whereas for a service monitoring pattern of life some loss of data may be tolerable.

The PESTLE analysis is modified to LESPET, as follows:

- Legal – includes certification of device type, safety and protection of sensitive data.
- Economic – affordability and the business case concerns of operating a service.
- Sociocultural – usability and acceptability.
- Political – the context of the wider healthcare system.
- Environmental – energy and waste management.
- Technological – influence of industry standards (note that technology requirements are covered under system requirements).

Table 2-2: LESPET 'PESTLE' Analysis

Stakeholder concern	Description/examples

Legal	• Duty of care demands (mandatory and best practice)
	  ○ Reasonably practicable risk mitigation
	• Certification of device type
	  ○ Medically certified (Part II – IV) versus consumer health-monitoring
	• Safety
	  ○ No undue toxicity
	  ○ Radiation levels within safe levels
	  ○ Accuracy
	  ○ Reliability
	  ○ Physical robustness – physical construction
	  ○ Robustness to noisy or missing data
	  ○ Alarms on detecting an error or outage (at both the device and service level)
	  ○ Data integrity (can detect or recover from corruption or malicious interference)
	  ○ Non-repudiation (both directions)
	  ○ Timeliness/Freshness
	  ○ Assurance of any on-board data processing

	• Protection of sensitive/personal health information
	  ○ Privacy
	  ○ Confidentiality (and encryption)
	  ○ Integrity
	  ○ Availability (resistance to denial-of-service attack)
	  ○ Assurance
	  ○ Authentication and Authorisation
	  ○ Consent management
	  ○ Identity management – users
	  ○ Identity management – devices, including onboarding

Economic	• Affordability
	  ○ Review and decision making – staff resources
	  ○ Training of wearers, carers/supervisors, support staff
	  ○ Integration of data into existing workflows
	  ○ Device purchase and ongoing support contracts
	  ○ Applications development and updates
	  ○ System administration
	  ○ Updates (development and deployment)
	  ○ Maintenance support
	  ○ Network costs
	  ○ Cloud processing and storage
	  ○ Disposal
	• Customisation and Flexibility
	  ○ User-specific tuning of monitoring
	• Scalability & Interoperability
	  ○ Integration with/modification of established workflows
	  ○ Elastic capacity
	  ○ Software and network scalable
	  ○ Open Standards
	  ○ Consistent semantics
	  ○ Density of devices

Sociocultural	• Acceptability

	○ Consent management of the use of PGHD • Usability, Mobility and Literacy ○ Convenient to wear near-permanently ○ Small form factor - size, weight and power (SWaP) ○ User-Centred Design (UCD)/ Disappearing User Interface (DUI)/ Silent Operations for Principal ○ UCD for Carers/Supporters ○ Support for local language/dialects if not DUI ○ Robustness ○ Geolocation of Principal ○ Low maintenance burden • Network Coverage ○ Operate away from home/base infrastructure and resources ○ Coverage appropriate to Principal's monitoring needs, environment and funding.
Political	• Integration with wider health system/medical records ○ Policy and Regulatory frameworks – including where liabilities reside ○ Data standards ○ Data quality ○ Filtering of data to avoid cognitive overload ○ Data sharing with third parties and privacy
Environmental	• Energy demands – device, network and data centre • Disposal

2.3.1 Duty of care

Many organisations are required to provide a duty of care to groups of their stakeholders. Consider the following examples.

- For an emergency responder (or team) dispatched into a wildfire, earthquake or other natural disaster, the command post may decide to frequently measure the

person's temperature, heart rate, oxygen levels to detect the early onset of exhaustion or other problems.

- A patient discharged from hospital to live with complex conditions in their own home may need regular monitoring. This monitoring may traditionally be carried out by carers or nurses visiting the principal and recording their physiological levels once a day. However, this only provides a very limited snapshot of the person's health at one moment in time, may be resource intensive due to the travelling of the carer and could miss peaks and lows of their daily levels.
- Public social service providers who have a mandatory duty of care to safeguard vulnerable people living in the community. Periodic status messages from a user may be required, and if such a message has not been received then the social services may be required to escalate initiation of contact via alternative means with the user.
- A company with employees working alone in remote areas, or schools with loco parentis for outward bound expeditions may need to justify why RHMs were not used should an incident. The company or school may decide a policy to sponsor RHM devices to track the condition of remote principals and determine that line managers/supervisors and teachers, as caregivers, are responsible for monitoring remote workers or young people (as principals). The sponsor may argue that the use of RHM demonstrates sufficient risk mitigation, and the ALARP principle ('as low as reasonably practicable') may allow the RHM coverage to have some blind-spots (e.g.: terrain masking), providing the system warns when a status update had not been received for longer than a given time period.

An organisation may have to explain the actions that it has taken are best practice to reduce risks to ALARP. The level of mitigation deemed adequate will depend upon the situation and the criticality of the conditions being monitored on the principal. Laws are rarely framed to mandate the use of a specific solution or technology as this could constrain innovation and give an unfair advantage to incumbent suppliers. In none of the examples above is there likely to be a mandated requirement to use an RHM solution. However, as the cost of such services decreases and their capability improves, in the medium-term it may become more difficult to justify why an RHM system has not been used for risk mitigation or improvement in care and monitoring. The more critical the

condition of the principal, or potential risk, then the more likely that an RHM solution would be expected to be used. However, in less critical circumstances, an RHM solution may be suitable for supporting carers who simply need to check that the user's normal 'pattern of life' is being followed in terms of eating, drinking, ablutions, sleeping, etc.

2.3.2 Certification of device type

As discussed in Section 2.1, there is a legal distinction between devices classed as medical grade and those that are consumer-grade health aids. There would be a higher expectation of a general medical device (categorised under Part II of the UK MDR 2002) to be reliable and accurate than a consumer device. Note that the criticality of active implantable medical devices (Part III) and in vitro diagnostic devices (Part IV) would be even greater.

Therefore, stakeholders need to determine whether their use case requires a certified medical device. Clinicians can be reticent to use RHMs where a suitable certified system is not available or cost-effective and they are concerned about their liability if they instead used a consumer grade system with unquantified reliability and accuracy. Drift and ambiguity in the sensors needs to be accounted for; mitigation may be 'by built-in redundancy with a large number of micro-fabricated sensors and reference electrodes' (Lo et al., 2016) or by periodically recalibrating sensors if this practical for remote wearers.

2.3.3 Safety and trustworthiness/reliability

There are several facets to safety of RHMs. Firstly, the device has to be built so as not to expose the user to any undue harm, such as through toxicity or radiation (see section 3.3.2).

Trustworthiness is termed by Ko et al (Ko et al., 2010) as the combination of data delivery and quality properties and they claim that medical sensing applications require high levels of trustworthiness. The level of trustworthiness in an RHM will include considerations such as accuracy, reliability, robustness (physical construction) and how the system reacts to missing measurements and errors (robustness to noisy data). '*Reliability directly influences the quality of patient monitoring. It can be life-saving in many situations and*

in a worst-case event; it can be disastrous when a life threatening incident has not been observed or detected' (Ghamari et al., 2016). Devices need to be reliable and not prone to system faults or generating false alerts (Ludwig et al., 2011); although this is an obvious statement, reliability can never be absolute, and the level of reliability and false positives required would need to be commensurate to the criticality of the system. If the wearer's device fails to measure and upload parameters for an extended period, then should the carer receive an alarm? If a measurement is well outside the expected range, how is it decided that this is a sensing problem rather than potentially a major health problem for the principal – is the carer or clinician involved or notified of these outliers in a timely manner?

The reliability of a service has to be considered at an overall system level, examining where problems may occur and how they could be mitigated. For example, in some areas, reliability may be impacted by unreliable networks and electricity outages (Aranda-Jan et al., 2014). Any service level agreement with the communications provider would need to specify key factors such as reliability, coverage, timeliness of data, managing user accounts and remote reconfiguration.

The most basic IoT communications protocols do not guarantee delivery of messages but simply rely on transmitting a message three times and accepting that not all messages will be successfully received. Such protocols may not be suitable for acute issues and assured health services, which may require confirmation of a message being received, generally requiring reliable message protocols using bi-directional connectivity.

The integrity of the PGHD needs to be assured, that is the data has not been corrupted or maliciously modified en-route from the sensor to the users who need to make decisions. Similar to this is assurance that the data originated from the device that is being claimed in message headers, rather than another device masquerading as the principal's device. In the return direction, where a carer, clinician or supervisor issues instruction or actuates the device, there must be non-repudiation so that they cannot later claim that they did not originate that instruction if there is a negative outcome.

The timeliness or freshness of recording and processing measurements needs to be considered in the context of the overall scenario. Several papers state the requirement for minimal latency time (Qadri et al., 2020), sometimes as short as milliseconds, which may be appropriate for time-critical applications where carers can intervene rapidly, but if the sensor is only used for occasional readings or the overall response time of a carer is hours for remote principals with tolerable health conditions, then a latency time of several seconds or minutes may be acceptable. For independent living/AAL applications, '*there is usually no immediate risk if data gets lost or corrupted. However, in the case of emergency response situations, substantial risks occur if the device is the only means to call for help'* (IEEE, 2015a). For diabetes management, a glucose meter may monitor blood glucose levels every 15 minutes and it may be important for high-priority readings to trigger (near) real-time alarms to carers supporting the principal (IEEE, 2015a).

An emerging issue is the level of trust in any algorithm on board the sensor that pre-processes the data. This is often done where there may be challenges in terms of battery life, network bandwidth or capacity to transfer all the PGHD to the core servers. The liability should the algorithm fail to correctly send the required PGHD or not raise a warning needs to be determined, often by legislators. This may require a scale of 'levels of autonomy' where at level 0, all PGHD is transmitted to the cloud, and at the other end the sensor may be totally relied upon to only pass warnings and alerts. This is analogous to the automated vehicles from level 0 to level 5.

As an example, the basic safety and essential performance requirements of some medical electrical equipment is specified in the ISO 80601 series of standards, which are also adopted as EN (European Standard) and BS (British Standard) standards. Pulse oximeter equipment is addressed in ISO 80601-2-61:2019 (ISO, 2019), with clause 201.12.1.101 specifying how accuracy is determined and that it should be less than or equal to 4.0% SpO_2 over the range 70% to 100% SpO_2 and clause 201.12.4.101 stating that '*there shall be an indication that SpO2 is not current when data update period is greater than 30s'* and that if equipped with an alarm at least a low priority alarm condition shall be provided when the data update period exceeds 30s.

2.3.4 Protection of Protected Health Information (PHI)

This section draws upon the following conference paper, included in Appendix B:

> Poyner, I. and Sherratt, S. (2018) *Privacy and security of consumer IoT devices for the pervasive monitoring of vulnerable people.* In: Living in the Internet of Things: Cybersecurity of the IoT, 28 - 29 March 2018, Savoy Place, London, doi: 10.1049/cp.2018.0043 (Poyner and Sherratt).

Extracts from the paper are highlighted in grey.

Healthcare professionals who have access to information must be confident that the patient's vital information is not tampered with or altered and did truly originate from the monitored individual. Furthermore, an overly secure system might disallow healthcare professionals from accessing vital health-related information in certain emergency events and thus jeopardize patient's life. Moreover, enriching the current systems with security and privacy mechanisms significantly increases the cost of energy for communication which results in more power drain from small batteries (Ghamari et al., 2016).

The potential impact of personal health information being accessed by inappropriate people or organisations may be much more harmful for a highly vulnerable individual than for people who routinely use and configure smart devices and may recognise when there is a problem. In terms of stakeholders, legislators and regulators are very active in protecting PHI, with bodies such as the Information Commissioner's Office (ICO), or equivalent, investigating and fining organisations for breaches of data security and privacy. The legal requirements to protect Personal Health Information (PHI) can be very demanding to reflect the potential impact of security vulnerabilities in a healthcare information system (Abouzakhar; et al., Blythe et al.), with potentially severe penalties for not effectively controlling data. Health systems need to demonstrate compliance to best practice, security and privacy of data guidelines (IOT Security Foundation, 2016).

Failure to consider safeguarding privacy at the outset of a project could result in significant re-design effort, recall of deployed devices, or inadvertent breaches resulting in large fines (potentially £17M or 4% of global turnover).

Careful consideration needs to be made for the users of any service. It cannot be assumed that end users are 'digital natives' who are familiar with technology and understand privacy options and how to configure settings. Instead, some users may be living with

cognitive or dexterity challenges, or for whom smartphones and other technologies are alien. For example, a complication with vulnerable users is that they are less likely to implement the most important protective behaviours (Blythe et al., 2017) such as updating devices and changing default passwords and may be more susceptible to falling victim to social engineering and counterfeiting.

Legal safeguarding of GDPR (EU General Data Protection Regulations) and HIPAA (US Health Insurance Portability and Accountability Act) mandate the protection of PHI (protected health information), including confidentiality, integrity and assurance of data whilst in transit and also authentication of the user and service provider. Claiming to offer a technology to track wellness and health does not exempt a company or project from needing to comply with the legislation to safeguard personal information applicable in any region in which they are operating. In addition to legal requirements, privacy concerns are often a barrier for users to adopt a new service and any loss of PHI would result in reputational damage for that company. Therefore, privacy and security must be a fundamental design consideration from the outset, often requiring some form of encryption.

Patients' concerns over the security and privacy of their data are also a barrier to the adoption of new technology (Dhukaram et al., 2011).

2.3.4.1 EU and UK General Data Protection Regulations

Within the EU, the GDPR (General Data Protection Regulation) came into force on 25 May 2018. In the UK, the Data Protection Bill implements most of GDPR with additional provisions; however, the Data Protection and Digital Information Bill was intended to reduce some of burdens on organisations, but its progress was paused in September 2022 to allow for further consideration.

Privacy by design has always been an implicit requirement of data protection, but the GDPR core principle of 'data protection by design and default', places a general obligation to implement technical and organisational measures to integrate data protection (Information Commissioner's Office, 2017).

The UK ICO advocates the seven 'foundational principles of privacy by design' developed by the Information & Privacy Commissioner of Ontario (Information & Privacy Commissioner of Ontario, 2013) which lists seven principles, four of which are very relevant to IoT:

- Privacy as the default setting,
- Privacy embedded into design,
- End-to-end security,
- Respect for user privacy (user-centric).

These have been further described against a number of user health data scenarios (Mihailidis et al., 2010) and design guidelines (Information and Privacy Commissioner Ontario et al., 2014). It is further worth noting that children are also afforded additional protection.

2.3.4.2 US HIPAA

In the United States, the Health Insurance Portability and Accountability Act (HIPAA) 1996, mandated the adoption of Federal privacy protections for individually identifiable health information. This has since been enhanced with a Privacy Rule (2003) and Security Rule (2005) setting national standards for compliance. It is applicable to any company that deals with Protected Health Information (PHI), including business associates and subcontractors who support treatment, payment or operations. PHI must be carefully controlled, and data has to be 'de-identified' by removing 18 specified identifiers or by using a statistical expert to determine that data cannot be associated with individuals.

2.3.4.3 Security and Privacy Requirements

The specific security and privacy requirements for an RHM service will depend upon the regulatory environment of the countries in which the service operates and processes PHI. However, the following general requirements are common to nearly all jurisdictions.

Security of data is often expressed in the three characteristics of confidentiality (privacy), integrity and availability. In a healthcare application, the following additional requirements may also be necessary (ITU, 2017, Continua, 2016).

i. Confidentiality – data is accessible only to those who have the right to know (ISO 17799). A user does not disclose information to unauthorised entities allowing the deduction of the state of the user (Darwish et al., 2017).
ii. Integrity - assurance that information has not been tampered with or modified in any way to undermine its authenticity. This involves two areas (Darwish et al., 2017):
    - Device Integrity - Information must be correctly collected and transferred by medical devices and sensors.
    - Data Integrity - Non-existence of information flows that may have been subject to modification by entities at different levels of integrity than the originating principal (e.g., integrity of data-in-flight).
iii. Availability – having timely access to information or the means to process the information must be available when requested/required. Some applications also require the ability to withstand attacks aimed at denying availability (e.g.: denial-of-service flooding the communications channel or rapidly depleting the battery).
iv. Identity management - management of user identities across the end-to-end architecture, hence associating health information with the right individuals.
v. Nonrepudiation of origin – is provided through the use of digital signatures and guarantees that the sender of information cannot later deny (or repudiate) having sent the information.
vi. Consent management - enables patients to provide and manage their consent preferences, which serves as a basis for governing access to and usage of their individual identifiable health information.

Additional security requirements may also be required (Malina et al., 2016, Zhang and Liu)

vii. Freshness – the data is still relevant to analysis and treatment.
viii. Audit – recording user activities of the healthcare system in chronological order, such as maintaining a log of every access to and modification of data. Enables prior states of the information to be faithfully reconstructed.
ix. Archiving - moving healthcare information to off-line storage in a way that ensures the possibility of restoring them to on-line storage whenever it is needed without the loss of information.

Privacy is defined in several different ways, such as (ITU, 2017, Malina et al., 2016):

> "*An aspect of system security (preventing undesired system use) that deals with providing access to the parties to which the information belongs and to parties that have explicitly been allowed access to certain information (also known as confidentiality)".*
> *"Privacy should protect a user's personally identifiable information and keep a certain degree of anonymity, unlinkability and data secrecy*"

Zhang and Liu argue that three principles are necessary in a cross-institutional patient records system to ensure privacy of patients and the content authenticity and source verifiability of electronic medical records:

- All electronic medical records should be guarded through ownership-controlled encryption, enabling secure storage, transmission, and access.
- The creation and maintenance of records should preserve not only content authenticity but also data integrity and customizable patient privacy throughout the record integration process.
- Access and sharing of records should provide end-to-end source verification through signatures and certification process against blind subpoena and unauthorized change in healthcare critical data content and user agreements.

2.3.4.4 Consent Management

A recurring principle in the literature is that a patient should be able to control (provide consent) as to what data is divulged to different care providers / health applications. Vulnerable patients may also need further safeguards, such as a guardian proxy. The individual, or their legal guardian, must be able to consent what information can be accessed by carers, supervisors and family members, and as importantly, withdraw or modify consent as circumstances change. There should be granularity as to the information shared with carers and supervisors as appropriate about their pattern-of-life and status, such as location, glucose level, blood pressure, temperature etc. However, carers should only have access to the information they require for their particular role to safeguard the individual and should not have access to continuous surveillance or privileged information that GPs or guardians may have.

A pre-requisite for consent is being able to identify and authenticate users and devices, to prevent an impersonator (person or device) gaining access to data. Fundamentally, this is a privacy issue to ensure the correct sharing of information among a group where membership may vary over time (Darwish et al., 2017). For an individual user and a trusted family supporter this may not be a major issue, but for a larger service provider with many users, carers and devices, onboarding and maintaining the correct configuration of access could become a significant administrative challenge.

2.3.4.5 Healthcare Frameworks

IoT system developers need to consider existing frameworks for connected health devices, especially where they wish to provide a service to an established healthcare provider who expects data to be secured in accordance with the frameworks with which they already use to protect other PHI data. A large organisation may place greater emphasis on the potential liabilities and reputational impact of a data breach than the opportunities offered by an innovative IoT solution.

Several health security models have resulted from hospital and community health care environments. Some of these have now been incorporated into international and open standards which help to create a stable market for devices compared to the wide diversity of embryonic technologies currently promoted across the IoT community. Below we discuss two established frameworks, HL7 and Continua.

2.3.4.5.1 HL7

HL7 (Health Level Seven International) provides international standards for the transfer of clinical and administrative data between software applications used by healthcare providers. Its name is derived from the application level of the ISO OSI (Open Systems Interconnect) network model, and it aims to provide semantic interoperability between systems. The NHS Direct Interoperability Tool Kit (ITK) is based upon HL7.

The HL7 privacy and security classification system (HL7 International et al., 2013) includes fields for confidentiality, sensitivity, integrity, compartment and handling caveat (purpose of use, obligations and refrain policies). The privacy rule is applied to composite

information extracted from a health record; for example, if a treatment is prescribed for HIV, then this will apply a restricted access rule to the record and require patient consent for onward release. A user accessing the system without a need to see the restricted information (for example a nurse entering temperature and blood pressure readings) would see only the incomplete information with details of HIV treatment either masked, encrypted or redacted (removed).

2.3.4.5.2 Continua

The Continua Design Guidelines (CDG) (Continua, 2017) are published and promoted by the Personal Connected Health Alliance (PCHAlliance). It is a framework based upon open standards to create a secure and interoperable health data exchange in personal connected health. The CDG are recognised by the International Telecommunication Union (ITU), which is the United Nations agency in the field of ICT, in ITU-T guidelines H.810 (ITU, 2017).

The CDG builds upon the ISO/IEEE 11073 series (IEEE, 2015b) for personal health device communication, HL7 standards and the technical specifications for USB, Bluetooth, Bluetooth LE, NFC and ZigBee.

In addition to the normal security C,I,A requirements of confidentiality (privacy), integrity and availability, the CDG/H.810 add requirements for identity management (management of devices across the end-to-end architecture), non-repudiation of origin (using digital signatures) and consent management (to individually identifiable health information).

2.3.4.6 Security Challenges for IoT Systems

Small, constrained IoT devices are not readily amenable to the security methods used in IT. For example, encryption is a common method of protecting data, but this is more challenging on very constrained IoT devices. This is being addressed by industry with the introduction of microcontrollers and other devices that include security and encryption capabilities, that reduce the processing and energy consumption overhead (for examples, see psacertified.org). Other security challenges to RHM IoT devices include (Darwish et al., 2017):

- Validation of the measurements due to device diversity, misuse, and mistakes.
- Valuableness of the information related to patient's health for data understanding.
- Heterogeneity, agreement and synchronisation among the sensors by different producers.
- Various aims of access to the data entries that may be used by doctors, carers, researchers and others.
- Different protocols and technologies of communication.
- Dynamic network topology, and multiplex data transfer between providers, home, central back-ends.
- Computational, memory, energy, mobility limitations.
- Special threats for the privacy because of big data collections.
- Dynamic security updates and tamper-resistant packages required for IoT devices.
- Complex human interactions, e.g., it is vital to adopt usable passwords due to patient's health.

Further challenges arise, such as:

- User interface - IoT devices often do not have a direct HMI (human-machine interface), such as a screen or keyboard. The use of DUI (Disappearing User Interface) helps to keep complexity and component costs low and to reduce size and power consumption. However, configuring devices and 'enrolling' them with credentials to a particular service can be more difficult. One challenge is how can users create and enter a unique password if there is no simple HMI, especially where a user may have visual, dexterity or cognitive challenges.
- Access - RHM monitoring devices often have to be kept on or around the user to be effective. There is greater potential for the devices to come into contact with more people compared to smart-home devices. The contact may be physical or within proximity of the RF signal (which can be up to 10's of metres). Therefore, there is greater opportunity for a malevolent actor to intercept data or to interfere with the security of the device.
- Non-traditional communications: Internet protocols such a as TCP/IP, enable the user of higher-level protocols including HTTPS and DTLS that provide security mechanisms such as handshaking. However, IoT communications generally do not use TCP/IP [and instead use MQTT or COAP]. In addition, IoT devices

generally act as servers, 'pushing' information to clients that request, or 'pull' data. Acting as a server means that the IoT device should respond to unsolicited requests from other devices, which is the opposite to how a PC or smartphone would typically request data from web servers, enabling it to block any unsolicited devices which it has not initiated communication with.

2.3.4.7 Guidelines on Potential Threats and Mitigations/Countermeasures

Following the above discussion on the constraints of IoT systems, fault-tolerant architectures need to accept that IoT devices, software and communications do not have the same reliability as that found within a clinical or hospital environment. Intermittent data loss has to be assumed and designed into the overall system.

Implementing security in a very constrained system requires very careful selection of the encryption algorithm to ensure that it remains robust even with intermittent connections and without needing end-user interaction (Yang et al., 2017), (Suo et al.), (ETSI, 2014). It may be that currently '*no security approach provides the perfect solution [for constrained devices*]' (IETF et al., 2017).

Identifying potential threats and introducing mitigations, or countermeasures, that can be deployed within the constraints of IoT systems should reduce their vulnerability. Two examples are given below (from (Mosenia and Jha, 2017) (Abdul-Ghani and Konstantas, 2019)):
(Abbreviations: C – Confidentiality; I – Integrity; A – Availability; AC – Accountability; AU – Auditability; TW – Trustworthiness; NR – Non-repudiation; P – Privacy; M – Manufacturer; D – Developer; C – Consumer; P - Provider).

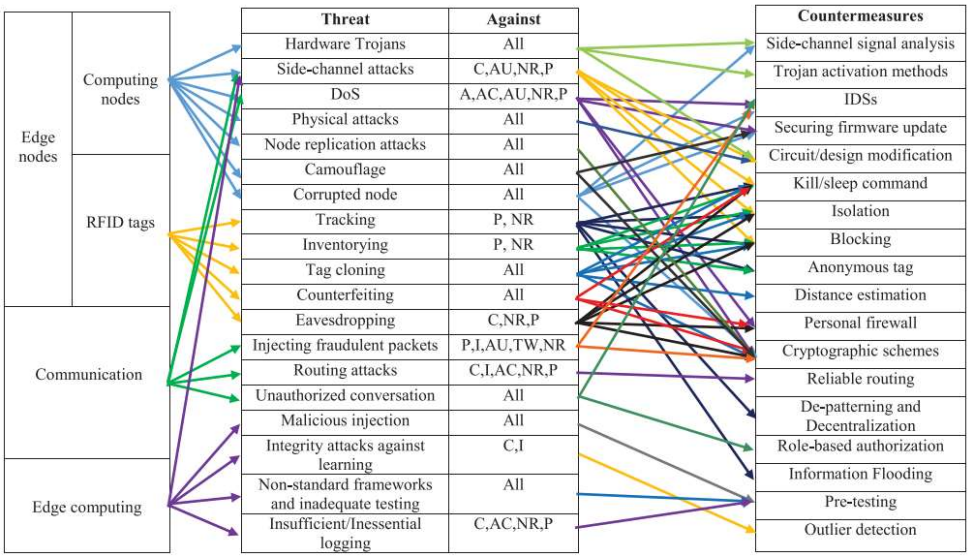

Figure 2-4: Summary of attacks and countermeasures (from Fig 3 of (Mosenia and Jha, 2017))

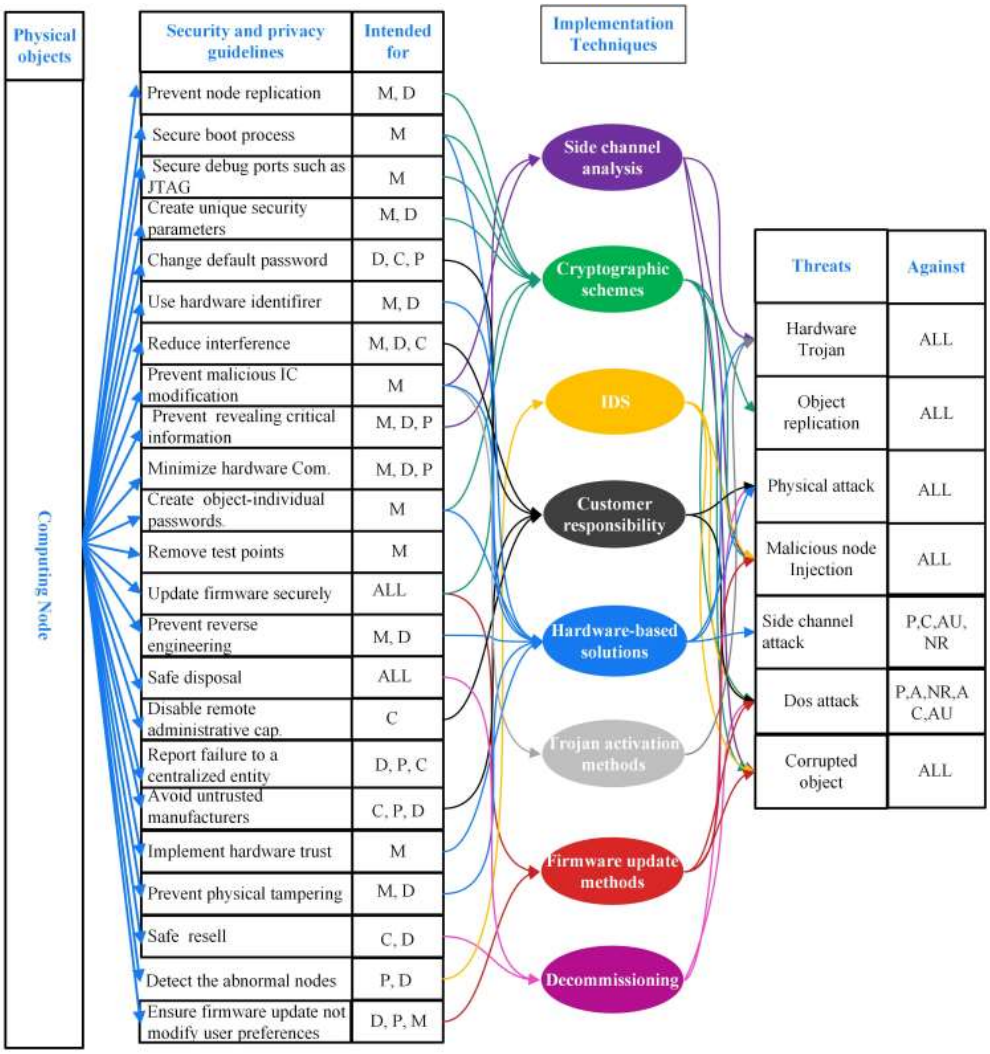

Figure 2-5: An overview of guidelines, stakeholders, attacks and countermeasures for computing nodes (from Figure 2 of (Abdul-Ghani and Konstantas, 2019))

Other initiatives in this area include:

- NCSC (National Cyber Security Centre) Secure by Default (NCSC, 2017) – see below.
- ETSI (European Telecommunications Standards Institute) Technical Specification for 'Cyber Security for Consumer Internet of Things' (ETSI, 2019)
- NIST (US National Institute of Science and Technology) lightweight security project for IoT (NIST, 2017),
- IETF working group addressing authentication in constrained environments (IETF, 2018a)

2.3.4.7.1 Secure by Default

'Secure by Default' (SbD) principles, as espoused by the UK National Cyber Security Centre (NCSC) (NCSC, 2017), are highly applicable in systems for vulnerable people. SbD addresses the entire system including hardware, firmware, software, services, applications and configuration. Also, security must be considered throughout the whole lifecycle of an end-to-end service, especially when devices are upgraded as this may be an opportunity to deploy more capable countermeasures. The SbD principles include:

- security should be built into products from the beginning, it can't be added in later,
- security should be added to treat the root cause of a problem, not its symptoms,
- security is never a goal in and of itself, it is a process – and it must continue throughout the lifetime of the product,
- security should never compromise usability – products need to be secure enough, then maximise usability,
- security should not require extensive configuration to work, and should just work reliably where implemented,
- security should constantly evolve to meet and defeat the latest threats – new security features should take longer to defeat than they take to build,
- security through obscurity should be avoided,
- security should not require specific technical understanding or non-obvious behaviour from the user.

The last point is particularly pertinent to vulnerable patients who should not be relying on untrained family or carers to configure their systems; this is a matter of balancing the advice that default configuration settings are set to the most secure possible, against the likelihood that this may frustrate users, especially where users have special needs (e.g.: need to configure large font, audio-visual accessories). As two of the principles state, security should not compromise usability and should not require non-obvious behaviour from the user, especially when users may not be familiar with technology or living with cognitive challenges.

2.3.5 Affordability including Maintainability and Battery Life

Affordability is a leading barrier to according to the ITU (ITU, 2018b). They provide that example that only 32% of Rwandan households would find mobile access affordable, based upon the assumption that cost is less than 5% of total expenditure. While globally two in three people are now online according to the ITU definition of internet access (any use of the internet at any time within the past three months) (ITU, 2020a), billions lack the meaningful connectivity they need to make the most of the internet (A4AI, 2022), defined by the Alliance for Affordable Internet (A4AI) as 4G-like speeds, smartphone ownership, daily use, and unlimited access at a regular location, like home, work, or a place of study.

Table 2-3: Comparison of Internet Use and Meaningful Connectivity (from (A4AI, 2022))

	ITU Definition of Internet Use	A4AI Meaningful Connectivity
Speed	No minimum speed	4G-like speed
Device	Any device	Smartphone ownership
Data Allowance	No minimum	Unlimited broadband connection
Frequency	At least once in the past three months	Daily use

On average, only one in ten people in LMICs have meaningful connectivity, compared with just under half who have basic internet access, according to official figures. The A4AI assert that maybe fewer than one in 160 Rwandans have meaningful connectivity, and that there are large inequalities, particularly for women.

There is also a distinction to be made between the use of a smartphone with data connection to the internet, and a more basic phone that primarily uses voice and SMS messaging. Mobile phones used in low-income communities are predominantly basic phones, although smartphone penetration is rapidly increasing within the middle-class communities. Therefore, it is likely that SMS will continue to play a key role in RHM where cost-effectiveness is important (Khazbak et al., 2017).

Even in developed countries, smartphones and mobile contracts may be perceived as unaffordable by people on the lowest incomes or pensions. However, for care providers, such as social services, remote monitoring services may be cost-effective if they enable users to remain independent for longer and not require residential care.

Most medical sensors have traditionally been too costly and complex to be used outside of clinical environments (Ko et al., 2010), whereas the integration of devices from the consumer market can enable a solution to serve the mass market (Fabbricatore et al., 2011). However, device cost is only one aspect as affordability encompasses many staff, capital and operating costs across the whole lifecycle, including:

- Staff time to review the PGHD and decision making as to what needs to be integrated into medical records or actioned as a priority.
- Training of wearers, carers, clinicians and system support staff.
- Integration with the care-service/clinician workflows and systems.
- Device purchase and any ongoing support contract.
- App development and upgrades, for both the principal and the carer/clinician/supervisory teams.
- Firmware and software updates required to protect security and address obsolescence over a potential 10-year lifespan. Unless a device can be updated over-the-air or over-the-network, this may require the device to be returned to the maintainer.
- Maintenance support, which may include sending a technician to change a battery.
- Network/telecommunication service costs.
- Cloud storage.

- System Administration to securely provision and update heterogenous remote sensors with limited memory, processing, cryptographic and receiving capabilities.
- Disposal, including secure erasure of PHI/sensitive data from the device, cloud and care services systems as required (note that these may need to occur at different times if a care service is required to maintain records for a number of years after the end of providing services).

For a carer service employing staff (as opposed to family or volunteer carers), staff costs are likely to be significant. Manual monitoring and assessment of RHM data will take staff time that needs to be paid for. Clinicians also need to be rewarded for their time to review any PGHD received and integration with medical records where appropriate. This may lead to the development of automated assessment software, but the output may still need to be reviewed to minimise the risk of liabilities should a principal's anomalous PGHD not be recognised by the software.

Integration with, or modification of, existing workflows can often entail specialised software and business process engineering effort, which will have a cost, both financially and in management time. There may also be a need for legal support to ensure compliance with statutory and professional standards and that potential liabilities are not being introduced.

The initial device cost may be a small fraction of the whole life costs of the service. Consumer health trackers can often be very cheap because of the use of mass-produced items and open-source libraries. Certified medical devices will usually be more costly due to the need to recoup the engineering costs of gathering the evidence required to undertake certification and the costs of independent assessment.

The level of maintainability of a system, which may incur extra design and manufacturing costs, needs to be analysed carefully during the design stage (Akmandor and Jha, 2018).

Ideally, an RHM solution should avoid the user sensor and communications needing to be mains-powered or recharged frequently during service. This is because a user may not

be able to, or simply forget, to routinely check the battery level. This could occur due to several reasons including diminishing cognitive ability or being bed-bound. Intermittent power supply is another potential barrier in remote communities and LDCs.

Relying on carers or service technicians to visit the user to recharge or replace batteries could add an additional cost and complexity to the solution. Battery life can be a significant disadvantage when using smartphone devices for processing and communications in a mHealth system because such phones typically discharge within only a few days. In contrast, some IoT devices are designed to be fielded for 10 years without changing the battery, albeit when operating in very low power mode (in sleep mode for much of the time, scheduled to 'wake-up' to perform readings/processing and then sending the results to a remote server). A more realistic target may annual battery recharge/replacement for an IoT mHealth device that includes remote communication.

Alternative energy sources, such as body heat, movement, and solar charging, are becoming more common in fitness wearables, but may not be appropriate for health monitoring if the wearer is sedentary or the device needs to be in a specific body area (e.g.: the chest for cardiac monitoring).

Of the three aspects of health monitoring: sensing, wireless communication and data processing, the wireless communication is likely to be the most power consuming. The power available in the nodes is often restricted. The size of the battery used to store the needed energy is in most cases the largest contributor to the sensor device in terms of both dimensions and weight. Batteries are, as a consequence, kept small and their energy capacity is limited, nominally around 1000J, although this will be dependent upon many factors including construction and chemistry, discharge profile and temperature. Therefore, to reconcile the requirements for both a small battery and a long battery operating life, the wearable needs to be highly energy efficient. In some applications, a sensor/actuator node should operate while supporting a battery lifetime of months or even years without intervention. For example, a pacemaker or a glucose monitor would require a lifetime lasting more than 5 years. Especially for implanted devices, the lifetime is crucial. The need for replacement or recharging induces a cost and convenience penalty

which is undesirable not only for implanted devices, but also for larger ones (Latre et al., 2011).

Abstraction methods for reducing the quantity of data to be transmitted, whilst maintaining sufficient precision, are described by Ganz et al.(Ganz et al., 2015). The emerging field of machine learning (ML) on a constrained IoT device, which can also reduce the amount of data needing to be transmitted, is explored in Section 3.7. However, data abstraction may hide subtle changes in medical conditions or make the data less useful for subsequent research and analysis (Fafoutis et al., 2018).

In summary, the costs of introducing and the ongoing operation of RHM services can be much more than the initial cost of the devices and network access. Labour costs may be more significant, and the upfront costs of developing some automation to process PGHD may be repaid within the lifetime of the service. However, another potential barrier to introducing RHM services is the perceived risk of liabilities for professionals when a significant health issue is missed, either because the RHM fails to record it or the PGHD is not processed and assessed properly.

2.3.6 Customisation and flexibility

Health indications, tolerance to treatments, and physiological responses to rehabilitations are user dependent. To provide accurate user-specific health monitoring, the RHM needs to be cognizant of these variations. Customisable RHM systems should make it easy to personalize their parameters to a particular user and health condition. This will lead to more accurate responses (Akmandor and Jha, 2018).

2.3.7 Scalability and Interoperability

For any organisation, including voluntary, that is monitoring more than a few people, it is important that any new RHM can be integrated into existing workflows and that adding additional users has a low marginal effort and cost in order to provide efficiencies of scale. Often, new technologies are introduced with very conservative estimates of future demand and then demand exceeds capacity early in the lifecycle. This may result in requiring the replacement of supporting infrastructure and support software at additional

cost. Instead, an RHM system should plan for elastic capacity to be able to cope with rapid growth or reduction in users as required. Cloud services are often promoted as being a very flexible method of managing capacity, although it may actually be more expensive than on-premises systems. Any software and networking systems also have to be planned for changes in demand. For example, databases may need to be able to process an order of magnitude more data points than they were originally designed for, whereas networking systems need to be able to grow without congestion resulting in loss of PGHD or perceptible delays to the users.

A further technical consideration may be the density of devices within a network and whether interference could reduce the reliability of the systems. This may occur where the principal has multiple devices or in a community where many users are using sensors (e.g.: independent living community, retirement villages). Note also that non-RHM systems (e.g.: industrial monitoring, environmental sensors, smart street infrastructure) may also be using the same frequency bands which can lead to congestion and loss of messages (see section 3.4).

Interoperability of the PGHD with existing systems is also important to avoid information becoming siloed and not being readily accessible to carers and clinicians when making decisions regarding a principal. Interoperability is also an important business requirement in that it avoids vendor lock-in or losing support if that vendor is no longer operating in your region; this could apply to device type, communications provider or other proprietary components or software. Using the same open standards and data formats are extremely important (Qadri et al., 2020, Fabbricatore et al., 2011), but it is vital that the semantics and calibration of data points are consistent across systems so that processing and decision-making services can incorporate a variety of heterogenous devices. This can avoid having to maintain separate monitoring and control systems for each different family of RHMs, potentially dispersing critical information across several databases (Akmandor)(Akpakwu et al.). Subtle differences, such as where a body temperature is measured by an RHM resulting in an offset, could lead to incorrect assessments of the principal.

Several health security models have resulted from hospital and community health care environments, or from personal health and fitness devices operated around the home. Some of these have now been incorporated into international and open standards which help to create a stable market for devices compared to the wide diversity of embryonic technologies currently promoted across the IoT community. Examples include the HL7 and Continua frameworks discussed in section 2.3.4.

2.3.8 Acceptability

Principals may not wish to be 'tagged' with RHM devices, and it is important that any monitoring is fully consensual. Adoption of an RHM will be personal to each individual's perception of the benefits versus the potential drawbacks, such as invasion of privacy. Examples of making the RHM less intrusive and increasing its functionality include incorporating the RHM into a walking stick, pendant or clothing. The emotional aspects of trust are important to adoption process and principals may have fears of illegal access, data transfer without consent and data loss due technical malfunctions (Ziefle et al., 2011).

Remote workers may feel that they are being overly tracked and their performance being unduly monitored. This has been experienced in industries such as courier and delivery drivers. A further consideration may be whether an RHM could lead to an increase in health insurance premiums.

There may also be cultural challenges from carers, clinicians and supervisors (Michell, 2017 est), including the following quotes:

- *Scepticism of new devices 'they will never work'*
- *Often culture and people change aspects are not addressed in IT changes – e.g., EPR –was expected to reduce paper notes, but clinicians are still writing paper notes as they have not altered their traditional culture practices.*
- *The mismatch between costs managed in one NHS department and benefits seen by a separate NHS department reduces incentive to purchase from separate departments.*

- *The clinician attitude (to IoT) is not managed (needs to change). Part is fear of losing control and lack of understanding (of IoT).*
- *There is not enough information on who it is for (IoT) and how it's used.*

2.3.9 Usability, Mobility and Literacy/Language

A benefit of ambulatory medical sensors, whose small form factor allowed them to be worn or carried by a person, is that they can observe cardiac irregularities, neural events and other symptoms that may not manifest during a short visit to the doctor (Ko et al., 2010). However, to achieve such benefits, the sensors may need to be worn for all or much of the time to capture transient events. Therefore, they need to be convenient to use and wear, and the user needs to accept that the benefit to them outweighs any perceived drawbacks (see also section 2.3.10).

Usability issues (including language and configuration) is the leading barrier to mHealth adoption in LDCs, according to the United Nations (Johnson M. ITU Deputy Secretary-General, 2018). For example, the ITU reports that in Rwanda, 92% of the population are within 3G coverage and internet should be affordable to around 32% of households, yet only 9% of households had access to the internet. The Government of Rwanda recognises that low computer literacy is a major obstacle to becoming a "Smart Nation" (less than 9% of over 15-years old are computer literate) (ITU, 2018a).

Technologies such as smartphones, low-cost computers and broadband communications are enabling novel healthcare services that can help people, young and old, reduce the impact of chronic conditions, such as diabetes, cognitive challenges and dementia. In some cases, a user may consent to share information in a controlled manner with other people who can provide support, such as family members, carers and organisations who have legal safeguarding obligations to the user, such as social services.

Usability also impacts security, and an RHM service should be designed so that some security mechanisms accomplish their objectives even if they are not used properly by users (including carers) (Darwish et al., 2017).

However, smartphones and fitness apps may be wholly inappropriate and difficult for some users, especially if they are coping with cognitive and physical challenges. In LDCs or remote areas, users may be further inadvertently excluded from the transformative benefits due to a smartphone solution potentially lacking support for the user's native language or script, are unaffordable, not supported by the network or requires too frequent charging, in communities where power is not always reliably available. User-centred design is needed to ensure that devices are usable for the individual, rather than a large demographic group (Goldberg et al., 2011).

End-user device management and ergonomics can also deter users. As a minimum, the ISO 9241 series of standards for Ergonomics of Human-System Interaction should be considered. For the principal who may be living with several conditions, such as failing eyesight and hearing, loss of dexterity and cognitive skills, this may result in a commercial fitness device being unusable. Instead, a more appropriate device may have a 'disappearing user interface' (Hui and Sherratt, 2017) or operates in a silent manner i.e., without seeking frequent inputs from the user (Akmandor and Jha, 2018).

Having an RHM that is not dependent upon a smartphone may have clear benefits for monitoring elderly people, supporting users unfamiliar with smartphones or the available languages, and people involved in sports and similar activities (Sanchez-Iborra, 2021). If the RHM requires frequent user interaction, it becomes less appealing and may lead to either incorrect decisions or discontinued operation. Edge-side processing of data and inference can help to make automated decisions without user interaction (Akmandor and Jha, 2018). This is explored further in section 3.7.

Related to lower and cognitive challenges, robustness of all the devices in the RHM system is required to ensure continuing performance. Tolerance to being dropped, poked with objects and submerged in water and other fluids need to be considered, such as using an EN/IEC 60529 IPxx rating for degrees of protection by enclosures.

Usability should also be considered for the carers/supervisors. They may be unfamiliar with technology and reluctant to use it, preferring instead to use their familiar tools and

processes. Part of the business change of introducing a new system should be training of all the users of the system and PGHD.

Literacy may also be a barrier for both the principals and carers (Aranda-Jan et al., 2014), especially where the system does not support the users' native language or local dialects. Graphical representations of data can assist in understanding, but it is vital that users are trained on the significance of what it is being presented to avoid incorrect conclusions.

Mobile devices almost always rely on a battery for power, and they are also the largest component in terms of weight and volume compared to other electronic components (Ghamari et al., 2016). As small size and weight may be a necessary, or at least desirable, user requirements, this drives the user need to minimise battery size, which then leads to the technical requirement of the device needing to be energy efficient.

For wearers with cognitive conditions, such as dementia, carers may need to be warned when the principal has left their home, so requiring geo-fencing and possibly subsequent periodic reporting of location. Sensors used by remote workers are very likely to be mobile and require a location capability so that the wearer can be located quickly if an alarm is triggered.

Ideally, a health monitoring system would operate effectively without any ongoing intervention or maintenance by the user or carer after the initial setup. There should be no need for technicians to configure or maintain a device. A constraint to be reduced is the burden of needing to repower the device(s), such as by recharging or replacement of batteries.

2.3.10 Coverage

For many users, there are significant benefits to exercising and socialising outside of the home. To be effective, the sensors need to continue functioning whilst untethered from any home infrastructure. As a minimum, any PGHD may need to be stored on the device and then uploaded later once the sensor reconnects to the home infrastructure.

It is easy to assume that coverage is no longer a problem when using mobile phone networks. However, although 93% of the world population has access to a mobile broadband network, this varies considerably with 23% of Africa's population having no access (ITU, 2020b). For supervising remote workers or monitoring groups in wilderness areas, coverage is likely to be an even greater issue.

The need for telehealth, and some of the same availability challenges, also exist in remote or isolated communities in the western world. The FCC programme 'Promoting Telehealth in Rural America' (FCC, 2017) recognises that rural health care is more vital than ever and provides examples of benefits to diabetes patients and mental health care. Only in the last couple of years have mobile services of at least 5Mbps has become near-ubiquitous in the USA, but still over 16% of the rural population do not have mobile coverage of 10Mbps as of April 2020, according to the Federal Communications Commission (FCC)(FCC, 2020a). As of August 2022, over 500,000 square miles of the US have no cellular coverage, as do large stretches of ocean (T-Mobile, 2022).

Although satellite coverage is available to most parts of the Earth, this usually requires specialised handsets which consume relatively high power and are expensive. For example, the FCC rejected a Starlink bid for broadband subsidies under the Rural Digital Opportunity Fund stating that users would be required to buy a $600 dish and would need nearly $900M in universal service funds until 2032 (FCC, 2022). There are plans to provide coverage to regular mobile handset via satellites in late 2023 (T-Mobile, 2022), which will depend upon launching a new fleet of satellites (Starlink v2) and the subscription costs are yet to be announced. Another company aiming to achieve the same outcome is AST SpaceMobile (AST SpaceMobile, 2022). However, as with all major satellite projects, there is a risk of cancellation or delay to service.

In the UK, Ofcom report that as of summer 2020, 590,000, or 2%, of the total residential and commercial premises cannot receive a 'decent' broadband service (10Mb/s download) and that 9% of the UK geographic area is not covered by any 4G operator (including 5% of all roads) and 5% of the UK is a 'not-spot' for calls and texting (Ofcom, 2020).

In China, over half the population (as of 2007), referred to as the bottom of the pyramid, do not take part in the formal economy and rely on C/VDs (community and village doctors) (Jiehui and Zhang, 2007). Although the trend of urbanisation may have reduced this, it is likely that many in rural areas still use C/VDs.

2.3.11 Integration with wider health system and medical records

There is great potential for new technologies, such as RHM, to greatly improve the level of monitoring of principals' health. However, *data quality and cognition overload can lead to 'issues managing, processing and analysing the continuous flow of data as we scale up' and clinicians 'having to ask the right questions in an ocean of possible data'* (Michell, 2017 est).

Further, lack of policy and regulatory frameworks might prevent the scaling up of e-health, particularly in LMICs. This may have resulted in many services in LMICs not being built for large-scale implementation, but rather for small pilot studies (Quaglio et al., 2016). Examples of where regulation may need to be established to support the adoption of PGHD include:

- Clarification of liability for inaccurate or missing data.
- Controls on the sharing of population-level 'Big Data' with researchers and commercial companies to avoid individuals being de-anonymised.

Exploiting the insights provided by PGHD will need to be managed at the national health care system level. Similar to how an individual clinician or hospital would need to integrate data-sources into their workflows, at the national level new policies and regulations may need to be introduced to enable wide-spread adoption. The primary factors may include standardisation of data formats and semantics, funding for devices and services, and agreement as to what is cost-effective. In the UK, NICE (National Institute for Health and Care Excellence) Office for Digital Health is leading on development of digital health policy and the Innovative Devices Access Pathway (NICE, 2022).

2.3.12 Environmental Concerns

Environmental concerns include the energy and resources consumed during manufacture and through-life operation. Replacement of batteries is an obvious concern, but also the energy consumed by data centres in processing and storing the vast quantities of data that will be generated by large scale adoption of RHMs.

Planning for device disposal at the end of life should be considered at the initial design stage. Where possible, much of the device should be re-used but how any personal data residing on the device could be sanitised should be an early design decision.

Chapter 3 The Use of IoT Technologies for RHM

3.1 System (Technical) Requirements

Managing a system based upon constrained devices deployed into a highly variable environment with limited communications will always be challenging. The system may need to manage security, FCAPS (fault, configuration and remote reprogramming, accounting, performance and security) (ITU, 1992), localisation of mobile devices, prioritisation of messages (e.g.: routine, high priority, emergency), reliability, remote calibration, scalability and interoperability across homogeneous devices.

IoT systems can be considered as a series of layers, with data moving back and forth across each layer. However, there is no universally agreed, or ideal, architecture and often architectures are developed for a specific use case, and often the different layers are not completely independent and will impact or constrain the other layers. Two useful architectures are presented by the ITU (ITU, 2012b) and ISO/IEC (ISO/IEC, 2018)

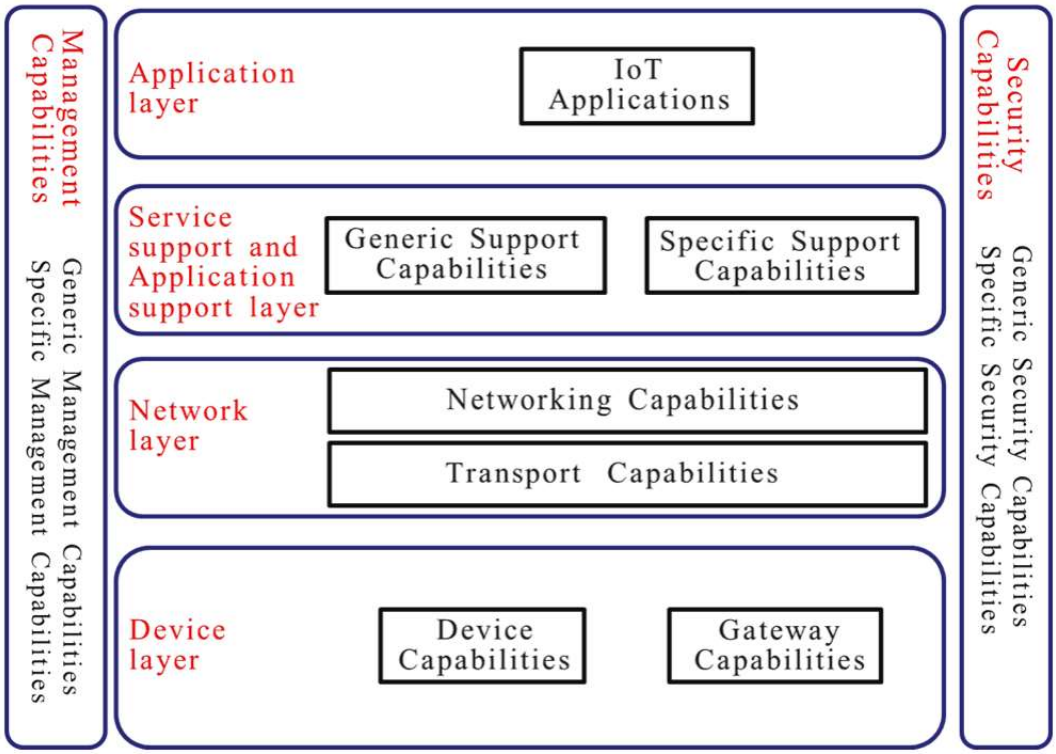

Figure 3-1: IoT Reference Model (from ITU Y.2060 figure 4)

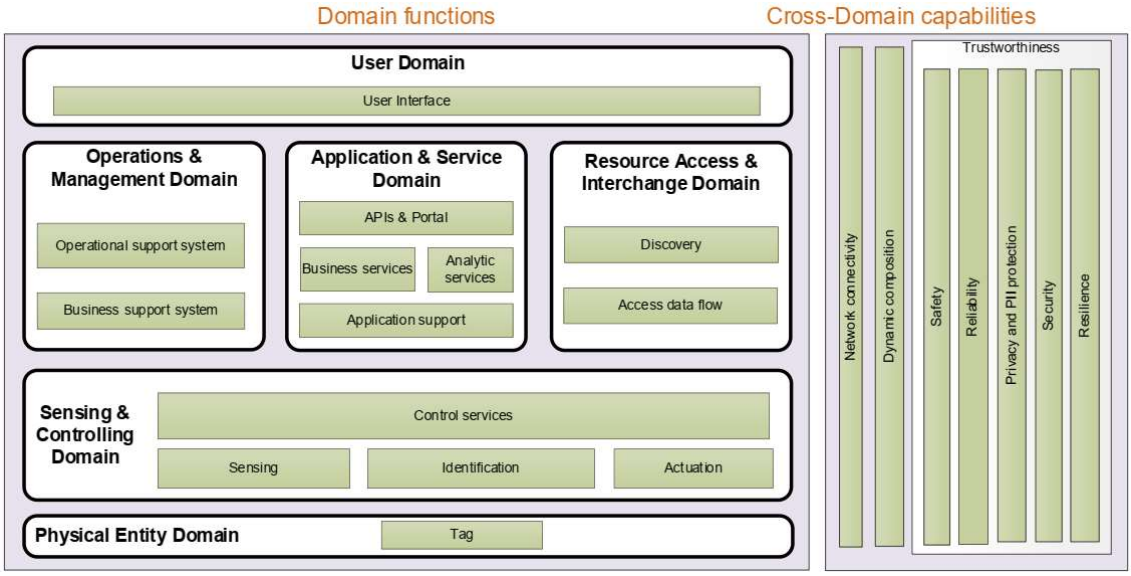

Figure 3-2: IoT Reference Architecture Functional View – decomposition of functional components (from ISO/IEC 30141 figure 15)

Variations on the terminology include the use of Perception Layer for the Device/Physical Entity and Business Layer for the User Domain. As can be seen in the two figures, there are differences between how some functions are considered – ISO/IEC consider the network connectivity to be cross-domain, whereas the ITU consider it to be a layer, and the Operations and Management is a function in the ISO/IEC model, but cross-cutting for the ITU. However, both models (and several other architectures) recognise that security/trustworthiness must be considered as a whole across the entire system to avoid any weak links being exploited.

For assessing the system requirements of an RHM service, the ISO/IEC model will be adopted in this paper, as the greater granularity of the cross-domain capabilities are important to discuss for an RHM service. Note that each component may be implemented using combinations of hardware, firmware, device drivers and software (ISO/IEC, 2018).

It is also necessary to recognise that some of the user requirements may be best considered at the overall system/service level as they may be satisfied by a combination of technical and business attributes from two or more components.

The following subsection includes extracts from (Poyner and Sherratt, 2019).

This section explores four segments of an IoT system (ITU, 2012a), namely organisational/carer services, the user device (sensor), wide area communications, and overall system requirements, including security.

Within a smart-home or healthcare environment there is generally ample opportunity to generate a richness of data from multiple sensors, either worn by or implanted within the user, and ambient sensors, such as sensors in chairs, bed and the bathroom. This can improve overall accuracy and reliability by correlating events across sensors. However, for this research, the use cases are more complicated by considering users who may be outside of their homes, and there is no gateway device or decision measuring unit (Ghamari et al., 2016) to provide processing power to perform data collection and protocol conversion from multiple heterogenous sensors, aggregation and filtering, local analysis of the data and triggering of alarms when required, implementation of security protocols and encryption, establishment of high-bandwidth communications, etc. Instead, these functions must be performed either within the low-power user device or by a remote system.

It is also important to consider the context and environment of how the system will be used. Accuracy, reliability and robustness can often be improved by designing a sensor specifically for RHM purposes, such as by including redundant or alternative measuring circuits, and calibrating the package around the human physiology.

3.2 Security & Protection of Personal Health Information (PHI)

This section investigates the security challenges faced by RHM service providers, the security frameworks that can assist in securing RHM services and where they may need to be tailored for the specific RHM use cases. It also looks at the principles of ‘Secure by Design’ the Data Protection by Design and Default’ and how these may need to be balanced against the health benefits for the principal users. Industry’s efforts to support cyber security is investigated, as is the emerging use of machine learning on devices to minimise the amount and granularity of data transmitted from the RHM device.

3.2.1 Security Challenges

As discussed in Section 2.3.4, users will not fully adopt IoT if there is no guarantee that it will protect their privacy (Hammi et al., 2018). There are many security challenges in protecting PHI due to the high privacy and confidentiality requirements in a healthcare system, service or application (ENISA, 2015), as shown in Figure 3-3.

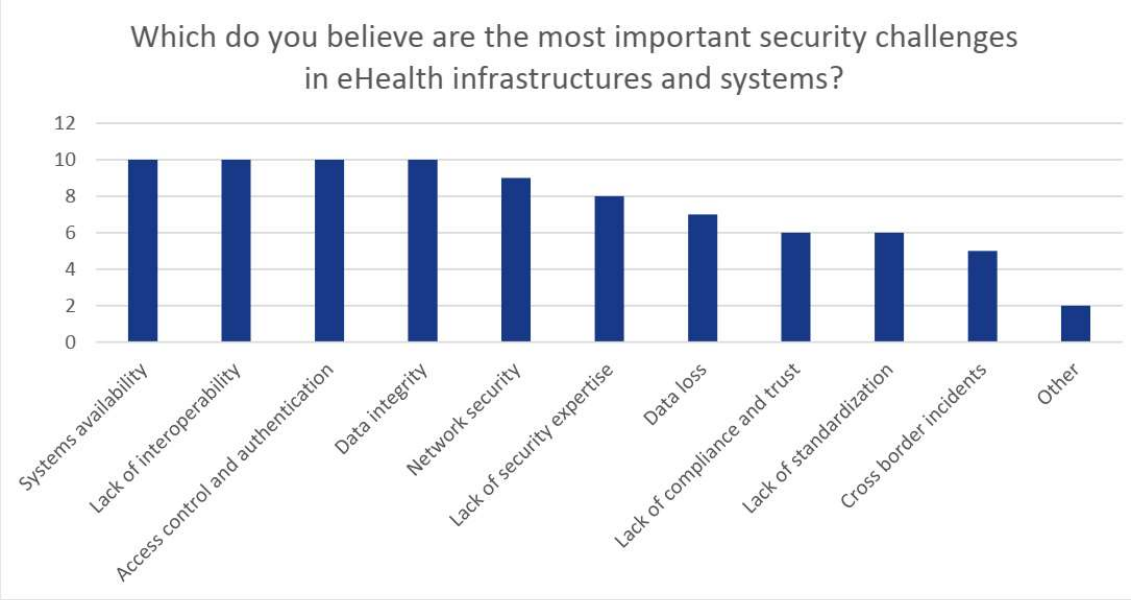

Figure 3-3: Security Challenges in eHealth (from (ENISA, 2015))

In addition to the security challenges, privacy also needs to be maintained with 4 key attributes of (ISO/IEC, 2008):

- Anonymity - a user may use a resource or service without disclosing the user's identity.
- Pseudonymity - a user may use a resource or service without disclosing its user identity but can still be accountable for that use.
- Unlinkability - a user may make multiple uses of resources or services without others being able to link these uses together.
- Unobservability - a user may use a resource or service without others, especially third parties, being able to observe that the resource or service is being used.

The use of constrained IoT devices for RHM can be particularly challenging. Malicious actors (and naiive designers or users) can target threats at the RHM device hardware level

or data link layer as a method of propogating access to the data stored and processed in the cloud (see Section 3.8 for discussion on machine learning on the device). Threat factors include (Fakhr and Fotouhi, 2016, Haddadpajouh et al., 2021, Hammi et al., 2018):

- Low-capacity power supply resulting from the need to be of small size and weight; lower speed processing and making them less capable of carrying out larger computations, and so needing to use light-weight cryptographic mechanisms with low storage overhead.
- There are few widely adopted security standards, but many proprietary protocols which may not be widely researched and examined.
- Homogeneity of devices may introduce problems in cryptographic exchanges and the many types of operating system used by IoT devices are often not as mature or tested as IT systems.
- Authentication protocol needing to be sufficiently fast to avoid being overwhelmed by a denial-of-service attack.
- Spoofing attack - spoofing can happen by Sybil nodes which impersonate themselves for degrading IoT environment functionality. This attack can cause a huge denial of service for the legitimate nodes due to consuming network resources by introducing fake nodes.
- Authorisation mechanisms that enable a non-expert user to allow an emergency clinician to access their medical records, without compromising the overall level of security.
- Insecure boot-up/initialization, which allows for the introduction of malware because the integrity of software updates is not guaranteed.
- Insecure interfaces – can be a route to attack the wider service.
- Wireless attacks - most of the RHM device communications are wireless, which presents opportunities for flooding, eavesdropping, replay attack, command injection and man-in-the-middle attacks. This may also create a vulnerability in the wider service, potentially exposing many users to attack.
- Exhaustion or sleep deprivation attack – any vulnerability that results in the IoT edge device constantly making measurements or transmitting data can lead to a rapid draining of the battery and so loss of service. This can occur when a large number of tasks get assigned to a single node.

Data confidentiality and integrity are typically addressed with encryption. There are two widely used classes of algorithms, namely symmetric encryption and assymetric or public key encryption algorithms.

Symmetric algorithms are computationally light, but rely on secure key exchange and storage, and if the same key is shared between more than two parties then this increases the risk of compromise which could impact all users within that group, including their historical data.

Public key encryption is computationally more expensive and requires a PKI (public key infrastructure - a trusted repository of a device's public key) where every user has a certificate generated by the certificate authority (CA), to bind the device's identity and public key. The management of certificates becomes more demanding as the number of users grows (He et al., 2016). Public-key encryption algorithms used for IoT include Rabin's Scheme, NtruEncrypt and elliptic curve cryptography (ECC), which offers good scalability and can provide the same security as traditional public key cryptogrpahy with a much smaller key size. However, the application of these algorithms in IoT environments is still being investigated (Hammi et al., 2018).

3.2.2 Security Frameworks

3.2.2.1 ETSI EN 303 645 Cyber Security for Consumer Internet of Things: Baseline Requirements

The European Telecommunications Standards Institute (ETSI) operates on a not-for-profit basis and is one of only three bodies officially recognized by the EU as a European Standards Organization (ESO) in the fields of telecommunications, broadcasting and other electronic communications networks and services. ETSI has a special role in Europe to support European regulations and legislation through the creation of Harmonised European Standards. Only standards developed by the three ESOs (CEN, CENELEC and ETSI) are recognized as European Standards (ENs) (https://www.etsi.org/about).

In February 2019, ETSI, the European Standards Organisation, launched the first globally-applicable industry standard on internet-connected consumer devices. ETSI Technical Specification 103 645 brings together what is widely considered good practice in consumer IoT security. ETSI European Standard 303 645 published in June 2020 establishes a security baseline for Internet-connected consumer devices and provides a basis for future Internet of Things product certification schemes (ETSI industry standard based on the Code of Practice, UK Department of Digital, Culture, Media and Sport (DCMS), July 2020). ETSI EN 303 645 largely originated in the DCMS Code of Practice (Department for Digital, 2018).

ETSI EN 303 645 (ETSI, 2020) advocates 13 major provisions for security of consumer IoT, namely:

1. *No universal default passwords – unless in factory default state, all consumer IoT device passwords shall be unique per device or defined by the user.* Unfortunately, *many IoT devices are sold with universal usernames and passwords, such as "admin, admin" which are readily available by searching on the internet. Devices should be manufactured with unique passwords or should enforce the user to choose a password during initialisation that follows best practice, such as NIST SP 800-63B. Authentication mechanisms used to authenticate users against a device shall use best practice cryptography, appropriate to the properties of the technology, risk and usage. When the device is capable, it should prevent brute-force authentication attacks, such as limiting the number of authentication attempts within escalating time intervals or locking out an account after a number of failed attempts*. An RHM service provider should review candidate devices and how passwords are instantiated.
2. *Implement a means to manage reports of vulnerabilities - The manufacturer shall make a vulnerability disclosure policy publicly available. This policy shall include, at a minimum:*
    a. *contact information for the reporting of issues; and*
    b. *information on timelines for:*
        i. *initial acknowledgement of receipt; and*
        ii. *status updates until the resolution of the reported issues.*

Coordinated Vulnerability Disclosure (CVD) is a set of processes for dealing with disclosures about potential security vulnerabilities and to support the remediation of these vulnerabilities. CVD is standardized by the International Organization for Standardization (ISO) in the ISO/IEC 29147. Although CVD has *been successful in some large software companies, in the IoT industry, CVD is currently not well-established as some companies are reticent about dealing with security researchers.*

Disclosed vulnerabilities should be acted on in a timely manner.

Manufacturers should continually monitor for, identify and rectify security vulnerabilities within products and services they sell, produce, have produced and services they operate during the defined support period. Unfortunately, many manufacturers do not specify any support period and products are not updated during their typical operational lifetime. This exposes devices (and their connected services) to vulnerabilities discovered after manufacture. Another major concern is understanding each component in the Software Bill of Materials (SBoM), as many parts of code are copied as (open-source) libraries with further embedded code which is not itemised or maintained by the original developer.

The IoTSF (IoT Security Foundation) publish a Best Practice Guide on implementing ISO/IEC 29147 and an annual report Consumer IoT Vulnerability Disclosure Policy Status. The 2023 report reviewed the practice of 332 IoT vendors and the main finding was that vulnerability disclosure practice remains at a '*disappointingly low level'* with just 27.1% of firms having a disclosure policy. An RHM service provider should review the CVD policy of potential suppliers to gauge how committed they are to maintaining security through the life of device; uncorrected software vulnerabilities could lead to the service provider being responsible for a large exposure of user PII, with potential fines by regulators and a loss of reputation with customers.

3. *Keep software updated - Developing and deploying security updates in a timely manner is one of the most important actions a manufacturer can take to protect its customers and the wider technical ecosystem. It is good practice that all software is kept updated and well maintained.* However, this is dependent upon the update capabilities of the device. *For example, the first stage boot loader on a device is written once to device storage and from then on is immutable*. The

capability of a device to have a secure boot function is critical to ensuring that only software and firmware updates from the manufacturer are processed, rather than malicious code. However, this may also be a major disadvantage should the manufacturer go out of business. Secure boot is discussed more in Section 3.2.4. In addition, any updates should be authenticated against the supplier and integrity checked (typically using hash codes such as SHA-3 (Secure Hash Algorithm 3). Users should be notified when software updates are available, and the updates should be easy to install, or with the option to install manufacturer updates automatically. For an RHM service provider, it should be a business priority to select hardware vendors that are committed to updating device software for the service life of the device; obsolete software with known vulnerabilities may necessitate the replacement of all devices which could be costly to implement.

4. *Securely store sensitive security parameters - sensitive security parameters in persistent storage (and memory) shall be stored securely by the device, such as in a Trusted Execution Environment (TEE) or encrypted storage. A hard-coded unique device identity needs to resist tampering by physical, electrical or software means. In addition, critical security parameters must not be contained within the device software source code.* There have been many instances where security researchers have accessed the device software (such as through exposed USB or internal JTAG ports), read the source code or firmware, and have identified undocumented credentials, such as those used for development or testing that have not been removed in the production environment.
5. *Communicate securely - the consumer IoT device shall use best practice cryptography to communicate securely. Appropriateness .....is dependent on many factors including the usage context. As security is ever-evolving it is difficult to give prescriptive advice about cryptography or other security measures without the risk of such advice quickly becoming obsolete. Cryptographic algorithms and primitives should be updateable.* Cryptography on constrained devices was difficult to implement due to the computing requirements requiring a large percentage of the processor time, reducing the time available for other core functionality and also draining the battery quickly. In recent times, many microprocessors have been designed and manufactured with specialised cryptography modules that dramatically reduce the burden. For RHM, selection

of hardware with on-board cryptography modules is essential. This is discussed further in Section 3.2.4.

6. *Minimize exposed attack surfaces - The "principle of least privilege" is a foundation stone of good security engineering, applicable to IoT as much as in any other field of application. All unused network and logical interfaces shall be disabled, and hardware should not unnecessarily expose physical interfaces to attack (e.g.: A micro-USB port meant to be used to power the device only is physically configured so as not to also allow command or debug operations). The manufacturer should follow secure development processes for software deployed on the device.* Although each of these provisions are sensible, they are unlikely to be achieved unless security is considered as a critical requirement at the outset. Attempting to close exposed features late in the development stage may inadvertently impact on core functionality or leave routes for malicious actors. Additionally, when a device is updated (in hardware, firmware or software), there is a chance that the closed exposed surfaces are re-enabled.
7. *Ensure software integrity - the consumer IoT device should verify its software using secure boot mechanisms. A hardware root of trust is one way to provide strong attestation as part of a secure boot mechanism. A hardware root of trust is a component of a system from which all other components derive their "trust" - i.e., the source of cryptographic trust within that system. To fulfil its function, the hardware root of trust is reliable and resistant to both physical and logical tampering, as there is no mechanism to determine that the component has failed or been altered. By utilizing a hardware root of trust, a device can have confidence in results of cryptographic functions, such as those utilized for secure boot.* An RHM service provider can readily review whether candidate devices use microprocessors from reputable manufacturers who are implementing hardware root of trust (also see Section 3.2.4). What is more difficult to determine is whether the critical software components are correctly using the root of trust during initialisation.
8. *Ensure that personal data is secure - the confidentiality of personal data transiting between a device and a service, especially associated services, should be protected, with best practice cryptography. The confidentiality of sensitive personal data communicated between the device and associated services shall be*

protected, with cryptography appropriate to the properties of the technology and usage. All external sensing capabilities of the device shall be documented in an accessible way that is clear and transparent for the user (e.g.: the incorporation of a microphone or camera within the device). By definition, an RHM device is almost certainly handling sensitive data. It may be that an RHM service provider has to select a more capable, more expensive and potentially larger device (battery) in order to implement security that safeguards against breaches of confidentiality over the lifetime of the device. Considering that a device may be operational for 10 years or longer, the advances in computing that will be available to adversaries in that timespan are likely to render cryptography that is adequate today as too weak within its designed life. Therefore, devices need to be designed with much stronger cryptography than that deemed sufficient for today's threats.

9. *Make systems resilient to outages - the aim is to ensure that IoT services are kept up and running, including in functions that are relevant to personal safety. Keeping services running locally if there is a loss of network is one of the measures that can be taken to increase resilience. Other measures can include building redundancy into associated services as well as mitigations against Distributed Denial of Service (DDoS) attacks.* Note that safety-related applications (e.g., medical devices regulating insulin) will have very stringent resilience regulations. More generally, an RHM service provider will need to plan for intermittent communications, batteries being depleted, devices being inadvertently switched-off or damaged simply because of the users and environments that devices will be in, especially for devices that are on the user and not protected within a cabinet. Contingency methods, such as using landline phones or physically visiting users may need to be available to deal with outages. Note that any contingency method may not be able to rely on the end user to make decisions or take actions.
10. *Examine system telemetry data - if telemetry data is collected from consumer IoT devices and services, such as usage and measurement data, it should be examined for security anomalies. Security anomalies can be represented by a deviation from normal behaviour of the device…. for example, an abnormal increase of failed login attempts. Examining telemetry, including log data, is useful for security*

evaluation and allows for unusual circumstances to be identified early and dealt with, minimizing security risk and allowing quick mitigation of problems. Provisions specific to protecting personal data when telemetry data is collected must be included. Telemetry can reveal a wide range of anomalies, whether the cause is a software, hardware or communications problem or a malicious attack. The RHM may also be able to identify a threat against an individual user or device (potentially to harm that individual), or a more concerted attack against the central service which may indicate a ransomware or divulgence of sensitive files. Security monitoring needs to be conducted 24/7 as the threats can be global, and automated tools can assist to alert when there is anomalous behaviour.

11. *Make it easy for users to delete user data - The user shall be provided with functionality such that user data can be erased from the device and associated services in a simple manner (i.e.: minimal steps involving minimal complexity are required to complete that action). Such functionality is intended for situations when there is a transfer of ownership, when the consumer wishes to delete personal data, when the consumer wishes to remove a service from the device and/or when the consumer wishes to dispose of the device. It is expected that such functionality is compliant to applicable data protection law, including the GDPR. Such functionality can potentially present an attack vector.* This provision may need to be tailored for users of an RHM service, as a user inadvertently, or a malicious actor deliberately removing data or a profile from the care service could be a worse outcome. Certainly, there needs to be a simple method by which the user (or their representative) can request that their data be deleted from the device and core databases and receive confirmation from the RHM service provider. The other situation where deletion of data is necessary is when a device is handed from one user (patient, carer, clinician etc.) to another; the new recipient should not be able to access data from a previous user.
12. *Make installation and maintenance of devices easy – this should involve minimal decisions by the user and should follow security best practice on usability. The manufacturer should provide users with guidance on how to securely set up their device, or ideally, a process which achieves a secure configuration automatically. The manufacturer should provide users with guidance on how to check whether their device is securely set up. Security issues caused by consumer confusion or*

misconfiguration can be reduced and sometimes eliminated by properly addressing complexity and poor design in user interfaces. Clear guidance to users on how to configure devices securely can also reduce their exposure to threats. In the general case, the average overhead of securely setting up a device is higher than the average overhead of checking whether a device is securely setup. The check of a secure setup, from a process standpoint, can be undertaken in large part by the manufacturer through an automated process that communicates with the device remotely. Part of such an automated process could include validation of the device's capacity to establish a secure communication channel. Again, for an RHM service provider, the users may not be familiar with setting up devices, not only principal users who may be living with cognitive or dexterity challenges, but also the carers and clinicians, especially when the instructions may not be available in their native language or dialect. A minimal pre-configuration of a device (with no personal data in case the device is lost during despatch) so that it can securely communicate with the RHM core servers, followed by remote onboarding and configuration of the device once there is confirmation that the user has received the device, may be the most reliable method of installing and maintaining devices.

13. *Validate input data - the device software shall validate data input via user interfaces or transferred via Application Programming Interfaces (APIs) or between networks in services and devices. Systems can be subverted by incorrectly formatted data or code transferred across different types of interface. Automated tools such as fuzzers can be used by attackers or testers to exploit potential gaps and weaknesses that emerge as a result of not validating data. For example, out of range data is received by a temperature sensor, rather than trying to process this input it identifies that it is outside of the possible bounds and is discarded and the event is captured in telemetry.* Such range bounding may need to be customisable per principal in an RHM service. For example, a user living with COPD (chronic obstructive pulmonary disease) will constantly have blood oxygen saturation levels that would be considered dangerously low for the general population, while other users may normally have irregular or unusually high or low heart rates that would trigger interventions for other users. RHM devices used by remote workers in very cold or hot environments may encounter temperatures

> unpredicted by the device designers. How a user's profile is adjusted to accommodate their specific conditions needs to be carefully controlled so that any deterioration is quickly identified and not assumed to be 'normal' for them.

This standard was developed by industry and was not commissioned by the EU and so is not a legal requirement. However, adoption of ETSI EN 303 645 helps to eliminate some of most common methods that hackers use to attempt to breach security, but does not address more advanced threats from the most persistent and advanced malicious actors. Using devices that are certified against the standards should reduce the RHM's exposure to security vulnerabilities and may help demonstrate to a regulator that care had been taken to reduce the risk of data breaches. Tailoring of the provisions may be necessary in recognition of any dexterity, cognitive or language challenges faced by the users. In LDCs, the increase in cost resulting from preparing the evidence and submitting the device to independent assessment may reduce affordability.

3.2.2.2 NIST IR 8259A IoT Device Cybersecurity Capability Core Baseline

The National Institute of Standards and Technology (NIST) was founded in 1901 and is now part of the U.S. Department of Commerce. Its Mission is "*To promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve our quality of life*" (nist.gov/about-nist). NIST's Computer Security Resource Center (CSRC) has for 20 years provided access to NIST's cybersecurity- and information security-related projects, publications, news and events. CSRC supports stakeholders in government, industry and academia, both in the U.S. and internationally (csrc.nist.gov/). NIST, as a US Government organisation, is highly influential in North America and therefore its practices are often incorporated by manufacturers and suppliers globally who wish to operate in the USA.

In 2016, NIST established the Cybersecurity for the Internet of Things (IoT) program to support the development and application of standards, guidelines, and related tools to improve the cybersecurity of connected devices, products and the environments in which they are deployed. A series of documents have been published, including NIST IR 8259A IoT Device Cybersecurity Capability Core Baseline (NIST, 2020). The features defined

are '*a set of device capabilities generally needed to support commonly used cybersecurity controls that protect devices as well as device data, systems, and ecosystem*s'. It draws upon 15 other IoT cybersecurity frameworks in a '*co-ordinated effort to produce a definition of common capabilities, not an exhaustive lis*t'.

The Baseline has six core capabilities, each with a number of common elements. They are:

1. *Device Identification: The IoT device can be uniquely identified logically and physically. This capability supports asset management, which in turn supports vulnerability management, access management. This is to include both (1) a unique logical identifier and (2) a unique physical identifier at an external or internal location on the device authorized entities can access.* This is an essential capability for identifying that the data purporting to come from a specific RHM is not being spoofed. Likewise, data being sent to the device needs to go to the correct user otherwise incorrect medication could be administered.
2. *Device Configuration: The configuration of the IoT device's software can be changed, and such changes can be performed by authorized entities only. This supports vulnerability management, access management, data protection, and incident detection.* It may be that an RHM needs to be updated quickly to address some configuration issue that has been identified across the population of users. Because of the remote nature of RHM principals, manual configuration updates may take a long time and be costly to implement. Therefore, configuration changes may need to be implemented remotely to reduce the time that users may be at risk of faults in the software, which, in some cases, could be a safety issue. It is also important that hackers cannot affect the operation of RHM devices to potentially null warnings of when a user may have a health issue, or conversely to avoid flooding the system with alarms.
3. *Data Protection: The IoT device can protect the data it stores and transmits from unauthorized access and modification. This capability supports access management, data protection, and incident detection. Protection may include the ability to use demonstrably secure cryptographic modules for standardized cryptographic algorithms (e.g., encryption with authentication, cryptographic hashes, digital signature validation).* This is fundamental to health monitoring

data security and privacy. An interesting point is that although devices may use well-known and researched cryptographic algorithms, it is the software implementation that sometimes leaves weaknesses that can be exploited, such as using predictable seeds for random numbers or using programming libraries that are later found to have vulnerabilities. Note that 'security through obscurity' should not be followed, and the adoption of open modules that have been widely scrutinised are more likely to be secure.

4. *Logical Access to Interfaces: The IoT device can restrict logical access to its local and network interfaces, and the protocols and services used by those interfaces, to authorized entities only. This capability supports vulnerability management, access management, data protection, and incident detection.* A typical example where this has failed is a USB port intended for charging a device, but which also allows an unauthorised entity to access the configurations, data files and software. Other examples include (JTAG) connectors on the circuit board used for development testing being left on production boards.
5. *Software Update: The IoT device's software can be updated by authorized entities only using a secure and configurable mechanism. This supports vulnerability management. This includes the ability to update the device's software remotely or locally and the ability to verify and authenticate any update before installing it. Authorised entities should also be able to roll back updated software to a previous version.* Similar to configuration updates, a device's software may need to be updated to address a functional issue or because a vulnerability has been identified. It is critical that only authorised and authenticated updates from the manufacturer (or other know support entity) can be loaded onto the device, in order to avoid malicious actors introducing malware.
6. *Cybersecurity State Awareness: The IoT device can report on its cybersecurity state and make that information accessible to authorized entities only. This capability supports vulnerability management and incident detection. The ability to differentiate between when a device will likely operate as expected from when it may be in a degraded cybersecurity state.* This can be more difficult to implement on constrained devices that are unable to detect intrusion detection. Also, the amount of information that can be logged and transmitted back to the core will be dependent upon the device and communications channel.

Overall, the NIST 8259A baseline requirements are less comprehensive than ETSI EN 303 645. However, they will still become influential if they are mandated to be used in US Government contracts and suppliers believe that they need to comply in order to retain access to these customers.

3.2.2.3 IoTSF Assurance Framework

The author of this thesis was an Editor of Release 3.0 and was particularly involved with the mapping of requirements to ETSI EN 303 645 and NIST IR 8259A and defining new requirements in the Assurance Questionnaire.

The Internet of Things Security Foundation (IoTSF) is a non-profit, global membership organisation striving to make the connected world ever-more secure (https://www.iotsecurityfoundation.org/). It was launched on 23rd September 2015 and in 2016 published the IoTSF Compliance Framework, one of the pioneering IoT security frameworks and from which many of the requirements found in the UK DCMS Code of Practice and ETSI EN 303 645 originated. In Release 3.0, it was changed to the IoT Security Assurance Framework (IoT Security Foundation, 2021), with the emphasis moving from compliance to internal assurance that can be followed throughout the whole lifecycle of an IoT device or service and which leads its user through a structured process of questioning and evidence gathering. This is based upon the premise that "*providing good security capability requires decisions upfront in design and use – often referred to as secure by design. In most cases, addressing the security of a product at the design stage is proven to be lower cost, and requiring less effort than trying to "put security" into or around a product after it has been created (which may not even be possible*)". Its intended audience includes managers, developers and engineers, logistics and manufacturing staff and supply chain managers. However, the scope of its applicability is much wider than ETSI EN 303 645 or NIST IR 8295A as it can also be used by buyers, service delivery organisations and users to assess how 'secure by design' has been incorporated into the supporting infrastructure including web user interface, mobile apps, cloud and network elements, configuration and device ownership transfer. The Assurance Framework (AF) is freely available and has had over 10,000 downloads, and there is an in-depth Assurance Questionnaire available to Members of the IOTSF.

A feature of the IoTSF AF is that it starts with a risk analysis and then determines the assurance class applicable to the product. Control features are then tailored to provide suitable protections against the security objectives.

It addresses the 14 major areas of development and in-service:

1. Business Security Processes, Policies and Responsibilities
2. Device Hardware & Physical Security
3. Device Software
4. Device Operating System
5. Device Wired and Wireless Interfaces
6. Authentication and Authorisation
7. Encryption and Key Management for Hardware
8. Web User Interface
9. Mobile Application
10. Data Protection and Privacy
11. Cloud and Network Elements
12. Secure Supply Chain and Production
13. Configuration
14. Device Ownership Transfer

The last 7 assurance areas ensure that not only is an IoT device built securely, but that it can be deployed and operated securely throughout its lifecycle, including disposal.

3.2.2.4 Alignment of Frameworks

This section refers to the following published reference document (see Appendix B):

> IoTSF (Ian Poyner) (2021): *IoTSF-Framework-to-IR8259A (1.0.0)*, dated 10th November 2021, published on NIST's National Online Informative References Program (OLIR)

As one of the main Editors of the IoTSF Assurance Framework Release 3.0, the author of this thesis was particularly involved with the mapping of requirements to ETSI EN 303 645 and NIST IR 8259A.

There is a risk that the rapidly evolving IoT market becomes fragmented and that security capabilities are not cost-effective for manufactures to incorporate. This risk increases if each jurisdiction or region issues similar but different, or even conflicting, security requirements for its market. Therefore, to encourage manufacturers to invest in a set of security capabilities that can be sold across many markets, there is a need to improve harmonisation and coherency across frameworks influential in Europe, North America and globally. Even though these frameworks may not be legal requirements, they would be influential if they are specified in contracts (particularly for government projects) and may form the basis of future regulatory legislation. As such, the IoTSF Assurance Framework Release 3.0 was intended to provide a super-set of requirements, using common terminology to ETSI EN 303 645 and NIST IR 8259A against which suppliers can assess whether their products are likely to comply with requirements across multiple regions.

As part of this aligning of frameworks, there was considerable engagement with NIST's National Online Informative References (OLIR) Program. OLIR is a NIST effort to facilitate subject matter experts (SMEs) in defining standardized online informative references (OLIRs) between elements of their documents, products, and services and elements of NIST documents like the Cybersecurity Framework Version, Privacy Framework Version, NISTIR 8259A, or NIST SP 800-53. A mapping of the IoTSF Assurance Framework to NIST IR 8259A was written by Poyner, approved by OLIR and published on NIST's OLIR catalogue.

As the IoTSF AF had evolved into ETSI EN 303 645 (along with a German framework), then for the first time, industry, service providers and users had a clear linkage between the requirements of the leading guidance in Europe and North America.

3.2.3 Secure by Design and Data Protection by Default

Secure by Design, referred to as SbD, has been led in the UK by the Department for Digital, Culture, Media & Sport (DCMS) with the National Cyber Security Centre providing the technical guidance. Following high-profile events, such as the 2016 Mirai and 2017 WannaCry attacks, in March 2018 DCMS published the Secure by Design report "*which advocated a fundamental shift in approach to securing IoT devices, by moving the burden away from consumers and ensuring that security is built into products*

by design. Central to the report was a draft Code of Practice (Department for Digital, 2018) primarily for manufacturers of consumer IoT devices and associated services. An informal consultation on the report and its proposed policy interventions was undertaken" (https://www.gov.uk/government/publications/secure-by-design-report).

The National Cyber Security Centre's (NCSC) Secure Design Principles (NCSC, 2019) are divided into five categories, loosely aligned with stages at which an attack can be mitigated:

1. *Establish the context - Determine all the elements which compose your system, so your defensive measures will have no blind spots.*
    1. *Understand what the system is for, what is needed to operate it, and which risks are acceptable.* In an RHM service, a breach of sensitive data would be highly unacceptable, as well as any interference with the integrity of the health data and any information or actions going from the core to the principal. Also, financial data is almost always a target for hackers.
    2. *Understand the threat model for your system.* This will need input from the service users (carers, supervisors etc.) as it may be different from many IoT deployments because the devices are usually on the person who may not have the cognitive skills to protect the device from loss or interference, even for workers in remote locations who may mislay a device or have higher priorities.
    3. *Understand the role of suppliers in establishing and maintaining system security.* Suppliers will need to understand that they may potentially be processing sensitive personal data, as defined under GDPR. Telemetry to the manufacturer 'to help understand any crashes' will need to be carefully defined and constrained; they should not be able to harvest whatever information they feel may be useful for their purposes.
    4. *Understand the system 'end-to-end'.* This may include how data is stored in the core and how it may also be shared with clinician or hospital systems.
    5. *Be clear about how you govern security risks.* An RHM service has a greater challenge than systems monitoring inanimate objects because the

trade-offs between usability and privacy, including potentially multiple users (carers, clinicians, supervisors) all needing access.

6. *Ensure there is no ambiguity about responsibilities.* This could potentially be an area that causes vulnerabilities in an RHM service, for example where a local social services or clinic contracts with a service provider who then further sub-contracts with a carer agency. The data ownership and who can access what, needs to be carefully defined.

2. *Making compromise difficult - An attacker can only target the parts of a system they can reach. Make your system as difficult to penetrate as possible.*
    1. *External input can't be trusted. Transform, validate, or render it safely.* Validating cryptographic signatures of data works for many systems, but there is a risk that a constrained battery-operated device will only support short keys (e.g., 128-bit) which may be trivial to break over a 10-year life cycle. Although core services may be able to transform data, in the reverse direction, it would be more difficult to implement transformation or rendering in a sandboxed environment for the end user device.
    2. *Reduce attack surface.* This refers as much to the logical interfaces as physical. The main area of concern may be at the support end (carers, clinicians, supervisors) who should, in general, only be able to view data, rather than manipulate the raw data and access the business systems.
    3. *Gain confidence in crucial security controls.* It is challenging for a care service or remote workforce company to take on the cyber security aspects, however, this is a fundamental requirement for any service handling PHI. Trying to understand how service providers and device suppliers have addressed security may require hiring or contracting expertise in that area.
    4. *Protect management and operations environments from targeted attacks.*
    5. *Prefer tried and tested approaches.* Using popular libraries of cryptographic algorithms should mean that they have been examined by the community and any new vulnerabilities are more likely to identified and corrected than a bespoke solution.
    6. *All operations should be individually authorised and accounted for.* This is a critical requirement for an RHM – group accounts should not be used

as any breach cannot be attributed to an individual. However, this may be more challenging when providing access to external systems, such as a GP's surgery.

7. *Design for easy maintenance.* This is more applicable to the core services, but any software on the user's device should be the minimum required, well-structured and updateable. It is a general principle that any safety critical device should have the minimum functionality and software required to conduct the critical operations.
8. *Make it easy for administrators to manage access control*
9. *Make it easy for users to do the right thing.* For an RHM device, this could be interpreted as make it easy for users not to have to do anything (disappearing interfaces).

3. *Making disruption difficult - Design a system that is resilient to denial-of-service attacks and usage spikes.*
    1. *Ensure systems are resilient to both attack and failure.* If a principal is of very high concern, then there may be a need to have two or more independent sensors with no common hardware, software or communications so as to make it more difficult for a random failure or attacker to disrupt the data.
    2. *Design for scalability.* This is usually associated with the core services and nowadays is usually addressed with cloud hosting contracts.
    3. *Identify bottlenecks, test for high load and denial of service conditions.* In an RHM service, this may be an attack on the IoT communications, especially if operating in the unlicensed ISM band (see Section 3.5).
    4. *Identify where availability depends on a third party and plan for the failure of that third party.* IoT is still a relatively young market, with a risk that some of the technologies and suppliers' business models may not survive. For example, the Sigfox communications network (see Section 3.5.1.6 was founded in 2010 but filed for bankruptcy in January 2022, eventually being acquired by Unabiz in April 2022. This was at least partly due to competition from the decrease in prices of cellular networks and the emergence of competing IoT networks such as LoRaWAN (see Section 3.5.1.1). The lesson is that if an RHM service relies on a

proprietary service or manufacturer then there is a risk that the service will be impacted if that supplier fails.

4. *Making compromise detection easier - Design your system so you can spot suspicious activity as it happens and take necessary action.*
    1. *Collect all relevant security events and logs.* Although this is excellent advice for a typical IT environment, there are constraints for RHM devices in how to identify and route security events to an audit facility. It may be that there is a stage of filtering incoming data at the core system.
    2. *Design simple communication flows between components.* This common-sense guidance will be complicated in situations where there is an urgent need to share information to third parties, such as hospitals, police or rescue teams for remote workers).
    3. *Detect malware command and control communications*
    4. *Make monitoring independent of the system being monitored.* This may be effective at the supporters end of the service. For example, data received from the RHM device could be viewed by carers or supervisors, but anything that controls the configuration or operation of the device may require an independent route.
    5. *Make it difficult for attackers to detect security rules through external testing*
    6. *Understand 'normal' and detect the abnormal.* This may be challenging detecting what is abnormal for the pattern of life of a remote worker or someone living with health conditions. However, the service provider may have to accept a high false alarm rate in order to reduce the opportunity of an attacker breaching the system.
5. *Reducing the impact of compromise - If an attacker succeeds in gaining a foothold, they will then move to exploit your system. Make this as difficult as possible.*
    1. *Use a zoned or segmented network approach.* Ideally, each principal's device and data would be segmented from all others. However, researchers and service providers may wish to analyse data across multiple users in order identify trends or better methods of treatment.

2. *Remove unnecessary functionality, especially where unauthorised use would be damaging.* As commented earlier, an RHM device should have the minimum functionality required for the principal user. If a principal develops new conditions, then the device could be modified or replaced as necessary.
3. *Beware of creating a 'management bypass'.* This is more likely to occur in the interfaces from the core to the carers, clinicians or supervisors.
4. *Make it easy to recover following a compromise.*
5. *Design to support 'separation of duties'.* The two-person rule for administrators should be followed, where if a person can copy, modify or delete data then they should have no access to auditing logs, and vice versa.
6. *Anonymise data when it's exported to reporting tools.* This a major challenge for PHI – it is difficult to demonstrate that a principal cannot be re-identified, especially if they are living with a relatively uncommon condition for that region. 'Big data' analysis may offer insights into medical treatments, but the privacy of individuals must be maintained.
7. *Don't allow arbitrary queries against your data*
8. *Avoid unnecessary caches of data.* This may occur where a principal's data is copied to a separate 'drop-box' in order to be analysed by a hospital or an emergency service. Here the benefits to the user could outweigh potential loss of privacy, but the data should be purged as soon as it is no longer required by the external party.

Overall, the NCSC Secure by Design Principles include sound advice for every service provider. However, the specific challenges of an RHM service will need to be carefully assessed to understand the risks and the balance between the benefits to the principal of sharing personal data against the risks of breaches of privacy for both the individual and multiple users. The important thing for any RHM service is to adopt the NCSC's SbD Principles as it is authoritative guidance, but where it is necessary or chosen on health benefits to deviate for from the guidance then the reasoning must be clearly documented together with any mitigations. This may offer some protection in the event of a data breach as to why not all of the NCSC's principles were followed.

A related concept is “Data protection by design and default”, a legal requirement under GDPR (Articles 25(1) and 25(2) of UK GDPR, although privacy be design has always been part of data protection law (ICO, 2022). It requires data protection to be integrated into processing activities and business practices, from the design stage right through the lifecycle.

Data protection by design is an approach that ensures that privacy and data protection issues are considered at the design phase of any system, service, product or process and then throughout the lifecycle. UK GDPR requires that:

- Appropriate technical and organisational measures are designed to implement the data protection principles effectively.
- Safeguards are integrated into processing so that it meets the UK GDPR's requirements and protect individual rights.

Data protection by default requires that only data that is necessary to achieve a specific purpose is processed. It links to the fundamental data protection principles of data minimisation and purpose limitation. For an RHM service, personal data needs to be processed to achieve health and safeguarding benefits. Data protection by default means that this data is specified before the processing starts, individuals (or their representatives) are appropriately informed and only the data needed to provide an RHM service is processed. It does not prohibit data sharing or having to have a ‘default to off’ solution; it depends on the circumstances and the services being provided and the risks posed to individuals. Considerations must include:

- Adopting a ‘privacy-first’ approach with any default settings of systems and applications.
- Ensuring that an illusory choice is not provided to individuals relating to the data that will be processed.
- Not processing additional data unless the individual decides it can.
- Ensuring that personal data is not automatically made publicly available to others unless the individual decides to make it so.
- Providing individuals with sufficient controls and options to exercise their rights.

A challenge for an RHM service provider using constrained IoT devices is that the stream of physiological data gathered by the principal's sensing device may be difficult to filter at the device and so all of it may need to be transmitted to the core servers and unnecessary data filtered (deleted) there. However, in some circumstances, such as a medical emergency, the benefits of enabling all data to be viewed by other parties, such as clinicians, may outweigh privacy concerns, unless the principal (or their representative) has specifically stated that it is not to be used. How that data is 'turned on' to be viewed by the clinicians, and at what point it is then filtered/deleted again will need careful business processes to decide.

3.2.4 Manufacturing Security into IoT Devices

Some device designers and manufacturers are actively responding to the need for security built into constrained IoT devices, such as microcontrollers.

One example is PSA Certified, which launched in 2019 and "*offers a framework for securing connected devices, from analysis through to security assessment and certification. The framework provides standardized resources to help resolve the growing fragmentation of IoT requirements and ensure security is no longer a barrier to product development*" https://www.psacertified.org/what-is-psa-certified/ . It seeks to create a *secure-by-design culture where security is implemented from the beginning of product development*. It builds-in 10 security goals (largely aligned to the physical control measures advocated in the IoTSF AF and ETSI EN 303 645) as shown in Figure 3-4 and incorporating a Root of Trust (RoT), which is a foundational security element, built into the silicon, that completes a set of implicitly trusted functions, upon which higher level system functions can attribute integrity and authentication.

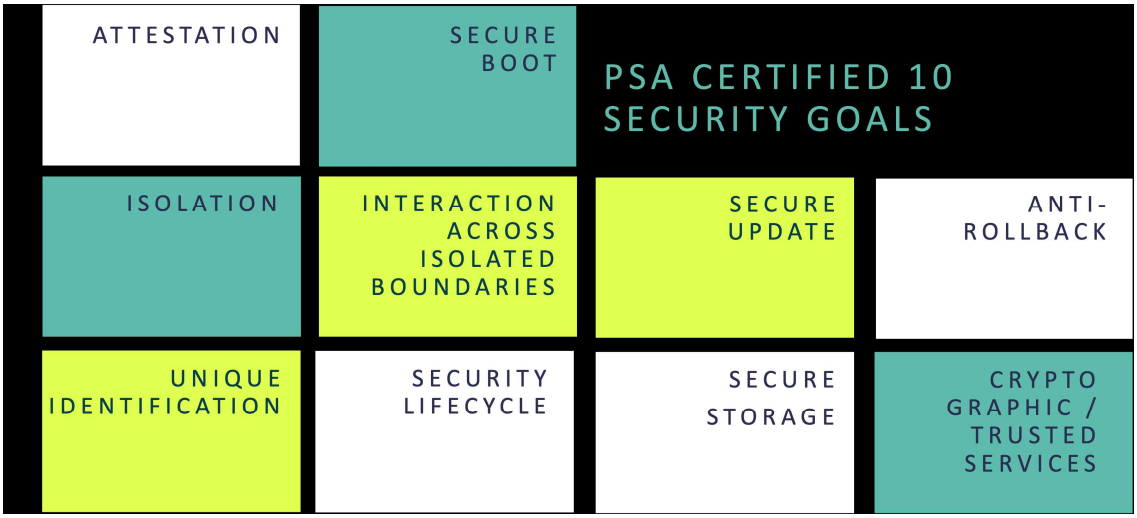

Figure 3-4: PSA Certified 10 Security Goals (from https://www.psacertified.org/what-is-psa-certified/our-approach/)

In RHM applications, the RoT and a "unique and tamper-resistant device identifier" (IoT Security Foundation, 2021) can provide a digital 'birth certificate' that could reduce the risk of a counterfeit or uncertified device being introduced into the service, potentially introducing an attack route for a malicious actor. Note that such an identifier may be linked to a specific microchip on the device, and a complex device may contain several microchips, each with their own identifier. This enables the microchip to be authenticated by the original equipment manufacturer, although it does not convey any assurance around the other components, software or assembly.

The PSA Certified framework has a four-step process for security design and implementation:

- **Analyse** the threats that have the potential to compromise the device and generate a set of security requirements based on these risks.
- **Architect:** Use the security requirements to identify and select components and specifications to architect the appropriate level of security for the product.
- **Implement:** Implement the trusted components and firmware, using high-level APIs to create an interface to the hardware Root of Trust (RoT).
- **Certify:** Following independent security evaluation, certify your device, system software platform or silicon.

Although PSA was originally led by ARM, there are now (as of January 2023) 85 chips, 27 devices and 26 system software products that have been certified from providers across the industry.

The advantage for an RHM service provider using products (chip, device or software) that have been certified (not necessarily by PSA if alternative schemes are available and appropriate) is that it provides evidence that 'industry best practice' has been adopted. For an RHM without a deep understanding of cyber security this can provide reassurance of at least a baseline level of security being incorporated into the product. However, the risk is that the RHM does not understand the limitations of the assurance and that it is not a panacea for all security measures. The RHM would still need to ensure that security has been correctly designed and implemented in the remainder of the device hardware, supporting infrastructure and services and the personnel who have access to the system.

3.2.5 Machine Learning for Security and Privacy

The use of machine learning (ML) embedded into the wearer's device may reduce the amount of personal data transferred to the processing centre and accessible to other users of the service (carers, clinicians, researchers and cloud service companies). An alternative model in widespread use with AAL and smartphone health apps is to upload the user data to cloud services that then perform ML health analytics on 'big-data' in order to generate valuable population-level insights. However, according to Amiri-Zarandi et al. (Amiri-Zarandi et al., 2020), the '*analytics provider must guarantee that they will not misuse data, and secondly, they will be able to prevent data leakage. However, several studies have shown vulnerabilities in these services that must be addressed. In addition, ML models sometimes can be the point of data leakage.....Many valuable data sources can be used and managed by ML to develop novel approaches to protect data. The main barrier to employ these data sources is the lack of standardization and interoperability. We need frameworks that can automatically clean the streaming data and convert them into a standardized format.'*.

Although this has benefits for security and privacy, there are disadvantages in terms of reliability of the ML decisions, energy and processing power required to perform ML and how the early aggregation and filtering of data may remove valuable information for

researchers and clinicians. The use of ML in RHM devices is discussed in more detail in Section 3.8.

3.3 RHM Communication Requirements

3.3.1 Authentication

As discussed in Section 3.2.4, there is a risk of a malicious actor introducing a false device to masquerade as a user or as a service provider. A trusted identifier is required to authenticate a device, such as a SIM card or an identity built upon a root of trust. This is necessary not only for sending data from the RHM, but also to authenticate any updates to the device to be verified and authenticated prior to installation, so reducing the opportunity for malware to be uploaded.

3.3.2 Safety

In a clinical or home environment, body-worn (battery-powered) sensors typically use low-power communications (e.g.: Bluetooth) to transmit data to a (mains-powered) hub; the hub then communicates the data to remote users. However, RHM devices may integrate the long-haul communications within the sensing device, which is usually kept in close proximity to the body. The proximity of the sensor to the body may cause damage to tissue and so power must be limited to an acceptably safe level.

The International Commission on Non-Ionizing Radiation Protection's (ICNIRP) view is that after several decades of RF-EMF research on numerous potential health effects, the only substantiated effect of RF EMF exposure relevant to human health and safety is heating of exposed tissue. RF EMF fields can penetrate into the body (the higher the frequency, the lower the penetration depth) and cause vibration of charged or polar molecules inside. This results in friction and thus heat. The human body has a strong ability to regulate its internal temperature and so can accommodate a small increase in heat. However, above a certain level (referred to as the threshold) depending on the duration of exposure, RF EMF exposure and the accompanying temperature rise can provoke serious health effects, such as heatstroke and tissue damage (burns). Acute and

long-term effects of RF EMF exposure below the thermal threshold have been studied extensively without demonstrating adverse health effects.

Unlicensed IoT communications often operate in the ISM Band ((Industrial, Scientific and Medical), around the 433MHz, 915MHz, 2.4GHz and 5.8GHz frequencies, depending on the ITU region (see Section 3.5). The ICNIRP has published guidelines for limiting exposure to RF fields based upon identified “adverse health effect threshold” - the lowest exposure level known to cause a substantiated (independently verified) health effect. They also set an “operational threshold” based on evidence that is independent from the radiofrequency health literature and which has (indirectly) shown that harm could occur at a lower level. These thresholds are based on the relation between the primary effect of exposure (e.g., heating) and health effect (e.g., pain), and provide restriction values in order to attain an appropriate level of protection. Reduction factors are required to account for biological variability in the population (e.g., age, sex), variation in baseline conditions (e.g., tissue temperature), variation in environmental factors (e.g., air temperature, humidity, clothing), dosimetric uncertainty associated with deriving exposure values, uncertainty associated with the health science, and as a conservative measure more generally (International Commission on Non-Ionizing Radiation Protection, 2020).

The Specific Absorption Rate (SAR) is a measure of the rate at which energy is absorbed per unit mass by a human body when exposed to a radio frequency (RF) electromagnetic field. The human body is a lossy medium where the electrical signals are absorbed by the surrounding tissues. This attenuates signals and increases the temperature of tissues, which poses serious health risks and must be avoided. Unfortunately, there is no absolute safe SAR level and some areas of the body (head and trunk) are more susceptible to damage than other areas (Ahmed et al., 2018).

ICNIRP basic restrictions are shown in Table 2 of (International Commission on Non-Ionizing Radiation Protection, 2020), reproduced here in Table 3-1.

Table 3-1: ICNIRP basic restrictions for EM exposure 100KHz to 300GHz (ICNIRP, 2020)

Table 2. Basic restrictions for electromagnetic field exposure from 100 kHz to 300 GHz, for averaging intervals ≥6 min.[a]

Exposure scenario	Frequency range	Whole-body average SAR (W kg^{-1})	Local Head/Torso SAR (W kg^{-1})	Local Limb SAR (W kg^{-1})	Local S_{ab} (W m^{-2})
Occupational	100 kHz to 6 GHz	0.4	10	20	NA
	>6 to 300 GHz	0.4	NA	NA	100
General public	100 kHz to 6 GHz	0.08	2	4	NA
	>6 to 300 GHz	0.08	NA	NA	20

[a]Note:

1. "NA" signifies "not applicable" and does not need to be taken into account when determining compliance.
2. Whole-body average SAR is to be averaged over 30 min.
3. Local SAR and S_{ab} exposures are to be averaged over 6 min.
4. Local SAR is to be averaged over a 10-g cubic mass.
5. Local S_{ab} is to be averaged over a square 4-cm^2 surface area of the body. Above 30 GHz, an additional constraint is imposed, such that exposure averaged over a square 1-cm^2 surface area of the body is restricted to two times that of the 4-cm^2 restriction.

Note that the basic restrictions are averaged over 6 minutes or longer, and there are local levels for the head/torso.

An implication of basic thresholds on RF exposure (see SAR limits above) is that RHM devices are limited on the power levels of the RF emissions on health grounds, as well as regulations (see Section 3.4). Even though the threshold is higher than what may be sustained by a battery powered device over a prolonged period, limitations in radiated power constrain range, throughput and error rates. The quality of the signal may fluctuate considerably depending upon the placement of the sensor and movement of the wearer, which may cause intermittent drop-outs in communications as the system adjusts RF power levels to compensate.

3.3.3 Data Timeliness

Medical imaging and specialised monitoring systems, such as EMG (electromyography) and ECG (electrocardiographs), generate data at very high rates, whilst other sensors do not produce as much data and only need sampling every few seconds (or minutes). The nominal data rate will depend upon the sensor and how the manufacturer has configured the sensor, but indicative data rates and sampling intervals are shown in Table 3-2 (Arnon et al., 2003, Cordeiro and Patel, 2007, Shu et al., 2015, Akbar et al., 2017, Alam et al., 2018).

Table 3-2: Indicative Data Rates

Application	Information Rate	Resolution (bits)	Sample Rate
SpO_2 Blood Saturation	16bps	8	
Skin Temperature	120 bps	8	< 60 sec
Respiratory rate	240 bps	12	0.1 – 20 Hz
Blood pressure	1200 bps	12	0-100 Hz
Glucose monitoring	1600 bps	16	< 240 sec
Motion sensor	35 kbps	12	25-100 Hz
EEG (12 leads)	43.2 kbps	12	350 Hz
ECG (6 leads)	71 kbps	12	1250 Hz
ECG (12 leads)	288 kbps	12	2500 Hz

In a hospital or nursing home environment that has resources such as system technicians, reliable power supplies and high-speed networks, data from a suite of sensors (and potentially from many principals) can be collected, processed and alerts can be triggered as required and without mutual interference creating errors or delays in data. Alerts for medical staff or carers can be generated with minimal (often sub-second) delay. However, the demand for the delay to not exceed 250 ms, as demanded by Shu et al. (Shu et al., 2015), for glucose monitoring or body temperature appears overly demanding compared with both the sampling rate, which is related to how slowly such physiological parameters change, and how quickly this would be responded to by a carer.

In contrast, such sampling and data rates may not be achievable for RHMs without a significant impact on battery duration. Compromises may need to be made with lower sampling rates and/or resolution which would reduce the amount of data to be transmitted yet would need to remain sufficient for remote carers to intervene.

An alternative approach is for the user device to perform processing and filtering of data. This may require a higher performance, and more costly, processor with higher energy consumption, but could significantly reduce the data transmitted, leading to an overall decrease in energy usage. A major challenge to such an approach is validating that the on-board processing of data does not create errors or misleads core systems and carers

and clinicians. Machine Learning (ML) may facilitate this as it becomes more common on constrained IoT devices. This is explored by Fafoutis et. Al.(Fafoutis et al., 2018) and discussed further in Chapter 5.

3.3.4 Data Integrity and Protection of Sensitive/Personal Health Information

Security has been a significant challenge for IoT devices for the reasons outlined in the ManySecured Secure IoT Gateways whitepaper (ManySecured, 2019), namely:

- Limited or no user interface, so making configuration challenging and limiting feedback to the user.
- Physically insecure due to location and mobility (depends upon the user's ability to safeguard the device). Secrets (tokens and private keys) may be more vulnerable.
- Power constraints, if battery powered, may limit sensor updates and transmission, as well as security (cryptographic) mechanisms.
- No IP (internet protocol) for IoT-focussed networks due to the overhead of the IP protocol.
- Connectivity to the internet is not always the norm. This impacts the security solution in that authentication and authorisation servers, and CA (Certification Authority) servers may not be available.

The IoTSF Assurance Framework (IoT Security Foundation, 2021) and the supporting Questionnaire guides stakeholders through the requirements of an IoT RHM device.

Data integrity, non-repudiation, authentication, authorisation, device integrity validation, access control, identity management and data confidentiality may all be enabled by strong device certificates and cryptography. Device certificates are discussed earlier in this section. Cryptography can be computationally intensive and so reduce battery life; however, many IoT microprocessors include cryptographic modules that can support these functions at a much-reduced energy cost. However, if each device has its own unique private key, then the device has to be programmed during manufacture of the microprocessor with that key and then claimed by the sensor manufacture once it is

assembled into a product. This requires a process for secure onboarding and offboarding of devices and owners, which can add more costs than can typically be absorbed into low-cost, consumer IoT gadgets. However, for an RHM sensor it would be difficult to justify not including these capabilities in order to protect both the integrity and confidentiality of the PGHD.

Being able to update the firmware on the device whilst it is still operating usually requires a second area of memory into which the new image can be installed and then device needs to failover to this new firmware image. If errors occur then the device may potentially be disabled (bricked), sometimes requiring return to the manufacturer or even disposal. Therefore, any image will require a hash check to demonstrate that the file has not been corrupted and it needs to be digitally signed to prove its origin.

3.3.5 Range, Coverage and Licencing

Range will impact the number of gateways required to provide coverage for a community and possibly the power consumption of the user devices. Unlike many other IoT deployments where the end-device is generally static, in a healthcare system the end-user can move considerable distances within their community. The maximum density of users may also necessitate multiple gateway devices to cope with the overall traffic flows. A related aspect is that of licensing of spectrum and the freedom for a community to establish their own infrastructure. For mobile phone technology, the applicable slices of frequency spectrum are generally auctioned by the government to mobile network operators. This usually confers a licence to deploy base stations operating at the assigned frequency to provide coverage within that country and forbids other entities from also operating within that spectrum. The result is that unless there is a mandatory condition attached to the licence, the coverage of the mobile network is determined by the economics for the licence holder of installing and operating an expensive base station against the projected revenue from users within that cell. Therefore, sparsely populated hilly areas that would require a high number of base stations per user are often not worth network operators providing coverage. Even in these situations, a local community would not be permitted to install their own communications base station operating in the same spectrum. However, IoT communications provides the opportunity for local communities to operate a network of base stations operating in the regulated yet unlicensed spectrum,

albeit at relatively very low data rates, too low for voice calls, and with constrained periods of operation (typically 1% of the time).

3.3.6 Reliability of communications

Reliability of the communications channel may be required, including assured delivery of messages, known freshness and error correction. Depending upon the system, it may be necessary for bi-directional communications, such as to provide reliable delivery, request different monitoring parameters or rates, patching software and renewing security associations (including revoking the device).

The bit error rate (BER) is a measure of the number of bit errors divided by the total number of bits communicated over a given time. Error-checking and error-correction techniques can be used to identify or correct errors in data, but they do incur an overhead of requiring extra bits to be transmitted, which effectively either reduces the useful information flow or requires an increase in data rates. The criticalness of the data will determine the BER required of a health system (Latre et al., 2011). For a remote health monitoring system, there may be a need to judge whether outlier data is due to communication errors or represents a real medical event.

A complication with body-worn communications is the absorption of signals by the body, and that this attenuation can vary significantly depending upon placement, movement of the limbs and orientation to the distant receiver.

A further complication for an RHM service provider is the heterogeneity of devices and protocols that can complicate interoperability. Using established standards, rather than proprietary protocols, can reduce the risk of the service infrastructure and devices being left stranded should the manufacturer cease to support the protocol.

3.3.7 Geolocation

Locating a user is often an important parameter for remote monitoring. Different methods for geo-locating a sensor are available (triangulation from base stations; GNSS - Global Navigation Satellite Systems) with co-ordinates can be provided in 9 bytes (GPS format

of 3 bytes each for latitude, longitude and altitude) or 7 bytes (AIS reports 28-bit longitude, 27-bit latitude) (United States Coast Guard, 2017). GNSS systems can be more power and processing hungry than geo-location services and deep indoor locations are difficult for GNSS signals to penetrate, but triangulation positioning often requires a database of towers to be stored on the device or for data to be sent to and from the device).

3.4 LPWAN communications options

3.4.1 Overview of Low-power wide area networks

Low-power wide area networks (LPWAN) potentially enable IoT devices to communicate over long distances (several kms) with low or no networking charges and at very low energy (e.g.: can run off an AA battery for several years). LPWAN network base-stations with a range of 10-15km, cost from around £400 (The Things Network, 2020) and often do not require a licence. A few base stations within a rural community could service dwellings over a wide area. The base stations could alert local carers and first responders and could backhaul to regional services either by landline (where available) or a network of base stations (although aggregating messages on LPWAN systems can lead to overload).

As identified in earlier sections, considerations in choosing an LPWAN will include:

- Overall capital and operating costs, including any licences and base stations.
- Range, and how it deals with hills or buildings blocking line of sight.
- Reliability, to assure important messages are delivered.
- Data rate for individual devices, constraints on the duty cycle of transmitting, and the overall system capacity.
- Energy consumption, to avoid having to recharge or replace batteries in user devices.
- Security to protect PHI in transit.
- Open standard versus proprietary solutions to avoid being locked into a vendor and potentially losing the investment should the vendor discontinue the system. An open system may also enable a community more freedom to deploy their own system rather than relying on having to buy a service.

RHM services need to plan for LPWAN communications being less reliable than those used in the home, such as Bluetooth and Wi-Fi. For example, the Bluetooth radio link is very robust, with spread spectrum digital codification in which the signal is distributed or expanded. The result is a more robust signal, which is less capable of deterioration because of electromagnetic noises and other sources of interferences. Additionally, frequency hopping makes the wireless transmissions more secure against interception (Bellido-Outeirino et al., 2008).

Table 3-3 provides an indicative performance of the leading current and planned LPWAN technologies. Compared to cellular technologies, LPWANs use more sensitive receivers to enable greater range (up to 40dBm, or 10,000 times more sensitive) (Link Labs, 2016). However, as with many networking protocols, there are trade-offs between range, data rate and energy consumption.

Table 3-3: Indicative LPWAN performance

LPWAN	Range (km)	Two-way	Data rate	Packet size (Bytes)
NB IoT	15	Y	60 - 250 kbps (30-170 kbps down)	Variable
EC-GSM-IoT	15	Y	70 – 240 kbps	Variable
eMTC (LTE Cat-M1)	15	Y	375 kbps (<1Mbps)	Variable
LoRaWAN	5 - 15	Y	250 bps to 50 kbps	59 to 250 B
Symphony Link (LoRa)	10	Y	Similar to LoRaWAN	256 B
MIOTY	5-15	Y	512 bps	10 to 192 B
Weightless -W	~ 2	Y	200 bps – 1 Mbps	10 B +
Dash 7	5	Y	167 kbps	256 B
SigFox	10	4/day	100 - 600 bps	12 (max 14 packets per day)
nWave		N	100 bps	2 – 20 B
Ingenu RPMA	5 – 6 (>500 km star)	Y	624 kbps up 156 kbps down	Variable up to 10Kb

3.4.2 IoT Security Protocols

Encryption of data on the end user device, as recommended by HL7, has considerable implications for an IoT system, including computation resources and key management where many heterogenous, remote devices may be joining or leaving the service. Identity management and non-repudiation may require devices to be enrolled into a service and the use of digital signatures and a certificate authority.

Traditional security and privacy algorithms are generally not constrained by processing power, memory or battery energy. When these are transferred into IoT systems the performance is often not acceptable in terms of response speed, reliability and robustness to loss of nodes or communication errors. Instead, a series of new protocols have been developed for IoT which are more suited to constrained devices. The security solutions in IoT (Malina et al., 2016) have to provide authentication and authorization of nodes (things, users, servers, objects) and data authenticity, confidentiality, integrity and freshness. The security solutions are usually implemented at network, transport and application layers in IoT. These include:

- Constrained Application Protocol (CoAP).
- IPSec.
- Host Identity Protocol (HIP).
- Transport Layer Security (TLS) protocol and Datagram Transport Layer Security (DTLS) protocol.
- Slim Extensible Authentication Protocol Over Local Area Networks (SEAPOL).
- Trust Extension Protocol for Authentication of New deployed Objects and sensors through the Manufacturer (TEPANOM).

A further complication is the distribution of secret cryptographic keys to users and devices that are geographically separated and may not be physically present to be authenticated when enrolled or registered onto a system. Revocation of a user or a device is also important, for example when a patient moves from a doctor or care provider, or when a device is lost, discarded or temporarily loaned. Several communication and authentication protocols are assessed by Malina et al (Malina et al., 2016) with a main finding that the choice of security scheme has to consider the constraints of the IoT nodes, which will impact other factors such as cost, battery life and the need for user interaction.

For example, asymmetric encryption may not introduce noticeable delays when implemented on a smartphone but is an order of magnitude worse when implemented on a constrained device. For very low-cost IoT devices, such as microcontrollers and smart cards, symmetric encryption and hashing can be performed in only a few milliseconds provided that the choice of algorithm is carefully chosen, such as AES-128b or RND 160b/ RND 560b random number generation, and there is sufficient RAM.

Other energy-saving techniques have been demonstrated using AES-128 and SHA-2 algorithms, such as encompression (encryption + compression) (Zhang et al., 2013). The authors claimed a reduction of up to 78% was achieved by increasing the compression ratio and can even be lower than an uncompressed system without any encryption.

3.4.3 Licenced Spectrum - 3GPP Protocols

The first three systems (NB-IoT, EC-GSM-IoT, eMTC/LTE CatM1) are 3GPP (3rd Generation Partnership Project) '5G' standards that operate from existing mobile phone base stations using licensed spectrum. 5G networks started to be rolled out in 2019, although initial deployments have been in dense urban areas to serve more customers. Coverage may eventually be better than the existing 3G/4G networks but is still unlikely to be economically viable in many rural areas.

The Global System for Mobile Communications Association (GSMA) argue that the benefit of using licensed solutions include (GSMA, 2016):

- Across more than 400 individual members, the 3GPP standards stipulate a minimum level of performance, regardless of vendor. Standards also ensure interoperability across vendors and mobile operators.
- 3GPP standards benefit from economies of scale due to the large number of companies that implement these standards.

The LPWA standards address applications requiring low mobility and low levels of data transfer:

- Low power consumption (to the range of nanoamp) that enable devices to last for 10 years on a single charge
- Data transfer optimised for small, intermittent blocks of data

- Low device unit cost
- Simplified network topology and deployment
- Improved outdoor and indoor penetration coverage compared with existing wide area technologies
- Secured connectivity and strong authentication
- Network scalability for capacity upgrade.

Mobile network operators (MNOs) can provide a managed communications infrastructure requiring minimal input from the RHM service, that is secure, generally reliable, and able to scale as required.

However, MNOs do not tend to provide infrastructure in areas of low population as it is not cost-effective (see section 2.3.10). The costs may also be prohibitive for some users in LMICs. Also, after a natural disaster, the infrastructure may be damaged and the local community may be without any communications until the MNO repairs the damage, which could be several weeks or months depending upon the priority of the community.

3.4.3.1 SMS Messaging

In areas where there is mobile network coverage, a cost-effective means of data transfer can be the SMS (short message service). SMS may also be appropriate in LMICs where a large percentage of the population have access to a basic 2G phone, even if they cannot regularly access a smartphone with internet service.

An SMS message is limited to 160 seven-bit characters (128 available letters, numbers and symbols), although multiple messages can be used to send parts of a large data transfer. The CellCheck system prototype demonstrated the use of a blood pressure monitor and pulse oximeter being collected on a cellphone via Bluetooth, and then the data is uploaded to carers. Using a medical record of around 100 bytes (patient's alias, and single pulse oximeter blood pressure readings, collection time and location), the end-to-end upload delay would be between 5 to 35 seconds (Khazbak et al., 2017).

A further use of SMS messaging is for clinicians, midwives and carers to send pre-natal advice or reminders to principals to take medicines or perform some other activity.

However, mobile phone networks use the Signaling System No 7 (SS7) protocol for interconnection, and SMS messages have been intercepted. SMS messages are not encrypted end-to-end, and so PHI can be vulnerable when transiting the core mobile network.

3.4.3.2 NB-IoT

NB-IoT (Narrowband-IoT), or Cat-NB, was approved by the 3GPP in June 2016 and is being rolled out by major network operators (e.g.: Vodafone within the UK). With links of up to 250kbps, it aims to be more cost-effective than 3G/4G networks for M2M (machine-to-machine) applications. NB-IoT fits into 180KHz of spectrum and so operators have the option of deploying on a GSM carrier spectrum where 4G has not been rolled out (Vodafone, 2018).

3.4.3.3 EC-GSM-IoT and LTE-M

EC-GSM-IoT (Extended Coverage Global System for Mobile Communications Internet of Things) and LTE-M (eMTC and LTE CatM1) provide high speed communications over mobile network spectrum. They are likely to be deployed only by existing carriers.

3.5 Unlicenced Spectrum - ISM Band

The following technologies use unlicenced spectrum, often in the ISM bands (Industrial, Scientific and Medical), as defined by the International Telecommunication Union (ITU, 2020c). Communications devices are permitted to operate in some of the ISM frequencies, subject to regulation and constraints, and having to accept the possible interference from other ISM devices. The frequencies and limitations on use are set by regional bodies – for example, ETSI in Europe (ETSI, 2018a) and the FCC in the USA under the Electronic Code of Federal Regulations, Title 47 Part 15 (Subpart 15.247) (FCC, 2020b). The constraints include parameters such as maximum radiated power, maximum duty cycle (<1% for ETSI), channel bandwidth. A growing problem with devices operating in the ISM bands is that as the density of devices becomes greater, then there will be increasing congestion due to collisions in messages, which will impact the quality of service. These ISM protocols, together with several more technologies, are all

vying for a growing, but as yet immature, market. It is probable that the market will coalesce around a few technologies for widespread adoption, with other technologies being adopted for niche application or industries and others withering away.

3.5.1.1 LoRaWAN

LoRaWAN (Long Range Wide Area Network) is promoted by the LoRa Alliance, but the LoRa radio modulation technology used within the transceiver chip is proprietary to Semtech. LoRa operates in the unlicensed ISM spectrum, is bi-directional, with data rates of up to 50 kbps. LoRaWAN aims for devices to operate for several years on a single battery, dependent upon the messaging rate and distances. Typical configuration is in a star-of-stars network to relay messages to a network server (IETF, 2018b) and offers good scalability (Bor et al., 2016).

There is a growing number of LoRaWAN networks, including The Things Network (TTN) with over 16,000 gateways globally and over 1,000 in the UK as of January 2021, with 17 gateways around Reading (The Things Network). One potentially attractive feature of LoRaWAN is that end nodes and base stations are available from several companies, costing from £400 (or can be built upon a Raspberry Pi for around £200) and so may be affordable for a community to create its own network in remote regions and in developing countries.

3.5.1.2 Symphony Link

Link Lab's Symphony Link is a proprietary system operating in the ISM band that uses the LoRa standard chipsets for the PHY (physical layer) but not the LoRaWAN MAC (media access control). It has an adaptive data rate, supports packet acknowledgement and provides privacy using AES and TLS encryption. It is more expensive than LoRaWAN products, but purports to offer improved performance, which could be suitable for a service which requires high quality of service.

3.5.1.3 MIOTY

MIOTY is a relatively recent IoT communications network also operating in the sub-1GHz ISM bands and based on the ETSI TS 103 357 Telegram Splitting Ultra Narrow Band (TS-UNB) family (ETSI, 2018b). MIOTY uses Fraunhofer's patented telegram splitting technology and is aimed at large industrial and commercial networks. The

MIOTY Alliance (https://mioty-alliance.com/) includes Fraunhofer, Diehl, BehrTech, Texas Instruments and several other companies. It is claimed that MIOTY performs better than LoRaWAN, as measured under the IEEE 802.15 LPWAN PHY Interference Model, when dense or co-existing IoT networks result in increasing packet error rate due to more packets colliding. The data rate is only 512 bit/s, and the data payload can be 10 to 192 bytes. Each message should only take 17.8 µWh to transmit, so with battery recovery periods, the battery life could reach 20 years. It also allows for up to 1M devices per network and each gateway can handle up to 1.5M messages per day. The range is similar to LoRaWAN, up to 15 km (5km more typical for urban and non-line of sight) and nodes and base stations can operate at speeds of up to 120km/h.

3.5.1.4 Weightless

Weightless is an open standard with three sub-standards; -N is unidirectional, -P and -W are both bi-directional. -N and -P operate in the sub-1GHz unlicensed spectrum, whilst -W operates in the TV spectrum, but has higher power consumption. Range is lower than other technologies above and so may not be suitable for dispersed communities.

3.5.1.5 Dash7

Dash7, derived from ISO18000-7, is a proprietary, open-source system. Originating from military logistics and RFID tags, it is suited for BLAST traffic (bursty, light, asynchronous, stealth and transitive). Due to limited range its use for community services is limited.

3.5.1.6 SigFox

SigFox is a proprietary standard and is widely deployed in Europe. It is connectionless, optimised for uplink communications and in ideal rural environments, the range can be over 30km. However, each node is limited to uplinking 150 messages of 12 bytes per day. The downlink channel is even more constrained to four messages of 8 bytes per day. In areas served by SigFox, it could provide simple health status, but it would be difficult to provide a reliable health service based on SigFox or a service requiring two-way messaging.

3.5.1.7 Nwave

Nwave's marketing is focussed on parking solutions, claiming to have twice the range and an order of magnitude lower power compared to LoRa. Similar to SigFox, it is too constrained to be useful for healthcare.

3.5.1.8 Ingenu RPMA

Ingenu RPMA (random phase multiple access) is a proprietary standard with a range of up to 50 km line of sight or 5-10 km without line of sight. Message acknowledgement improves reliability and helps to fulfil duty-of-care requirements. The number of access points required to serve an area is purported to be significantly lower according to Ingenu, although its energy consumption may be higher. However, roll-out appears to depend upon Ingenu and its partners, which may prevent community-led deployments. Its initial roll-out is to cover rural Texas oilfields and it is possible that it will focus on high value industrial segments rather than target the consumer market.

3.6 Satellite IoT communications

Satellite-based IoT communications shows great potential, with over 18 different constellations at varying stages of development. However, many of these proposed systems may not achieve operations and of those that do, the pricing structure may be prohibitive for all but the most affluent of rural residents. Power consumption may also be problematic. As an example, the US FCC rejected a Starlink bid for Rural Digital Opportunity Fund Subsidies because the user would be required to purchase a $600 dish to obtain high-speed broadband (FCC, 2022). Costs for Starlink high-speed broadband internet was £75 per month plus £460 for hardware as of September 2022 and average power usage is 50-75W (https://www.starlink.com/), so too expensive and power intensive for the use cases in this research.

There are at least four design considerations for using satellite IoT communications for remote health monitoring:

- Cost of service – although service costs have not been announced, they are expected to be higher than their terrestrial equivalent due to the costs of launching a constellation of satellites.

- Reliability and revisit times
- Power consumption
- Antenna design

The first three factors are likely to result in hybrid solutions for the foreseeable future, where the system uses terrestrial networks where available, switching to satellite communications when necessary to send (or receive) priority messages.

As an example of the potential services offered by satellites, Kineis enables 30-Byte messages to be transmitted with an average revisit time of 15 minutes (planned for 2023 onwards). For very small messages (7 Bytes), transmission is at a very low data rate of 200bps with transmission power of 100mW from a 7mm x 7mm modem. Kineis suggest that a 3000mAh battery would support two 30B messages per day for over 10 years or one 30B message per hour for over 1 year, although other parts of the system would lower the overall battery life of the system. Future modulations and a new constellation 'will significantly lower' power consumption (Kineis, 2022).

As of September 2022, Keneis and Lacuna Space are the two European companies that are operating in-orbit satellites for testing and development of IoT satellite connectivity in sub-GHz bands.

A second example is FOSSA who build and operate picosatellites (which weigh 0.1 to 1kg) capable of receiving LoRa compatible transmissions of 150 bytes at a data rate of 300bps. Existing LoRa sensors can operate a 2-way connection to their satellites with minor antenna and firmware modifications with 25mW of transmission power. Revisit time is approximately 24 hours with a single satellite, reducing to 5 minutes with a constellation of 60 satellites.

Other recent entrants to the satellite IoT communications market include:

- Wyld Networks Connect
- EchoStar Mobile
- Astrocast

- Swarm, owned by SpaceX, operates 120-150 nanosatellites (1 to 10kg), provides low-bandwidth (1 kbps one-way) connectivity for $5/month and sells an asset tracker for $89 (https://swarm.space/).

3.7 Sensor Edge Processing

Many AAL systems make use of several sensors around the home with the data being sent to a local hub, or smartphone, for processing and aggregation. This has several advantages such as:

- Simpler, cheaper devices can be used as sensors as they do not need as much processing capability or storage.
- Battery-powered end-devices being able to use short-range, energy-efficient communications, so prolonging battery life.
- Greater privacy in that much of the pattern-of-life data can be analysed locally, rather than sent to the cloud for processing.
- Potentially lower communications costs and continued operation in the event that communications with the cloud are interrupted. When using low-bandwidth communications, such as LPWANs, it may be necessary to filter or prioritise the data sent to the cloud due to network constraints. Reduced latency is also often cited as a benefit, although many applications can tolerate a sub-second response and local constrained devices may introduce their own latency in waking-up from sleep mode and slower processing times. GSMA (GSMA, 2016) estimate that about 30 seconds is the tolerated latency for people interacting with tracking devices, although 2-5 seconds may be required in some cases. In the overall response times of a carer being able to reach a user, such delays are a very small fraction, but a rapid indication that the service has been notified and is responding to an event may be very reassuring to a user in urgent need.

Edge computing means different things across industry specialisations and there is no consensus (Iorga et al., 2018). Following their nomenclature, an AAL system based in a smart home could be described as mist or fog computing (see Figure 3-5), which is distinct from edge computing that encompasses end devices and their users and is often referred to as the IoT network.

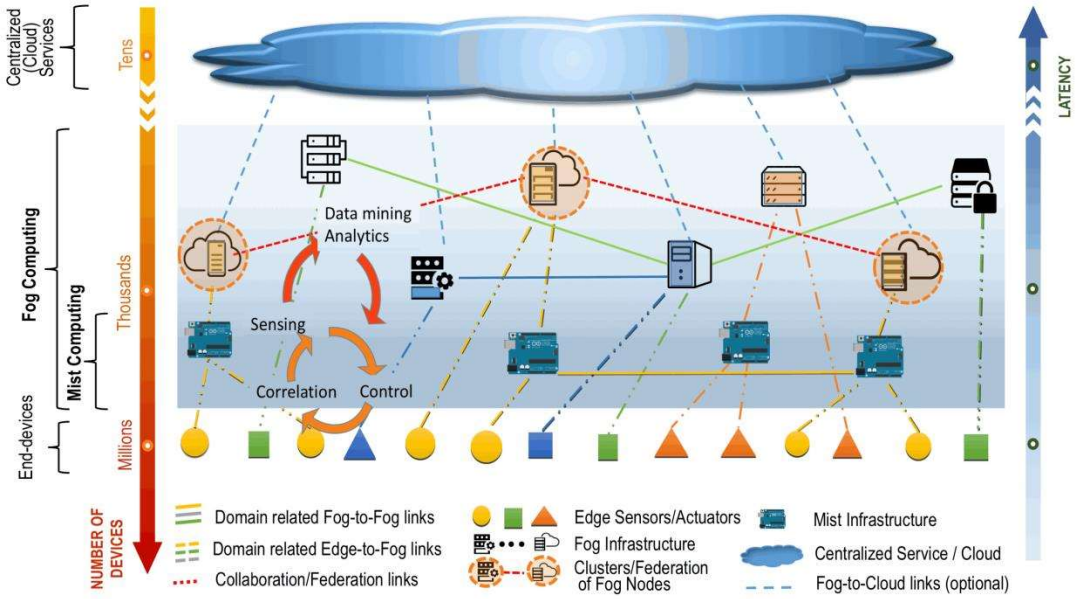

Figure 3-5 – Context of a cloud-based ecosystem for smart end-devices (from NIST SP 500-325 Fig 1)

A constraint on using a hub within a smart home is that it is generally not mobile, and so cannot support the user when they are out of range. A modern smartphone does offer mobility and may have the computing and communications performance to act as a hub (or mobile edge computer) if the user and provider can cope with the limitations expressed earlier, such as regularly recharging the battery, affordability of the device and communications, and the user interface.

One method of 'untethering' health devices from the mist/fog hubs is to perform more processing on the edge devices (sensors) and only sending important information to the central processing system (cloud). This reduces the energy used for communications and can improve privacy as much less raw data leaves the sensor. However, historically sensors have had very limited computing power and so could only perform very specific tasks.

3.8 Embedded Machine Learning

Edge Machine Learning is rapidly progressing with a growing number of devices being optimised to run ML algorithms on highly constrained processors (Murshed et al., 2019). Three of the key benefits are:

- The energy required to transmit all data to a remote data centre can be greater than the total energy required to use ML models on the raw data at the edge device and only transmitting the key information to remote users (Gómez-Carmona et al., 2020, Suresh et al., 2018). The determination of what is 'key information' will depend upon the application, but often includes anomalous or outlier data and any trends away from the average.
- To reduce network congestion, which may become more important as the number of devices using the unlicensed ISM bands is expected to rise rapidly (Merenda et al., 2020).
- To support security and privacy in IoT data (Amiri-Zarandi et al., 2020, Tahsien et al., 2020, Xiao et al., 2018, Sun et al., 2020).

3.8.1 Benefits and Examples of Embedded Machine Learning

The use of ML in healthcare is expected to bring improvements such as detecting falls (Torti et al., 2018, Tsinganos and Skodras, 2018), responding to heart rate variability, improving quality of service and being able to scale healthcare to many more users where resources (professionals) are constrained (Qadri et al., 2020). One example is the SPHERE project (SPHERE, 2018), where it was determined that the sensor battery lifetime could be extended with the use of embedded machine learning. However, in the SPHERE smart-home environment, the benefits of having the richness of a complete dataset outweighed the sparser information obtained through embedded machine learning (Fafoutis et al., 2018).

A review of deep learning techniques, in particular recurrent neural networks, to detect falls indicates a higher accuracy and adaptability to a wider range of situations can be achieved compared to threshold-based fall detection or SVM classification (Queralta et al., 2019). Further, deep learning can reduce alert latency compared to needing to transmit full sequences of raw data to be processed in the cloud. The use cases proposed by Queralta et al. uses sensor nodes which can collect health data such as electroencephalography (EEG) electrocardiography (ECG), electromyography (EMG), and blood pressure, together with contextual data including temperature, humidity, and air quality. The data is sent by BLE to an edge gateway for processing and then transmitted over LoRa to the cloud. Using a neural network with three hidden layers and

two dropout stages, a precision of 90.1% (+- 3%) and a recall of 95.3% (+- 0.8%) was achieved. Significantly, this resulted in only 10 bytes of data needing to be transmitted from the edge gateway to indicate detection of a fall, so reducing potential congestion in the network and battery usage. Testing was conducted in an urban environment with the LoRa access point over 4km away over a hill.

3.8.2 Guidance on ML Practice and AI Risk Management

Embedded ML makes it more complicated to assess the reliability of a service because the response of the system is less predictable as it is dependent upon the inferences of the algorithms and training data sets (Akmandor and Jha, 2018).

In October 2021, the U.S. Food and Drug Administration (FDA), Health Canada, and the United Kingdom's Medicines and Healthcare products Regulatory Agency (MHRA) jointly issued 10 guiding principles to inform the development of Good Machine Learning Practice (GMLP). The aim is to *promote safe, effective, and high-quality medical devices that use artificial intelligence and machine learning (AI/ML*) (FDA et al., 2021). The 10 Guiding Principles are:

- Multi-Disciplinary Expertise Is Leveraged Throughout the Total Product Life Cycle
- Good Software Engineering and Security Practices Are Implemented
- Clinical Study Participants and Data Sets Are Representative of the Intended Patient Population
- Training Data Sets Are Independent of Test Sets
- Selected Reference Datasets Are Based Upon Best Available Methods
- Model Design Is Tailored to the Available Data and Reflects the Intended Use of the Device
- Focus Is Placed on the Performance of the Human-AI Team
- Testing Demonstrates Device Performance During Clinically Relevant Conditions
- Users Are Provided Clear, Essential Information
- Deployed Models Are Monitored for Performance and Re-training Risks Are Managed

A recent document that may become influential is NIST's Artificial Intelligence Risk Management Framework (NIST, 2023). It identifies the characteristics of trustworthy AI systems

Figure 3-6: Characteristics of trustworthy AI systems (from (NIST, 2023))

Chapter 4 Opportunities from the Literature Review

Three major aspects are evident from the user and technical requirements of using IoT devices for RHM, and they are:

1. Integrity and reliability of messages being received by the core service so that the correct response can be actioned (see Sections 2.3.3 and 3.3.6).
2. Privacy and security of data, especially for PHI (see Sections 2.3.4 and 3.2).
3. The need to minimise energy consumption in order to extend the operation of battery-operated devices for a year or longer, so reducing reliance on users or maintainers to recharge or replace the battery (see Sections 2.3.5).

Machine Learning (ML) on the user's device has the potential to support all three of these aspects. The widespread growth of TensorFlow Lite for Microcontrollers and TinyML is resulting in '*technologies and applications including hardware, algorithms and software capable of performing on-device sensor data analytics at extremely low power, typically in the mW range and below'* (https://www.tinyml.org/).

RHM systems typically operate on very limited resources in terms of energy, memory and processing power. IoT devices are now available that are very energy efficient across each aspect of the device, from the sensing elements, duty-cycling low power networks, energy-efficient security, and low-power operating systems. However, a sensing system is only as efficient as its least efficient subsystem, and the efficiencies of the physical elements would not achieve overall system efficiency unless the principles of resource-efficient design are also adopted in the data layers. Indeed, despite energy-efficient sensing elements, the system energy requirements will also depend on the amount of data to be handled. Any data that is generated, transferred, stored, or processed unnecessarily is a potential waste of precious energy (Fafoutis et al., 2018).

A key question is whether the trade-off between increasing data processing on the edge device in order to reduce the overall amount of data transmitted actually results in energy savings. Other considerations include:

- Security and privacy benefits of processing raw data on the device, and only transmitting the outcomes to the gateway and core databases.
- Assurance of the ML algorithms that they would detect any anomalous readings
- The impact to clinicians and researchers in the loss of raw data that may provide a greater insight where a user has multiple conditions and where the edge processing only addresses the primary symptoms of the user.

4.1 Context - SPHERE Wearable

The SPHERE wearable will be the subject of this research. SPHERE is a Sensor Platform for Healthcare in a Residential Environment, designed for a user in a smart-home that has been provisioned with several sensors and Bluetooth (https://www.bristol.ac.uk/engineering/research/digital-health/research/sphere/).

Although SPHERE is a smart-home project, rather than for users roaming in the environment, this device was chosen as it has been researched extensively within the academic community (Fafoutis et al., 2017b, Beach et al., University of Reading, 2018, Ghamari et al., 2016, Sherratt and Dey, 2020). Measuring the energy consumption of Bluetooth Low Energy (BLE) messages also provides more consistent results compared to LPWAN communications which may change their power level and data rate to accommodate changes in distances and received power from remote gateways.

4.2 Energy Comparison of Embedded ML Processing to Raw Data Transmission

A key paper for this section is (Fafoutis et al., 2018), which assessed how embedded ML could be used to extend the battery lifetime of the SPHERE wearable sensor. Two approaches to processing data were compared:

1. All data generated is transmitted via BLE to a server for processing and storage.
2. Embedded ML extracts key information (knowledge) on-board the device, and only extracted data is transmitted to the core.

The two concepts are illustrated as Fig. 1 in the paper, reproduced here as Figure 4-1:

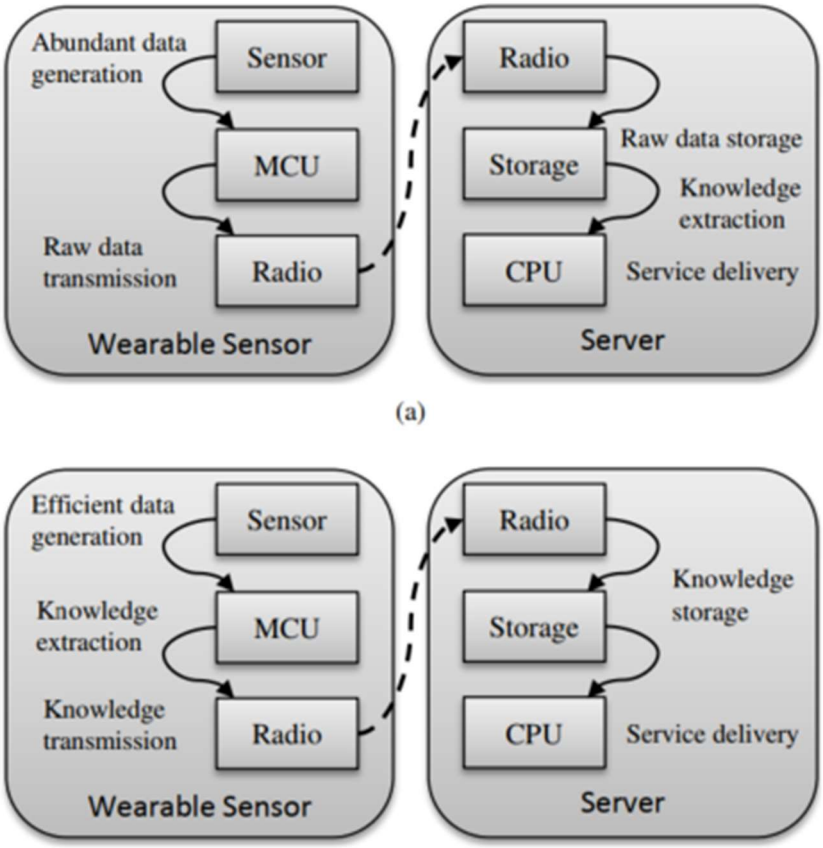

Figure 4-1: Comparison of raw data approach to embedded ML (from (Fafoutis et al., 2018))

Embedded ML has been demonstrated in many different health care research environments (see (Greco et al., 2020) for an extensive survey), including detection of atrial fibrillation (Mcdonagh et al., 2022), arrhythmia in ECG signals (Yang et al., 2022), cardiovascular disease (Ukil et al., 2021), myocardial infarction (Sopic et al., 2018a, Uchiyama et al., 2022), diabetes prediction (Ramesh et al., 2021), and detection of epileptic seizures (Sopic et al., 2018b).

Fafoutis et al. took a full-system perspective when assessing the benefits of embedded ML, and specifically using ML as a means to reduce the radio duty cycle of the processor and the radio, which are typically the two most energy consuming components in a low-power IoT device. Using previous energy measurements (Fafoutis et al., 2017a) and the datasheet of the accelerometer, an estimate of the energy used to perform additional processor cycles to represent embedded ML was compared to the energy required to send the whole data over BLE.

In its basic setup, the SPHERE wearable device was used to classify the wearer's activity level (rigorous - running, exercising; moderate - house cleaning, walking; sedentary – typing, watching TV). Feature extraction and classification used the Integral of Modulus of Acceleration (IMA) and acceleration measurements in each of the three axes. A Support Vector Machine (SVM) was used to classify the IMA results to the three levels of activity, with an average classification accuracy of 93.2% and standard deviation of 6.6%. Assuming the wearable battery provided 1000J, then the energy consumption corresponded to a battery lifetime of approximately 13 days.

A strategy to reduce energy consumption was to reduce the sampling frequency and the sampling resolution of the accelerometer. The nominal sampling frequency is 50Hz and the ADXL362 accelerometer has a fixed resolution of 12 bits, which are stored in the MCU by two 8-bit registers. By reading and transferring the most significant byte only, the energy required to transfer the data to the MCU can be halved. Analysis indicated that a classification accuracy of > 90% could be achieved at a sampling frequency of 0.39 Hz and a bit resolution of 5 bits. However, the ADXL362 lowest frequency is 12.5 Hz, and one byte (8 bits) is the smallest resolution that can be transferred to the MCU, but the MCU was configured to only poll the MCU at 0.39Hz. This configuration of optimised sampling predicted a significant (3 orders of magnitude) reduction in the energy required to transfer samples with the board. The battery lifetime was predicted to increase from 13 days to 771 days.

The optimised sampling configuration then became the baseline for assessing whether to perform embedded feature extraction and embedded classification.

Embedded feature extraction requires additional processing cycles but reduces the amount of data to be transmitted and hence the radio duty cycle. The balance depends on the specific use case. In the configuration outlined above, using embedded feature extraction extended the estimated battery lifetime from 771 days to 989 days.

Embedded classification was determined by testing the extracted feature against the IMA two SVM thresholds that separate sedentary from moderate, and moderate from vigorous.

The result is encoded as two bits, reducing the potential data transmission by a factor of 8, corresponding to an estimated battery lifetime of 997 days.

However, although the power consumption estimates show that battery lifetime can be extended considerably, the limiting factor may become the current draw while the device is in idle mode.

4.3 Privacy Implications

Transmitting only the embedded classification removes specific data regarding the specific activities that the wearer was engaged in.

What may be more important for the carer or supervisor is whether the class of activity fits in with the principal's expected pattern of life. For example, sedentary behaviour in the afternoon may be normal for an elderly wearer, but abnormal for a remote worker or rescue team. Similarly, vigorous activity in the middle of the night is unusual for most elderly wearers (but not for some individuals) but may be expected for shift workers.

Another example of preserving privacy could be classifying whether a wearer is in their expected location, or if they are wandering. By configuring the device with an expected geo-boundary for a specific user, the device could determine if the user is within the area, within a short distance, or is moving away. This could alert the carer to a potential problem, but any malicious actor intercepting the transmissions would not receive any location data.

4.4 Reliability of Classification

A weakness with ML is that it can give incorrect results when faced with situations that it has not encountered during training. For example, an embedded classification function may fail to recognise anomalous behaviour on which it has not been trained, such as a wearer falling or having a seizure being classified as a vigorous activity. It is not expected that a carer would be able to decode accelerometer data to determine the correct behaviour, but transmitting all the data to the core server allows three further improvements in assessment:

1. A greater number of models can be performed against the user's data. The constrained wearable may only have the memory and processing power to run one specialised classification model against the data, whereas servers in the core could assess the wearer's data against a wide range of potential conditions. This may improve the service's reliability in identifying an event where the principal needs support.
2. A constrained wearable device may have limited long-term memory or may not have the processing power to compare current readings to previous data. The core servers may retain historical data (subject to GDPR retention constraints) against which to compare current events. If previous events have been annotated with the impact or a clinical diagnosis, then the RHM service would be in a better position to reliably determine the severity of the wearer's current condition.
3. The core system can learn across its population of users in order to continually evolve and refine its models. This is not available to an embedded classifier without frequently downloading new models, which may be energy intensive and could interrupt the real-time monitoring or require user intervention.

Learning across the population of users is very important in researching health conditions, but data privacy and (pseudo-)anonymity must be carefully controlled.

Chapter 5 Comparison of Energy Consumption for Edge Processing and Transmitting Data

5.1 Introduction

As discussed in Chapter 4, Fafoutis et al. (Fafoutis et al., 2018) showed theoretically that the battery lifetime of the SPHERE smart-home wearable could be extended from a nominal 13 days to 997 days by the use of reduced sampling and embedded feature extraction and classification. This research seeks to practically validate the results published in the above paper.

The novel extension of this research compared to the original paper is to investigate the energy consumption of encrypting the BLE messages with AES-128.

Instead of performing embedded ML processing on the device, which may yield variable results, a proxy method is used of configuring the device to run a defined number of processor cycles (incrementing a counter). This allows for incremental increases in the processor duty cycles and yields repeatable results.

A Texas Instruments (TI) TI CC2652RB microcontroller (Texas Instruments, 2021) with Bluetooth LE (BLE) communications is used as the subject device instead of the nRF51822 system-on-a-chip used in the SPW-1, the first iteration of the SPHERE Wearable, which was investigated detailed by Fafoutis et al. (Fafoutis et al., 2017a). The TI datasheet states that the CC2652RB uses a 48 MHz Arm Cortex M4F processor and the board's power consumptions are 7.9 mA for active transmission mode at 0 dBm, sensor controller in low power mode of 30.8 µA, and a standby current draw of 0.94 µA.

An evolution of the SPW-1 software is used, namely SAPPHIRE3 developed for the Next Generation Wearable – Sapphire (University of Reading, 2021). The SPW-1 was based upon the Contiki RTOS (real-time operating system), but SAPPHIRE uses TI-RTOS-MCU (formerly named SYS/BIOS) within Code Composer Studio (CCS) IDE 10.1.0 which is wholly event driven with no 'main' programme running. An advantage of TI-RTOS-MCU is the advanced TI-RTOS Power Manager which, according to TI, "*provides pre-implemented, ultra-low power modes and can automatically determine the optimal*

low-power mode when the CPU becomes idle. TI-RTOS drivers are power-aware and communicate with the Power Manager to ensure peripherals are powered-down when not in use".

The TI CC2652RB board provides a True Random Number Generator (TRNG), AES 128- and 256-bit cryptographic accelerator, ECC and RSA public key hardware accelerator and SHA2 accelerator (up to SHA-512). The measurements will provide insight as to whether there is a significant energy cost to encryption, which may inform security and privacy approaches for smart wearables.

The CC2652RB block diagram is reproduced in Figure 5-1 and a photograph of the development kit is shown in Figure 5-2.

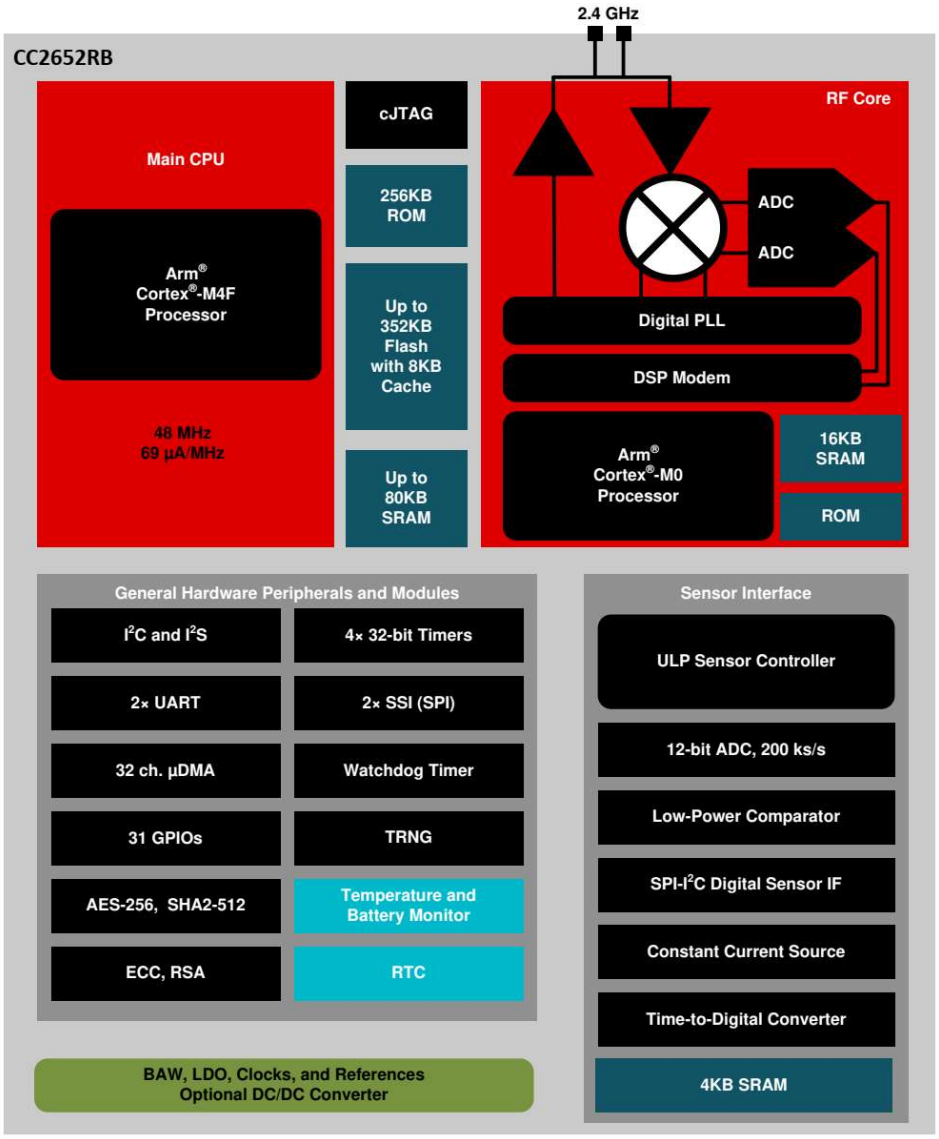

Figure 5-1: CC2652RB Block Diagram (from (Texas Instruments, 2021))

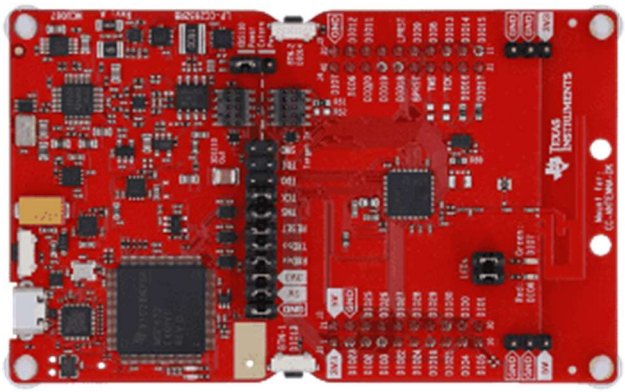

Figure 5-2: CC2652RB Development Kit (from https://www.ti.com/tool/LP-CC2652RB#tech-docs)

5.2 Experimental Setup

The methodology is based upon the approach described in (Fafoutis et al., 2017a), but using a TI LP-CC2652RB development board with a CC2652RB 32-bit Arm Cortex-M4F MCU. This provides accessible IO pin headers and removes any SPW-1 specific circuitry, so providing for more generally applicable results. Note that the results will not be directly comparable to the original paper, but it is the overall ratio of energy benefits for embedded processing that is being assessed.

The board is configured as a peripheral with BLE advertisements, and a Raspberry Pi 3 Model B v1.2 (RPi) is used as the host. Note that the RPi can only support Bluetooth 4.2, which has an advertising window of 31 bytes.

In this setup, the BLE connection interval, the time between two data transfer events, is available in 3 speeds of multiples of 1.25ms:

- Fast rate – 6 x 1.25 ms = 7.5 ms. Notifications packets of 247 bytes are sent at 100Hz, which may be required for multiple sensors and a fast IMU.
- Medium rate – 300 x 1.25 ms = 375 ms.
- Slow rate – 400 x 1.25 ms = 500 ms.

Note that the maximum interval could be 4 seconds. The slower the rate, the less energy will be expended in transmitting, but there will also be a lower maximum data rate. Note that the host can move to a slow rate if there is no IMU data. A timer triggers the LP-

CC2652RB peripheral to waken the CPU, data is transferred over BLE, and then the CPU returns to sleep mode.
When enabled, AES-128 encryption is used.

5.2.1 Hardware Setup

1. The CC2652RB GPIO DIO28 pin is connected by jumper cable to a 15.2Ω resistor in series to ground. This measures the "MCU_SPI1_CS_pin" (see software configuration).
2. A Textronix TDS1002 2-channel digital oscilloscope with TDS2CMA Communications Module was connected as follows:
   - Channel 1 measures the voltage across the 15.2 Ω resistor, from which the current drawn by the CC2652RB is calculated. The CH1 Volts/Div scale is typically set at 50mV.
   - Channel 2 measures voltage of the DIO28 pin relative to supply ground. This shows when the CPU is active (SUCK_DELAY = 1). The CH2 Volts/Div scale is typically set at 2V. The absolute voltage of DIO28 is not important; it is the rise and fall indicating the CPU activity that is being observed.
   - Timescale should be set to around 5 ms/division (between 2.5 to 10 ms/div), so as to capture 1-5 cycles of the CPU being activated.
   - The RS-232 connection allows the screen data to be sent to an external device, such as a computer. A RS-232 to USB convertor was used to capture the oscilloscope measurements.
   - The TDS1002 has a sample rate of 1.0 GSample/s (1µs interval).
3. The CC2652RB is powered by a 3.3V supply.
4. A Raspberry Pi (RPi) is used to configure the CC2652RB over BLE. In this experiment, the RPi was physically close to the CC2652RB so enabling BLE transmissions at very low power levels.

5.2.2 CC2652RB Configuration

In Sapphire3, GPIO DIO28 is called "MCU_SPI1_CS_pin". It indicates the status of the timer, SUCK_DELAY, which keeps the CPU active:

- DIO28 = 1; CPU SUCK_DELAY active
- DIO28 = 0; CPU SUCK_DELAY complete

SUCK_DELAY is a configurable parameter on "COMMAND_SUCK_DELAY=0xe9" and keeps the CPU active by cycling through the following "for-loop":

```
PIN_setOutputValue(Sapphire_GPIO_Handle, MCU_SPI1_CS_pin, 1);
      volatile int delay=0;
      volatile int i=0;
      for (delay=0; delay<suck_delay; delay++)
      {
        i=i+1;
      }
PIN_setOutputValue(Sapphire_GPIO_Handle, MCU_SPI1_CS_pin, 0);
```

The interval between the start of each SUCK_DELAY cycle is 20ms.

The connection interval is configured to fast (7.5 ms), but with a SLAVE_LATENCY = 3, which results in skipping 3 connections if there is no data to send, resulting in a CONNECT request being sent every 30 ms.

If there is no data to be sent for a prolonged period, the host can poll the peripheral and negotiate the slow connection interval (500 ms), but with the SLAVE_LATENCY = 3, may result in a connection request every 2 seconds.

5.2.3 RPi Commands

At the RPi terminal window the following commands are used:

RPi command	Comments
`sudo hcitool lescan`	Obtains MAC address of CC2652RB, e.g. 77:88:99:00:AA:BB.
`sudo gatttool -b [MAC_ADDRESS] – I`	gattool can be used to manipulate attributes with a BLE device, where -b is used to specify the remote BT MAC address, and -I specifies interactive mode

`Connect`	Configures the device into connection mode, the host and peripheral negotiate rates. Expected response is “connection successful”
`MTU 256`	Sets message Maximum Transmission Unit. Expected response “MTU set to 251”, with 247 bytes available due to addressing.
`char-write-req 0x23 0100`	Notify on CMD channel 22 (little-endian Byte; big-endian Bit)
`char-write-req 0x27 0100`	Notify on CMD channel 26
`char-write-req 0x2b 0100`	Write 247 bytes (flash updates)
`char-write-req 0x22 e60000`	Cryptography OFF
`char-write-req 0x22 e60100`	Cryptography ON
`char-write-req 0x22 e9000000xxxx`	xxxx denotes the suck_delay value in hexadecimal. e.g., 0500 = 800 counts

5.3 Methodology

The RPi host was used to command the CC2652RB peripheral SUCK_DELAY counter using the command:

```
char-write-req 0x22 e9000000xxxx
```

where xxxx denotes the value of SUCK_DELAY in hexadecimal.

SUCK_DELAY values were incrementally increased to keep the CPU active. The voltage measured over the resistor was captured by the oscilloscope (1Gsample/sec) and transferred to a laptop as 2,500 data points for analysis. A screen capture was also recorded to provide a visual representation of the duty cycle and overlapping BLE connections.

For each value for SUCK-DELAY, the encryption was toggled between ON and OFF in order to measure the additional energy required to encrypt the data, using the commands:

```
char-write-req 0x22 e60100      Encryption ON
char-write-req 0x22 e60000      Encryption OFF
```

The total energy used during a SUCK_DELAY cycle was calculated by summing the energy consumed during each individual data point, using the formula:

$$e = \sum (V^2/R * t)$$

where:

e = energy consumed for the cycle

V = Voltage measured on CH1 (note that measurements are in mV)

R = Resistance = 15.2 Ω

t = time interval of 1 sample = 1 µs

∑ = summation of each individual data point over the period of interest

The energy consumption was measured for both the processor cycles and also the energy used during a BLE connection. The idle current was extremely low, much lower than the processor activity, but was included in the power measurements.

The energy used during the period of time that the CPU was transitioning to active and then returning to sleep after a SUCK_DELAY cycle was included in the measurements.

5.4 Results

The detailed measurements are recorded in Appendix A, but a summary table is shown in Table 5-1 (cells marked N/R were not measured).

Table 5-1: Summary of Results

SUCK_DELAY (Hex) (CPU cycles)	SUCK_DELAY (Dec) (CPU Cycles)	Energy Consumption over 20 ms - Unencrypted (nJ)	Energy Consumption over 20 ms - Encrypted (nJ)
2000	8,192	509	507
3000	12,288	741	756
4000	16,384	982	965

5000	20,480	1,192	1,212
5500	21,760	1,269	1,263
6000	24,576	1,431	N/R
9000	36,864	2,127	2,137
9500	38,144	2,303	2,187
9a00	39,424	2,274	2,277
a000	40,960	2,359	N/R
b000	45,056	2,595	N/R

An example of the results are as follows which captures 100 ms of activity:

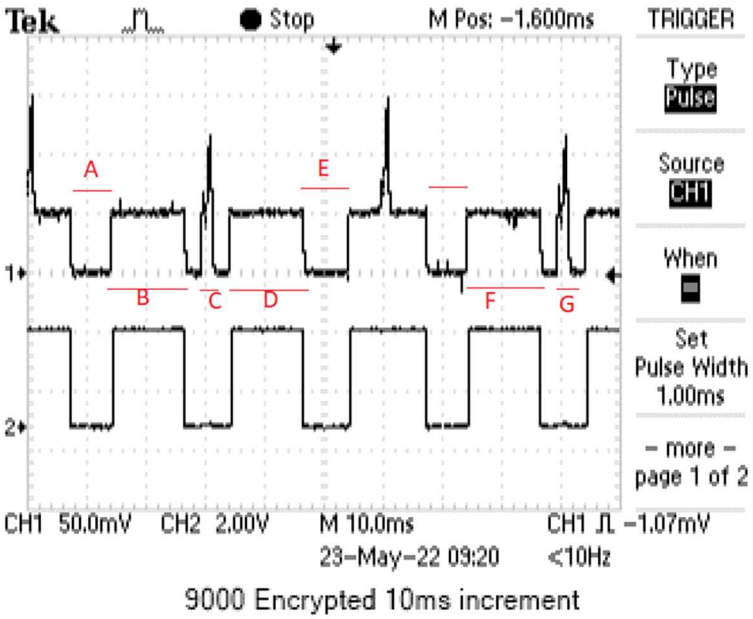

Figure 5-3: Example Waveform

The lower section of the oscilloscope screen is Channel 2 recording the status of GPIO DIO28, indicating when the SUCK_DELAY for-loop is active. As can be seen, this repeats every 20 ms.

The upper section of the screen is Channel 1 recording the voltage across the 15.2 Ω resistor, which represents the power consumption of the CC2652RB.

The square-wave patterns marked as B, D and F represent the power consumed when the CPU is active, corresponding to SUCK_DELAY = 1, but with additional time to transition from sleeping to active and also to transition from active to sleeping. The CPU is sleeping during the periods marked A and E.

The triple peaks marked as C and G are the BLE connections transferring 247 Bytes of data. They recur at 30 ms intervals.

Note that a BLE connection coincides with the CPU being active between E and F.
A more detailed view of a BLE triple advertisement is shown in Figure 5-4. Note that the timescale has been expanded to 1 ms/div to capture more detail.

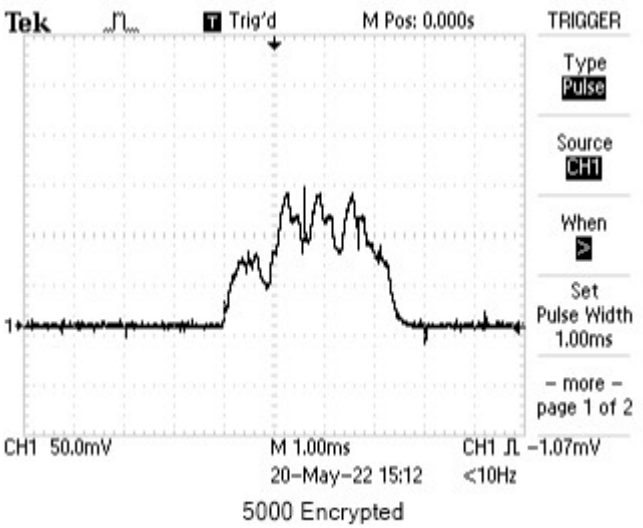

Figure 5-4: BLE Connection

An example of the measurements of the energy consumption for each part of the cycle is shown in Figure 5-5, where the values are in nJ.

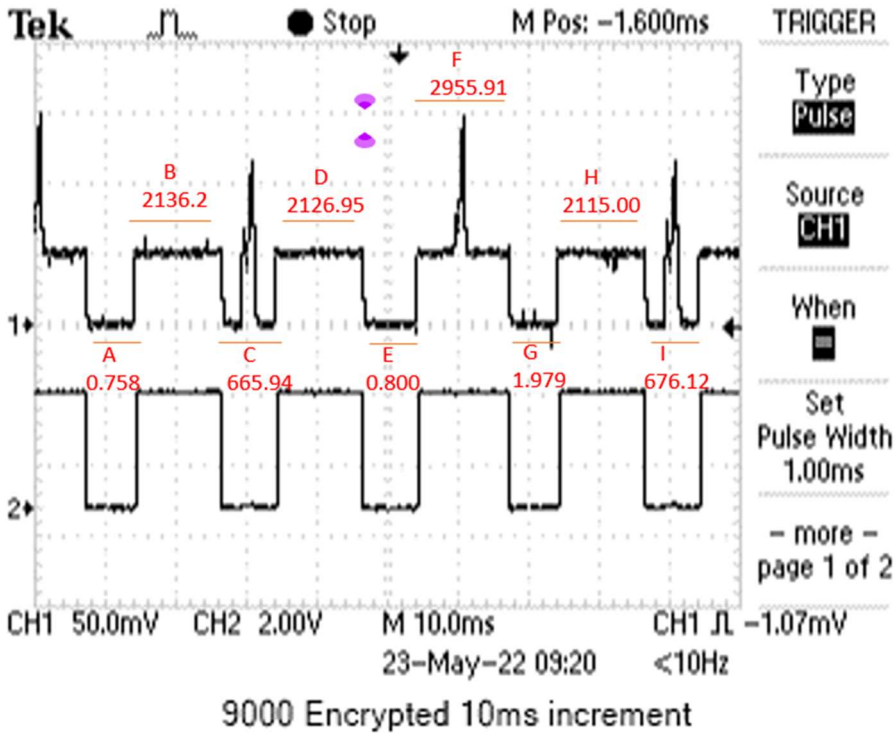

Figure 5-5: Example Waveform with Energy Consumption

5.5 Analysis of Results

The collated results from Appendix A are shown in Figure 5-6.

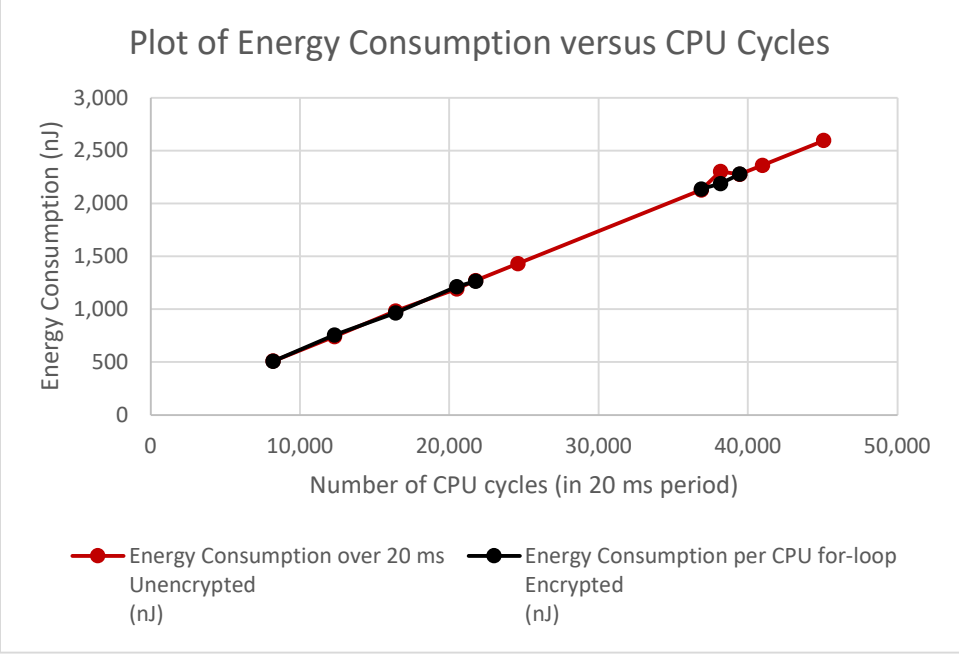

Figure 5-6: Correlation of Energy Consumption to CPU Cycles

As would be expected, the energy consumption increases linearly with the value of SUCK_DELAY, which is the number of CPU cycles (for-loop counter) in a 20 ms period.

A clear result is that when configured to encrypt the data there is no significant change in energy required; for some values, the measured energy consumption was slightly lower for encryption compared to being configured for unencrypted operations, although this may be an artefact of other activity within the board or due to the measuring precision of the oscilloscope.

Calculating the energy per 'for-loop' yields the following data Table 5-2 (note that the values are in pJ):

Table 5-2: Energy Consumption per 'for-loop' (pJ)

CPU Cycles - Decimal (SUCK_DELAY)	Energy Consumption per CPU for-loop Unencrypted (pJ)	Energy Consumption per CPU for-loop Encrypted (pJ)
8,192	62	62
12,288	60	62
16,384	60	59
20,480	58	59
21,760	58	58
24,576	58	
36,864	58	58
38,144	60	57
39,424	58	58
40,960	58	
45,056	58	

Examining the detail of the consumption in Figure 5-7, demonstrates a decrease in the energy consumption per for-loop as the duty cycle increases. This is expected as energy is required to activate the processor from Standby or Idle mode into Active mode. In Standby mode, only the always-on (AON) domain is active. An external wake-up event, RTC event, or Sensor controller event is required to bring the device back to active mode. All GPIOs are latched in standby mode (Texas Instruments, 2021).

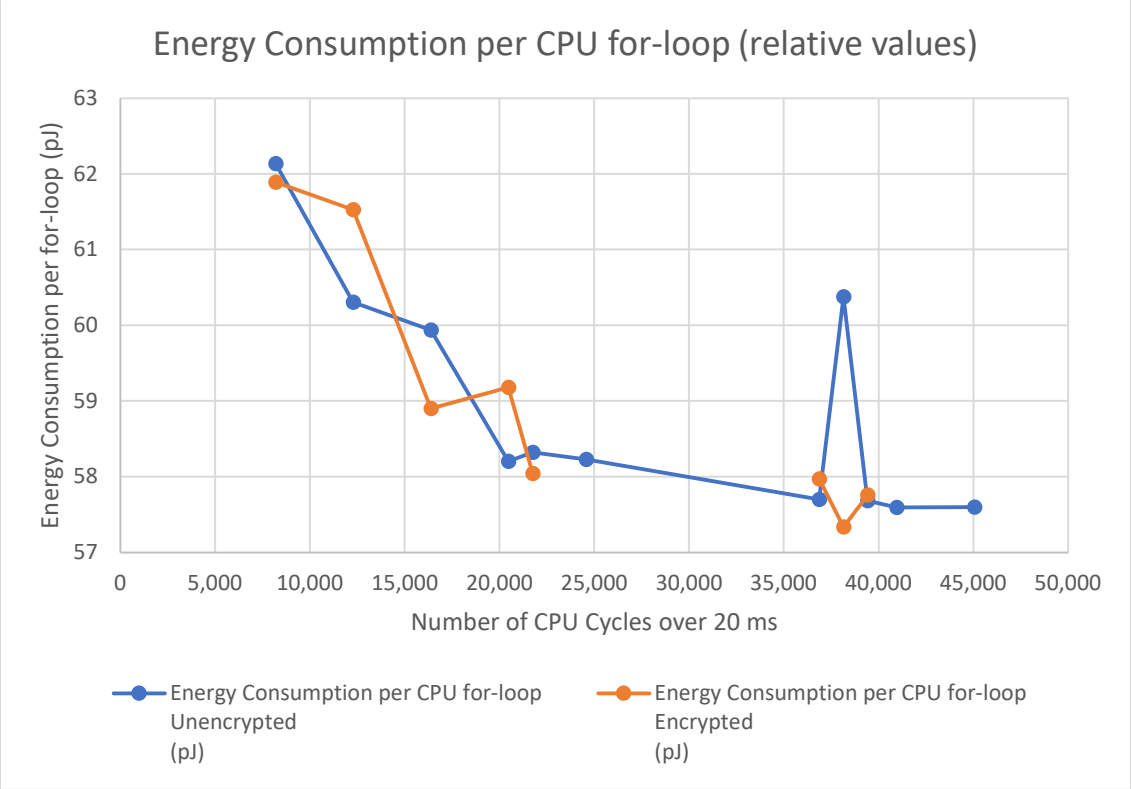

Figure 5-7: Plot of Energy Consumption per 'for-loop' – Relative Values (pJ)

In absolute terms, the reduction in energy consumption per for-loop is approaching a steady minimum at approximately 58 pJ (see Figure 5-8).

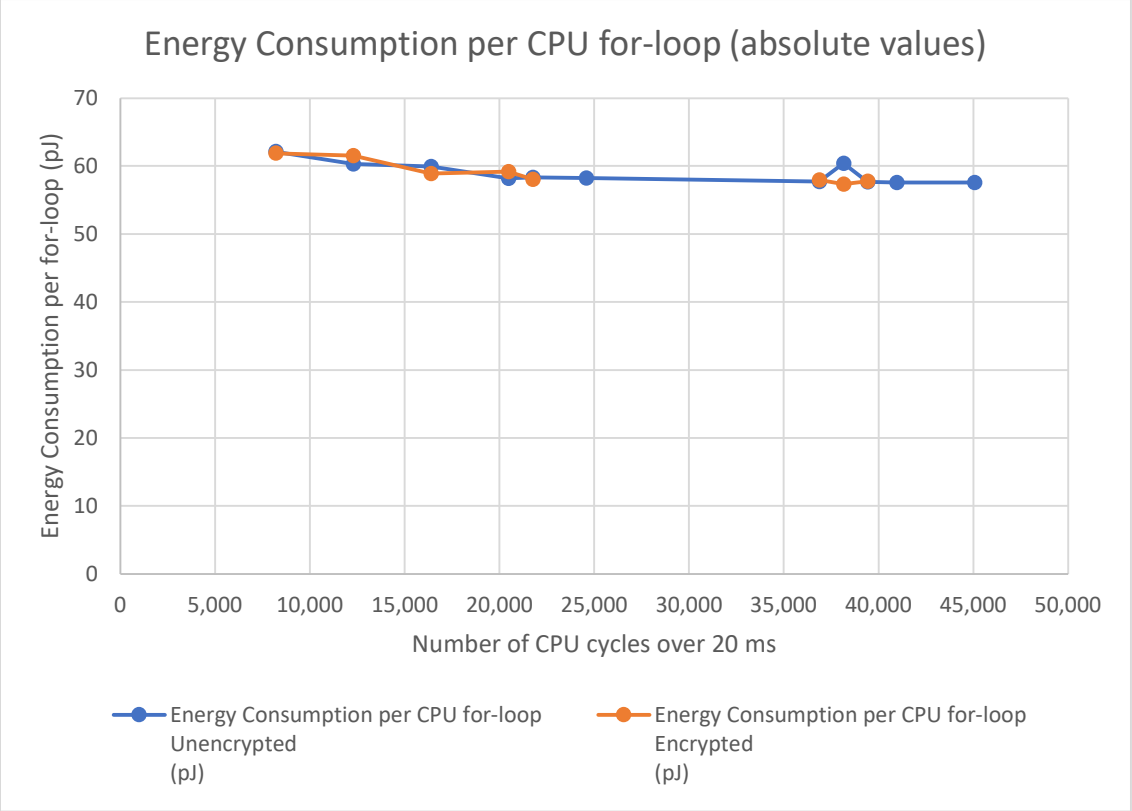

Figure 5-8: Plot of Energy Consumption per 'for-loop' – Absolute Values (pJ)

5.6 Discussion

5.6.1 Energy Consumption

As shown in Figure 5-5 and Section A-2, a BLE message with MTU = 251 consumes approximately 660 – 676 nJ, which will also be dependent upon transmitted signal strength. For comparison, this is approximately 50 times lower than the energy consumption of a BLE triple advertisement by the SPHERE SPW-1 wearable which consumes between 37 µJ (at -20dBm) and 60 µJ (at 4 dBm) (Fafoutis et al., 2017a). This may be partially due to the CC2652 board being more energy efficient than the nRF51822 used in the SPW-1, with the TI-RTOS Power Manager optimising the energy consumption on the CC2652 board, but other differences in the experimental setup will also have influenced the differences in energy consumption.

The BLE message (660-676 nJ) is approximately equivalent to the CPU processing 11,380 – 11,655 for-loops, assuming that the processor is already in Active mode due to receiving a wake-up event from the Sensor Controller. This provides a metric by which

specific on-board processing and machine learning strategies can be assessed as to their energy efficiencies compared to offloading the raw data for processing. Advancements in ML for edge devices, such as TinyML and TensorFlow Lite for Microcontrollers, may enable very specific models to be run on the device within this energy budget.

5.6.2 Encryption

A striking result was that enabling the AES-128 encryption had no discernible impact on energy usage. This may be due to:

- Hardware accelerators - according to TI, '*the CC2652RB device comes with a wide set of modern cryptography-related hardware accelerators, drastically reducing code footprint and execution time for cryptographic operations. It also has the benefit of being lower power and improves availability and responsiveness of the system because the cryptography operations runs in a background hardware thread*' (Texas Instruments, 2021).
- The SUCK_DELAY cycle not fully invoking the encryption of data. This would require a deeper analysis of the SAPPHIRE code which is not available for this research project.

As part of assessing the suitability of a system such as the CC2652RB for an RHM wearable, a full energy analysis would need to be undertaken with all the cryptography-related hardware accelerators, to determine which cryptography scheme provides the optimal balance between energy consumption and the strength of security required.

Chapter 6 Conclusion

The mass markets for both industrial and consumer IoT devices has made relatively powerful processor readily available at a low cost (a few Pounds/Euros/US Dollars per chip). Some of these devices have been optimised to be highly energy efficient, especially for industrial applications where it would be costly to replace a battery. These advances have opened up an opportunity for low-cost wearables for remote health monitoring where smartphones are not appropriate, for both low- and middle-income countries and western users with physical or cognitive conditions challenges that make smartphones challenging.

In parallel with the development of IoT devices, the emergence of IoT communications provides an opportunity for under-served communities to install a communications network infrastructure to serve remote citizens. The literature review identified several trials where LPWAN communications demonstrated potential benefits to users outside of a care- or smart-home environment. LoRaWAN was the most frequently used technology, principally based upon data rates, range, power consumption, availability and cost. However, no one IoT solution will be best suited against all use cases.

In addition, 5G IoT communications have become a reality and are being rolled-out in several regions, including the UK. Although users are tied to procuring these services from licensed network operators, they offer many benefits in terms of throughput, reliability, quality and security in a managed service. Where affordability and coverage are not primary constraints, licenced 5G networks may be the preferred solution for care services and organisations who need to remotely monitor workers.

However, the primary conclusion of this research is that technology should not be the primary focus when assessing the viability of an RHM service. The technology by itself is of no practical benefit to the principal users – there has to be a whole-service approach as to how the user's physiological and environmental data can help support each individual user, as well as being usable by the carers and supervisors providing care and support. As with all systems engineering projects, a clear understanding of the business and user needs, from a technology-agnostic viewpoint, must be the starting point for any development. In Chapter 2, this research identified the key stakeholders and primary

workflows for an RHM service, followed by a structured approach to identifying the high-level user requirements for a service. An examination of the key regulatory requirements, including GDPR and HIPAA, demonstrated that they need to be understood and incorporated into the design of the service from the outset, especially for factors such as privacy, security and reliability/duty-of-care of RHM services, which can be much more stringent than those normally associated with wellness or health apps for consumers.

Chapter 3 examined the technical requirements derived from the user requirements. Again, security and privacy are prevailing drivers for any technological solution. Frameworks such as ETSI/EN 303 645 and NIST IR 8295A provide high-level requirements for IoT products and services and may form the basis for future regulation. Other frameworks, such as the IoTSF Assurance Framework and Questionnaire provide advice as to how some of these requirements can be satisfied.

There is growing excitement at the emerging opportunities offered by embedded machine learning. The TinyML movement and tools such as TensorFlow Lite for Microcontrollers are enabling researchers to demonstrate relatively powerful models being run on very constrained and energy-efficient microcontrollers, albeit those models may be very specific to a particular task and data-set. However, again, technological advancements need to work within the wider user service requirements. The need for ‘understandable’ AI, including where the reliability of the outcomes of the ML the model can be determined, is likely to rise in importance for RHM applications. The UK MHRA, Health Canada and the US FDA have already issued “Good Machine Learning Practice for Medical Device Development: Guiding Principles”, and further scrutiny by regulators is highly likely.

Chapter 7 Plan for further research

Due to Covid-19, there was no opportunity to undertake field trials of the IoT LPWAN communications identified in Section 3.4. Further research will be practical assessments of the coverage, reliability, latency and energy efficiency of LoRaWAN to support representative mHealth applications, such as detection of arrhythmia using ML on an IoT device (Sanchez-Iborra, 2021, Yang et al., 2022). A factor to be assessed will be the placement of the device on the subject's body to determine the impact of attenuation or absorption of signal due to the human body.

The privacy requirements for the IoT devices discussed in Sections 2.3.4 and 3.2 will be compared to the current and emerging IoT devices to identify the key enabling technologies that should be included as part of the 'Secure by Design' principles in any mHealth system.

The IoT devices identified as being capable of 'Secure by Design' will then be compared against their ability to support the most suitable LPWAN communications identified by the earlier simulation and also their ability to support ML algorithms. The research will propose guidelines as to how trade-offs between battery life, coverage, privacy can be assessed against a range of potential mHealth applications.

Chapter 8 References

A4AI 2022. Advancing Meaningful Connectivity: Towards Active and Participatory Digital Societies. Alliance for Affordable Internet.

ABDUL-GHANI, H. A. & KONSTANTAS, D. 2019. A Comprehensive Study of Security and Privacy Guidelines, Threats, and Countermeasures: An IoT Perspective. *Journal of Sensor and Actuator Networks,* 8**,** 22.

ABOUZAKHAR;, N. S., JONES;, A. & ANGELOPOULOU, O. 2017. Internet of Things Security: A Review of Risks and Threats to Healthcare Sector.

AHMED, G., ISLAM, S. U., SHAHID, M., AKHUNZADA, A., JABBAR, S., KHAN, M. K., RIAZ, M. & HAN, K. 2018. Rigorous Analysis and Evaluation of Specific Absorption Rate (SAR) for Mobile Multimedia Healthcare. *IEEE Access,* 6**,** 29602-29610.

AKBAR, M. S., YU, H. & CANG, S. 2017. IEEE 802.15.4 Frame Aggregation Enhancement to Provide High Performance in Life-Critical Patient Monitoring Systems. *Sensors (Basel, Switzerland),* 17**,** 241.

AKMANDOR, A. O. & JHA, N. K. 2018. Smart Health Care: An Edge-Side Computing Perspective. *IEEE Consumer Electronics Magazine,* 7**,** 29-37.

AKPAKWU, G. A., SILVA, B. J., HANCKE, G. P. & ABU-MAHFOUZ, A. M. A Survey on 5G Networks for the Internet of Things: Communication Technologies and Challenges. *IEEE Access,* 6**,** 3619-3647.

ALABOUD, K., SHAHREEN, M., ISLAM, H., PAUL, T., RANA, M. K. Z., MORRISON, A., KUMAR, A. & MOSA, A. S. M. 2022. Clinicians' Perspectives in Using Patient-Generated Health Data to Improve Ischemic Heart Disease Management. *AMIA Annu Symp Proc,* 2022**,** 112-119.

ALAM, M. M., MALIK, H., KHAN, M. I., PARDY, T., KUUSIK, A. & MOULLEC, Y. L. 2018. A Survey on the Roles of Communication Technologies in IoT-Based Personalized Healthcare Applications. *IEEE Access,* 6**,** 36611-36631.

AMIRI-ZARANDI, M., DARA, R. A. & FRASER, E. 2020. A survey of machine learning-based solutions to protect privacy in the Internet of Things. *Computers & Security,* 96**,** 101921.

ARANDA-JAN, C. B., MOHUTSIWA-DIBE, N. & LOUKANOVA, S. 2014. Systematic review on what works, what does not work and why of implementation of mobile health (mHealth) projects in Africa. *BMC Public Health,* 14**,** 188.

ARNON, S., BHASTEKAR, D., KEDAR, D. & TAUBER, A. 2003. A comparative study of wireless communication network configurations for medical applications. *IEEE Wireless Communications,* 10**,** 56-61.

AST SPACEMOBILE. 2022. *The first space-based cellular broadband network for mobile phones* [Online]. Available: https://ast-science.com/spacemobile/ [Accessed 2022].

BEACH, C., KRACHUNOV, S., POPE, J., FAFOUTIS, X., PIECHOCKI, R. J., CRADDOCK, I. & CASSON, A. J. 2018. An ultra low power personalizable wrist worn ECG monitor integrated with IoT infrastructure. *IEEE Access,* 6**,** 44010-44021.

BELLIDO-OUTEIRINO, F. J., ARIAS, J. M. F., CALVO, R. R. & ROLDAN, M. T. LR-WPAN technologies. An approach to industrial applications. 2008 First Conference on IT Revolutions, 17-19 Dec. 2008 2008. 1-4.

BELLIDO-OUTEIRINO, F. J., GONZÁLEZ R., M., MORENO M., A. & DE LA CRUZ F, J. L. "Low Power Wireless Technologies: An Approach to Medical Applications". 2009 Berlin, Heidelberg. Springer Berlin Heidelberg, 14-20.

BHATTI, A., SIYAL, A. A., MEHDI, A., SHAH, H., KUMAR, H. & BOHYO, M. A. Development of cost-effective tele-monitoring system for remote area patients. 2018 International Conference on Engineering and Emerging Technologies (ICEET), 22-23 Feb. 2018 2018. 1-7.

BLYTHE, J. M., MICHIE, S., WATSON, J. & LEFEVRE, C. E. 2017. Internet of Things in Healthcare: Identifying key malicious threats, end-user protective and problematic behaviours. *Frontiers in Public Health*.

BOR, M. 2020. *Towards the efficient use of LoRa for wireless sensor networks.* Lancaster University.

BOR, M. C., ROEDIG, U., VOIGT, T. & ALONSO, J. M. 2016. Do LoRa Low-Power Wide-Area Networks Scale? *Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems.* Malta, Malta: Association for Computing Machinery.

CENTELLES, R. P., FREITAG, F., MESEGUER, R., NAVARRO, L., OCHOA, S. F. & SANTOS, R. M. A LoRa-Based Communication System for Coordinated Response in an Earthquake Aftermath. Multidisciplinary Digital Publishing Institute Proceedings, 2019. 73.

CONTINUA 2016. H.810 Interoperability design guidelines for personal connected health systems. 2016 ed.

CONTINUA 2017. Introduction to the Continua Design Guidlines 2017.

CORDEIRO, C. & PATEL, M. 2007. Body area networking standardization: present and future directions. *Proceedings of the ICST 2nd international conference on Body area networks.* Florence, Italy: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

DARWISH, S., NOURETDINOV, I. & WOLTHUSEN, S. D. 2017. Towards Composable Threat Assessment for Medical IoT (MIoT). *Procedia Computer Science,* 113**,** 627-632.

DEPARTMENT FOR DIGITAL, C., MEDIA AND SPORT (DCMS) 2018. Code of Practice for Consumer IoT Security.

DHUKARAM, A. V., BABER, C., ELLOUMI, L., BEIJNUM, B. J. V. & STEFANIS, P. D. End-User perception towards pervasive cardiac healthcare services: Benefits, acceptance, adoption, risks, security, privacy and trust. 2011 5th International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth) and Workshops, 23-26 May 2011 2011. 478-484.

ENISA 2015. Security and Resilience in eHealth Infrastructures and Services Security Challenges and Risks.

ETSI 2018a. ETSI EN 300 220-2 V3.2.1 Short Range Devices (SRD) operating in the frequency range 25 MHz to 1 000 MHz; Part 2: Harmonised Standard for access to radio spectrum for non specific radio equipment.

ETSI 2018b. TS 103 357 Short Range Devices; Low Throughput Networks (LTN); Protocols for radio interface A.

ETSI 2019. ETSI TS 103 645 CYBER Cyber Security for Consumer Internet of Things.

ETSI 2020. ETSI EN 303 645 V2.1.1(2020-06) CYBER; Cyber Security for Consumer Internet of Things.

FABBRICATORE, C., ZUCKER, M., ZIGANKI, S. & KARDUCK, A. P. Towards an unified architecture for smart home and Ambient Assisted Living solutions: A focus on elderly people. Digital Ecosystems and Technologies Conference

(DEST), 2011 Proceedings of the 5th IEEE International Conference on, May 31 2011-June 3 2011 2011. 305-311.

FAFOUTIS, X., JANKO, B., MELLIOS, E., HILTON, G., SHERRATT, R. S., PIECHOCKI, R. & CRADDOCK, I. SPW-1: A Low-Maintenance Wearable Activity Tracker for Residential Monitoring and Healthcare Applications. 2017a Cham. Springer International Publishing, 294-305.

FAFOUTIS, X., MARCHEGIANI, L., ELSTS, A., POPE, J., PIECHOCKI, R. & CRADDOCK, I. Extending the battery lifetime of wearable sensors with embedded machine learning. 2018 IEEE 4th World Forum on Internet of Things (WF-IoT), 5-8 Feb. 2018 2018. 269-274.

FAFOUTIS, X., VAFEAS, A., JANKO, B., SHERRATT, R. S., POPE, J., ELSTS, A., MELLIOS, E., HILTON, G., OIKONOMOU, G. & PIECHOCKI, R. 2017b. Designing wearable sensing platforms for healthcare in a residential environment. *EAI Endorsed Transactions on Pervasive Health and Technology,* 3.

FAKHR, S. M. & FOTOUHI, F. 2016. A Survey on Security Issues in Internet of Thing.

FCC 2017. Promoting Telehealth in Rural America. *In:* COMMISSION, F. C. (ed.). Washington, D.C.

FCC 2020a. 2020 Broadband Deployment Report.

FCC. 2020b. *Electronic Code of Federal Regulations Title 47* [Online]. Available: https://www.ecfr.gov/cgi-bin/retrieveECFR?gp=&SID=a6f874d4ab6e8c17e18ba8cdea5154f4&mc=true&r=SECTION&n=se47.1.15_1247 [Accessed 4 January 2021 2021].

FCC 2022. FCC Rejects LTD Broadband, Starlink Bids for Broadband Subsidies. Washington, DC.

FDA, HEALTH CANADA & MHRA 2021. Good Machine Learning Practice for Medical Device Development: Guiding Principles. *In:* EXCELLENCE, T. D. H. C. O. (ed.).

FERNANDEZ, F. & PALLIS, G. C. Opportunities and challenges of the Internet of Things for healthcare: Systems engineering perspective. 2014 4th International Conference on Wireless Mobile Communication and Healthcare - Transforming Healthcare Through Innovations in Mobile and Wireless Technologies (MOBIHEALTH), 3-5 Nov. 2014 2014. 263-266.

GANZ, F., PUSCHMANN, D., BARNAGHI, P. & CARREZ, F. 2015. A Practical Evaluation of Information Processing and Abstraction Techniques for the Internet of Things. *IEEE Internet of Things Journal,* 2**,** 340-354.

GHAMARI, M., JANKO, B., SHERRATT, R., HARWIN, W., PIECHOCKIC, R. & SOLTANPUR, C. 2016. A Survey on Wireless Body Area Networks for eHealthcare Systems in Residential Environments. *Sensors,* 16**,** 831.

GOLDBERG, L., LIDE, B., LOWRY, S., MASSETT, H. A., O'CONNELL, T., PREECE, J., QUESENBERY, W. & SHNEIDERMAN, B. 2011. Usability and accessibility in consumer health informatics: current trends and future challenges. *American journal of preventive medicine,* 40**,** S187-S197.

GÓMEZ-CARMONA, O., CASADO-MANSILLA, D., KRAEMER, F. A., LÓPEZ-DE-IPIÑA, D. & GARCÍA-ZUBIA, J. 2020. Exploring the computational cost of machine learning at the edge for human-centric Internet of Things. *Future Generation Computer Systems,* 112**,** 670-683.

GRECO, L., PERCANNELLA, G., RITROVATO, P., TORTORELLA, F. & VENTO, M. 2020. Trends in IoT based solutions for health care: Moving AI to the edge. *Pattern Recognition Letters,* 135**,** 346-353.

GSMA 2016. GSMA: 3GPP Low Power Wide Area Technologies. *In:* GRANT, S. (ed.).

HADDADPAJOUH, H., DEHGHANTANHA, A., M. PARIZI, R., ALEDHARI, M. & KARIMIPOUR, H. 2021. A survey on internet of things security: Requirements, challenges, and solutions. *Internet of Things,* 14**,** 100129.

HAMMI, B., KHATOUN, R., ZEADALLY, S., FAYAD, A. & KHOUKHI, L. 2018. IoT technologies for smart cities. *IET Networks* [Online], 7. Available: https://digital-library.theiet.org/content/journals/10.1049/iet-net.2017.0163.

HARIKRISHNAN, H. 2017. *IoT and Healthcare: Redesigning Care Pathways* [Online]. Available: https://www.iotforall.com/health-monitoring-using-iot [Accessed 30 August 2022 2022].

HAYES, C. 2022. Islands of Connectivity [communications - networks and services]. *Engineering & Technology,* 17**,** 32-33.

HE, D., ZEADALLY, S., KUMAR, N. & LEE, J. H. 2016. Anonymous Authentication for Wireless Body Area Networks With Provable Security. *IEEE Systems Journal,* PP**,** 1-12.

HEALTHCARE AND PUBLIC HEALTH SECTOR COORDINATING COUNCIL 2019. Medical Device and Health IT Joint Security Plan.

HL7 INTERNATIONAL, JOHN "MIKE" DAVIS, BERND BLOBEL, JOHN MOEHRKE & WILLIAMS, T. 2013. Guide to the HL7 Healthcare Privacy and Security Classification System (HCS).

HUI, T. K. L. & SHERRATT, R. S. 2017. Towards disappearing user interfaces for ubiquitous computing: human enhancement from sixth sense to super senses. *Journal of Ambient Intelligence and Humanized Computing,* 8**,** 449-465.

ICO. 2022. *Data protection by design and default* [Online]. Available: https://ico.org.uk/for-organisations/guide-to-data-protection/guide-to-the-general-data-protection-regulation-gdpr/accountability-and-governance/data-protection-by-design-and-default/ [Accessed 12 Decemer 2022].

IEEE 2015a. IEEE Std11073-00103:2015 Health informatics Personal health device communication Part 00103: Overview.

IEEE 2015b. IEEE Std 11073-10101a-2015 Health informatics--Point-of-care medical device communication --Part 10101: Nomenclature Amendment 1: Additional Definitions. *IEEE Std 11073-10101a-2015 (Amendment to ISO/IEEE 11073-10101:2004)***,** 1-116.

IETF 2018a. An architecture for authorization in constrained environments draft-ietf-ace-actors-07.

IETF 2018b. RFC 8376 Low-Power Wide Area Network (LPWAN) Overview Draft.

INFORMATION & PRIVACY COMMISSIONER OF ONTARIO 2013. Privacy by Design.

INFORMATION AND PRIVACY COMMISSIONER ONTARIO, ANN CAVOUKIAN, STUART SHAPIRO & R. JASON CRONK 2014. Privacy Engineering: Proactively Embedding Privacy, by Design.

INFORMATION COMMISSIONER'S OFFICE. 2017. *Data protection by design and default* [Online]. Available: https://ico.org.uk/for-organisations/guide-to-the-general-data-protection-regulation-gdpr/accountability-and-governance/data-protection-by-design-and-default/ [Accessed 14 December 2017].

INTERNATIONAL COMMISSION ON NON-IONIZING RADIATION PROTECTION 2020. Guidelines for Limiting Exposure to Electromagnetic Fields (100 kHz to 300 GHz). *Health Physics,* 118**,** 483-524.

IORGA, M., FELDMAN, L., BARTON, R., MARTIN, M. J., GOREN, N. & MAHMOUDI, C. 2018. Fog Computing Conceptual Model. Special Publication (NIST SP)-500–325. *National Institute of Standards and Technology, L. Erazo et*

al.: A Domain-Specific Language for Modeling IoT System Architectures that Support Monitoring.
IOT SECURITY FOUNDATION. 2016. *Establishing Principles for Internet of Things Security* [Online]. Available: https://iotsecurityfoundation.org/ [Accessed].
IOT SECURITY FOUNDATION 2021. IoT Security Assurance Framework. 3 ed.
ISLAM, S. M. R., KWAK, D., KABIR, M. H., HOSSAIN, M. & KWAK, K. 2015. The Internet of Things for Health Care: A Comprehensive Survey. *IEEE Access,* 3**,** 678-708.
ISO 2019. ISO 80601-2-61:2019 Medical electrical equipment. Particular requirements for basic safety and essential performance of pulse oximeter equipment.
ISO/IEC 2008. ISO/IEC 15408-2:2008 Information technology. Security techniques. Evaluation criteria for IT security.
ISO/IEC 2018. Internet of Things Reference Architecture 2018.
ITU. 2012a. *ITU-T Y.2060 Overview of the Internet of things* [Online]. Available: http://www.itu.int/ITU-T/recommendations/rec.aspx?rec=y.2060 [Accessed 8 July 2016].
ITU 2012b. ITU-T Y. 4000/Y. 2060 (06/2012). *Overview of the Internet of things (06 2012)*.
ITU 2017. H.810 Interoperability design guidelines for personal connected health systems: Introduction.
ITU 2018a. ICTs, LDCs and the SDGs Achieving universal and affordable Internet in the least developed countries.
ITU 2018b. ITU Council Contribution to the High-Level Political Forum on Sustainable Development (HLPF). 8 March 2018 ed.: ITU.
ITU 2020a. Manual for measuring ICT access and use by households and individuals, 2020 Edition. ITU.
ITU 2020b. Measuring digital development Facts and figures 2020. Geneva.
ITU 2020c. Radio Regulations Edition of 2020. *Vol 1.*
JIEHUI, J. & ZHANG, J. 2007. Remote patient monitoring system for China. *IEEE Potentials,* 26**,** 26-29.
JOHNSON M. ITU DEPUTY SECRETARY-GENERAL. The Commonwealth of Nations and United Nations Perspectives: Bridging the Health and Technology Sectors with the Global Goals. 71st World Health Assembly "Creating a digital Health Dynamic for Universal Health Coverage 2030", 22 May 2018 2018 Geneva.
KARAGEORGOS, G., ANDREADIS, I., PSYCHAS, K., MOURKOUSIS, G., KIOURTI, A., LAZZI, G. & NIKITA, K. S. 2019. The Promise of Mobile Technologies for the Health Care System in the Developing World: A Systematic Review. *IEEE Reviews in Biomedical Engineering,* 12**,** 100-122.
KHAZBAK, Y., IZZ, M., ELBATT, T., FAHIM, A., GUIRGUIS, A. & YOUSSEF, M. 2017. Cost-Effective Data Transfer for Mobile Health Care. *IEEE Systems Journal,* 11**,** 2663-2674.
KINEIS. 2022. *Kineis FAQ* [Online]. Available: https://www.kineis.com/en/faqs/ [Accessed 16 September 2022 2022].
KO, J., LU, C., SRIVASTAVA, M. B., STANKOVIC, J. A., TERZIS, A. & WELSH, M. 2010. Wireless sensor networks for healthcare. *Proceedings of the IEEE,* 98**,** 1947-1960.
KOHN, L. T., CORRIGAN, J. & DONALDSON, M. S. 2001. Institute of Medicine (US). Committee on Quality of Health Care in America.(2000). *To err is human: Building a safer health system*.

LATRE, B., BRAEM, B., MOERMAN, I., BLONDIA, C. & DEMEESTER, P. 2011. A survey on wireless body area networks. *Wireless Networks,* 17**,** 1-18.

LINK LABS 2016. A comprehensive look at Low Power, Wide Area Networks.

LO, B. P. L., IP, H. & YANG, G. Z. 2016. Transforming Health Care: Body Sensor Networks, Wearables, and the Internet of Things. *IEEE Pulse,* 7**,** 4-8.

LUDWIG, W., WOLF, K. H., DUWENKAMP, C., GUSEW, N., HELLRUNG, N., MARSCHOLLEK, M., WAGNER, M. & HAUX, R. 2011. Health-enabling technologies for the elderly--an overview of services based on a literature review. *Comput Methods Programs Biomed,* 106**,** 70-8.

MAAß, L., FREYE, M., PAN, C.-C., DASSOW, H.-H., NIESS, J. & JAHNEL, T. 2022. The Definitions of Health Apps and Medical Apps From the Perspective of Public Health and Law: Qualitative Analysis of an Interdisciplinary Literature Overview. *JMIR mHealth and uHealth,* 10**,** e37980.

MALINA, L., HAJNY, J., FUJDIAK, R. & HOSEK, J. 2016. On perspective of security and privacy-preserving solutions in the internet of things. *Computer Networks,* 102**,** 83-95.

MANYSECURED 2019. Secure IOT Gateways Whitepaper on the need for industry collaboration. *In:* PROJECT, M. G. (ed.). IoTSF.

MCDONAGH, S. T. J., RHODES, S., WARREN, F. C., KEENAN, S., PENTECOST, C., KEELING, P., JAMES, M., TAYLOR, R. S. & CLARK, C. E. 2022. Performance of the imPulse device for the detection of atrial fibrillation in hospital settings. *Cardiovascular Digital Health Journal,* 3**,** 171-178.

MERENDA, M., PORCARO, C. & IERO, D. 2020. Edge Machine Learning for AI-Enabled IoT Devices: A Review. *Sensors,* 20**,** 2533.

MICHELL, V. 2017 est. IoT & Health.

MIHAILIDIS, A., BOGER, J., LOCKTON, V. & CHIBBA, M. 2010. *Sensors and In-Home Collection of Health Data: A Privacy by Design Approach*.

MOSENIA, A. & JHA, N. K. 2017. A Comprehensive Study of Security of Internet-of-Things. *IEEE Transactions on Emerging Topics in Computing,* 5**,** 586-602.

MURSHED, M. G. S., MURPHY, C., HOU, D., KHAN, N., ANANTHANARAYANAN, G. & HUSSAIN, F. 2019. Machine Learning at the Network Edge: A Survey. *ArXiv,* abs/1908.00080.

NCSC. 2017. *Secure by Default* [Online]. Available: https://www.ncsc.gov.uk/articles/secure-default [Accessed 9 January 2018].

NCSC. 2019. *Secure design principles* [Online]. Available: https://www.ncsc.gov.uk/collection/cyber-security-design-principles [Accessed 8 August 2022 2022].

NICE. 2022. *Office for Digital Health* [Online]. Available: https://www.nice.org.uk/about/what-we-do/digital-health/office-for-digital-health#innovative-devices-access-pathway [Accessed 1 November 2022 2022].

NIST 2017. Report on Lightweight Cryptography.

NIST 2020. NISTIR 8259A IoT Device Cybersecurity Capability Core Baseline. *In:* SMITH, M. F. K. M. K. S. M. (ed.). Gaithersburg, MD, USA: National Institute of Standards and Technology.

NIST 2023. Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology.

OFCOM 2020. Connected Nations Update Summer 2020.

OH, H., JADAD, A., RIZO, C., ENKIN, M., POWELL, J. & PAGLIARI, C. 2005. What Is eHealth (3): A Systematic Review of Published Definitions. *Journal of Medical Internet Research,* 7.

PASLUOSTA, C. F., GASSNER, H., WINKLER, J., KLUCKEN, J. & ESKOFIER, B. M. 2015. An Emerging Era in the Management of Parkinson's Disease: Wearable Technologies and the Internet of Things. *IEEE Journal of Biomedical and Health Informatics,* 19**,** 1873-1881.

PETERSEN, C., ADAMS, S. A. & DEMURO, P. R. 2015. mHealth: Don't Forget All the Stakeholders in the Business Case. *Med 2 0,* 4**,** e4.

POYNER, I. K. & SHERRATT, R. S. 2018. Privacy and security of consumer IoT devices for the pervasive monitoring of vulnerable people. *IET Conference Proceedings.* London: Institution of Engineering and Technology.

POYNER, I. K. & SHERRATT, R. S. 2019. Improving access to healthcare in rural communities — IoT as part of the solution. 3rd IET International Conference on Technologies for Active and Assisted Living (TechAAL 2019), 25-25 March 2019 2019 London. London: IET, 1-6.

QADRI, Y. A., NAUMAN, A., ZIKRIA, Y. B., VASILAKOS, A. V. & KIM, S. W. 2020. The Future of Healthcare Internet of Things: A Survey of Emerging Technologies. *IEEE Communications Surveys & Tutorials,* 22**,** 1121-1167.

QUAGLIO, G., DARIO, C., KARAPIPERIS, T., DELPONTE, L., MCCORMACK, S., TOMSON, G. R., MICHELETTI, G., BONNARDOT, L., PUTOTO, G. & ZANABONI, P. 2016. Information and communications technologies in low and middle-income countries: Survey results on economic development and health. *Health Policy and Technology,* 5**,** 318-329.

QUERALTA, J. P., GIA, T. N., TENHUNEN, H. & WESTERLUND, T. Edge-AI in LoRa-based Health Monitoring: Fall Detection System with Fog Computing and LSTM Recurrent Neural Networks. 2019 42nd International Conference on Telecommunications and Signal Processing (TSP), 1-3 July 2019 2019. 601-604.

RAMESH, J., ABURUKBA, R. & SAGAHYROON, A. 2021. A remote healthcare monitoring framework for diabetes prediction using machine learning. *Healthc Technol Lett,* 8**,** 45-57.

RASHIDI, P. & MIHAILIDIS, A. 2013. A survey on ambient-assisted living tools for older adults. *IEEE J Biomed Health Inform,* 17**,** 579-90.

SANCHEZ-IBORRA, R. 2021. LPWAN and Embedded Machine Learning as Enablers for the Next Generation of Wearable Devices. *Sensors (Basel),* 21.

SECRETARY OF STATE FOR HEALTH 2002. The Medical Devices Regulations 2002. United Kingdom.

SHAPIRO, M., JOHNSTON, D., WALD, J. & MON, D. 2012. Patient-generated health data. *RTI International, April*.

SHERRATT, R. S. & DEY, N. 2020. Low-Power Wearable Healthcare Sensors. *Electronics,* 9**,** 892.

SHU, M., YUAN, D., ZHANG, C., WANG, Y. & CHEN, C. 2015. A MAC protocol for medical monitoring applications of wireless body area networks. *Sensors (Basel, Switzerland),* 15**,** 12906-12931.

SOPIC, D., AMINIFAR, A., AMINIFAR, A. & ATIENZA, D. 2018a. Real-Time Event-Driven Classification Technique for Early Detection and Prevention of Myocardial Infarction on Wearable Systems. *IEEE Transactions on Biomedical Circuits and Systems,* 12**,** 982-992.

SOPIC, D., AMINIFAR, A. & ATIENZA, D. e-Glass: A Wearable System for Real-Time Detection of Epileptic Seizures. 2018 IEEE International Symposium on Circuits and Systems (ISCAS), 2018-01-01 2018b. IEEE.

SPHERE. 2018. *SPHERE* [Online]. Available: https://www.irc-sphere.ac.uk/allpublications [Accessed].

SUN, Y., LIU, J., WANG, J., CAO, Y. & KATO, N. 2020. When Machine Learning Meets Privacy in 6G: A Survey. *IEEE Communications Surveys & Tutorials*, 1-1.

SURESH, V. M., SIDHU, R., KARKARE, P., PATIL, A., LEI, Z. & BASU, A. Powering the IoT through embedded machine learning and LoRa. 2018 IEEE 4th World Forum on Internet of Things (WF-IoT), 5-8 Feb. 2018 2018. 349-354.

T-MOBILE 2022. T-Mobile Takes Coverage Above and Beyond With SpaceX.

TAHSIEN, S. M., KARIMIPOUR, H. & SPACHOS, P. 2020. Machine learning based solutions for security of Internet of Things (IoT): A survey. *Journal of Network and Computer Applications,* 161**,** 102630.

TEXAS INSTRUMENTS 2021. Datasheet CC2652RB SimpleLink Crystal-less BAW Multiprotocol 2.4 GHz Wireless MCU.

THE THINGS NETWORK. 2020. *List of Gateways* [Online]. Available: https://www.thethingsnetwork.org/docs/gateways/start/list.html [Accessed].

THE THINGS NETWORK. 2021. *The Things Network Reading* [Online]. Available: https://www.thethingsnetwork.org/community/reading/ [Accessed 1 January 2019].

TORTI, E., FONTANELLA, A., MUSCI, M., BLAGO, N., PAU, D., LEPORATI, F. & PIASTRA, M. Embedded Real-Time Fall Detection with Deep Learning on Wearable Devices. 2018 21st Euromicro Conference on Digital System Design (DSD), 29-31 Aug. 2018 2018. 405-412.

TSINGANOS, P. & SKODRAS, A. 2018. On the Comparison of Wearable Sensor Data Fusion to a Single Sensor Machine Learning Technique in Fall Detection. *Sensors,* 18**,** 592.

UCHIYAMA, R., OKADA, Y., KAKIZAKI, R. & TOMIOKA, S. 2022. End-to-End Convolutional Neural Network Model to Detect and Localize Myocardial Infarction Using 12-Lead ECG Images without Preprocessing. *Bioengineering,* 9**,** 430.

UKIL, A., SAHU, I., MAJUMDAR, A., RACHA, S. C., KULKARNI, G., CHOUDHURY, A. D., KHANDELWAL, S., GHOSE, A. & PAL, A. Resource Constrained CVD Classification Using Single Lead ECG On Wearable and Implantable Devices. 2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), 2021-11-01 2021. IEEE.

UNITED STATES COAST GUARD. 2017. *AIS Standard Position Report* [Online]. Available: https://www.navcen.uscg.gov/?pageName=AISMessagesB [Accessed].

UNIVERSITY OF READING. 2018. *SPHERE Sensor Platform for HEalthcare in a Residential Environment* [Online]. Available: https://www.reading.ac.uk/biologicalsciences/res/bm/bme/research/sphere/bioscsi-bme-sphere.aspx [Accessed 5 December 2021 2021].

UNIVERSITY OF READING. 2021. *Next Generation Wearable – Sapphire* [Online]. Available: https://research.reading.ac.uk/wearables/research/next-generation-wearables/ [Accessed 5 December 2021 2021].

VODAFONE. 2018. *Narrowband IoT* [Online]. Available: https://www.vodafone.co.uk/business/iot [Accessed 1 January 2019].

WHO GLOBAL OBSERVATORY FOR EHEALTH 2011. mHealth: new horizons for health through mobile technologies: second global survey on eHealth. Geneva: World Health Organization.

XIAO, L., WAN, X., LU, X., ZHANG, Y. & WU, D. 2018. IoT Security Techniques Based on Machine Learning: How Do IoT Devices Use AI to Enhance Security? *IEEE Signal Processing Magazine,* 35**,** 41-49.

YANG, J., LI, J., LAN, K., WEI, A., WANG, H., HUANG, S. & FONG, S. 2022. Multi-Label Attribute Selection of Arrhythmia for Electrocardiogram Signals with Fusion Learning. *Bioengineering,* 9**,** 268.

ZHANG, M., KERMANI, M. M., RAGHUNATHAN, A. & JHA, N. K. Energy-efficient and Secure Sensor Data Transmission Using Encompression. 2013 26th International Conference on VLSI Design and 2013 12th International Conference on Embedded Systems, 5-10 Jan. 2013 2013. 31-36.

ZHANG, R. & LIU, L. Security Models and Requirements for Healthcare Application Clouds. 2010. IEEE.

ZIEFLE, M., ROCKER, C. & HOLZINGER, A. Medical Technology in Smart Homes: Exploring the User's Perspective on Privacy, Intimacy and Trust. 2011 IEEE 35th Annual Computer Software and Applications Conference Workshops, 18-22 July 2011 2011. 410-415.

A Appendix A – Measurements

A-1 Unconnected BLE Triple Beacon

Prior to connection to the RPi, the BLE triple beacon waveform was measured.

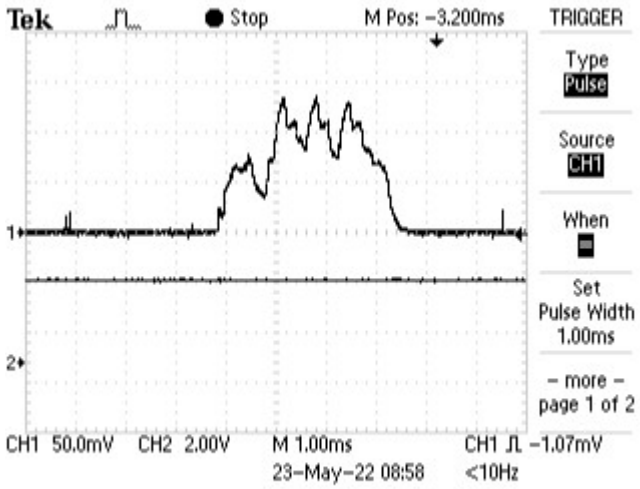

Unconnected BTLE triple beacon

Unconnected BTLE triple beacon.xlsx

The energy expended during the period of the triple beacon was 1730 nJ.

A-2 Connected with MTU = 251B

The BLE packet starts transmission at -4.524 ms on the display and continues until approximately -2.292 ms (2.23 ms duration).

The total energy used during the transmission was 659.9 nJ.

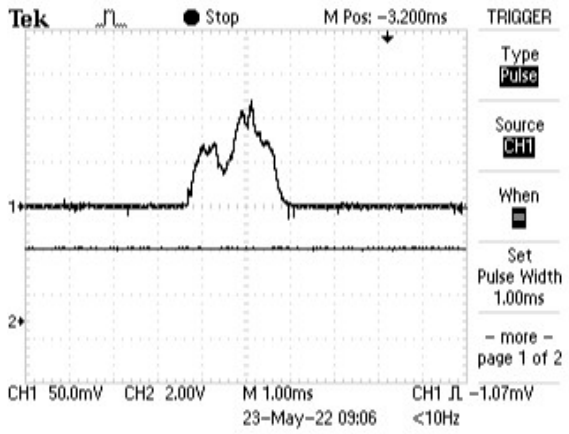

A-3 Energy Consumption over 20 ms CPU Cycles

The energy consumption over a complete 20 ms CPU cycle where it does not coincide with a BLE transmission was measured for different values of SUCK_DELAY and both with and without encryption.

In the embedded spreadsheets, for the data points indicated by 'Cycle', the energy consumed over each period (20 µs or 40 µs dependent upon the time-base) is summed (1,000 or 500 data points are used to provide 20 ms duration). Example images of the output are provided below.

SUCK_DELAY (Hex) (CPU cycles)	SUCK_DELAY (Dec) (CPU Cycles)	Energy Unencrypted (nJ)		Energy Encrypted (nJ)	
2000	8,192	Worksheet	509	Worksheet	507
3000	12,288	Worksheet	741	Worksheet	756
4000	16,384	Worksheet	982	Worksheet	965
5000	20,480	Worksheet	1,192	Worksheet	1,212
5500	21,760	Worksheet	1,269	Worksheet	1,263

6000	24,576	Worksheet	1,431		N/R
9000	36,864	Worksheet	2,127	Worksheet	2,137
9500	38,144	Worksheet	2,303	Worksheet	2,187
9a00	39,424	Worksheet	2,274	Worksheet	2,277
a000	40,960	Worksheet	2,359		N/R
b000	45,056	Worksheet	2,595		N/R

A-4 SUCK_DELAY = 5500_H ($21{,}760_D$) Unencrypted

In this measurement, the cycle is measured from the rise of the leading edge.

The total energy consumed over the cycle is 1,269 nJ.

This is LESS than the energy measured for the same duration cycle for encrypted data, although it the discrepancy is within measuring error.

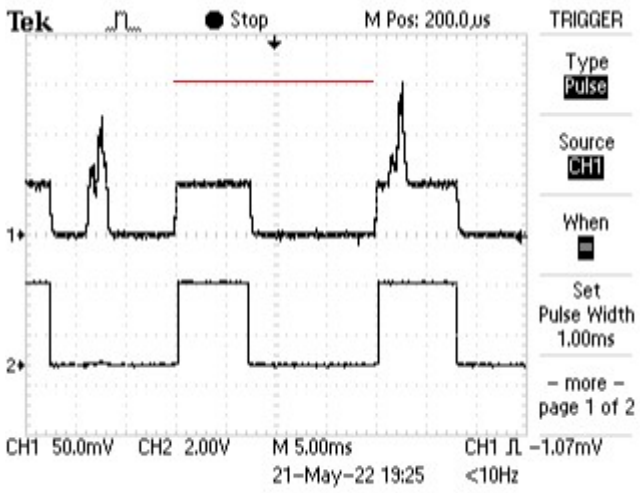

5500 Unencrypted

5500 Unencrypted.xlsx

A-5 SUCK_DELAY = 5500_H ($21{,}760_D$) Encrypted

In this measurement, a complete cycle without a BLE packet transmission is available starting at -13.60 ms (when the previous cycle returns to 0.0V) to 4.4 ms (when the CPU returns to sleep mode). The energy consumed during this cycle was 1,262 nJ.

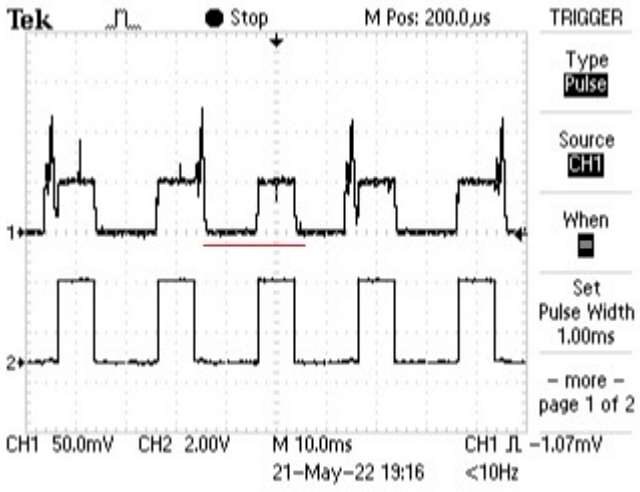

5500 Encrypted.xlsx

5500 Encrypted

A-6 SUCK_DELAY = 8000_H ($32{,}768_D$) Unencrypted

In this measurement, a 20 ms cycle consumed 525 nJ of energy.

Microsoft Excel
97-2003 Worksheet

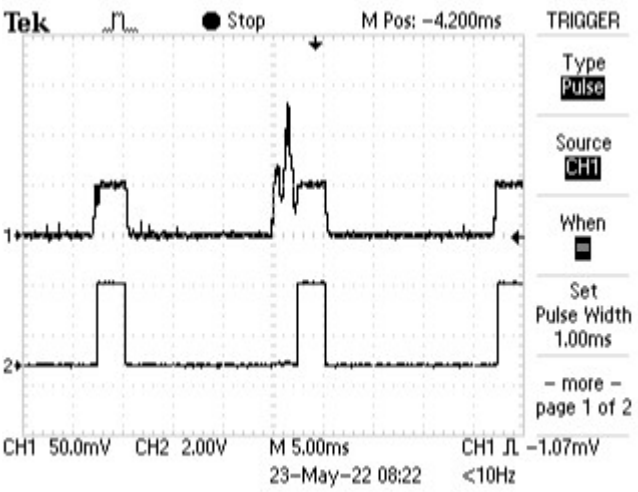

8000 Unencrypted

A-7 SUCK_DELAY = 8000_H ($32{,}768_D$) Encrypted

In this measurement, a 20 ms cycle consumed 514 nJ of energy.

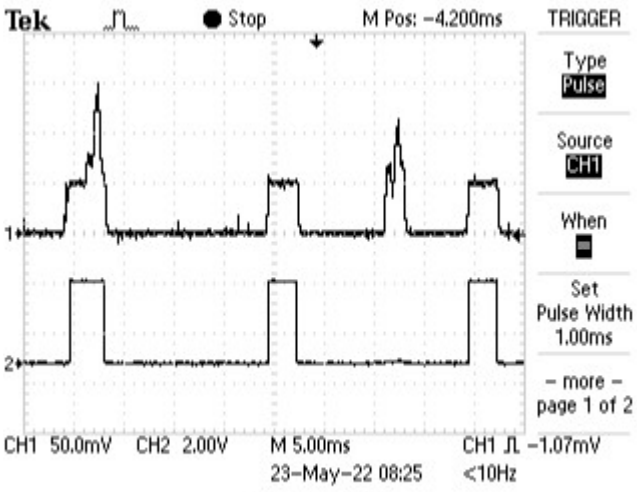

8000 Encrypted.xlsx

8000 Encrypted

A-8 SUCK_DELAY = $9a00_H$ ($39{,}424_D$) Unencrypted

In this measurement, a 20 ms cycle consumed 2,274 nJ.

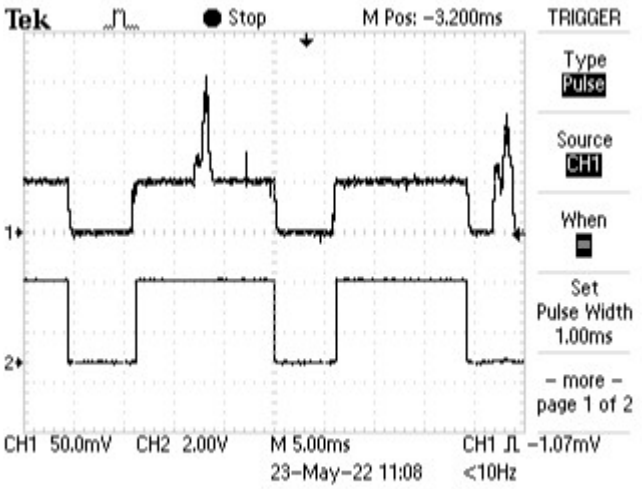

9a00 Unencrypted

Microsoft Excel
97-2003 Worksheet

A-9 SUCK_DELAY = $9a00_H$ ($39{,}424_D$) Encrypted

In this measurement, a 20 ms cycle consumed 2277 nJ.

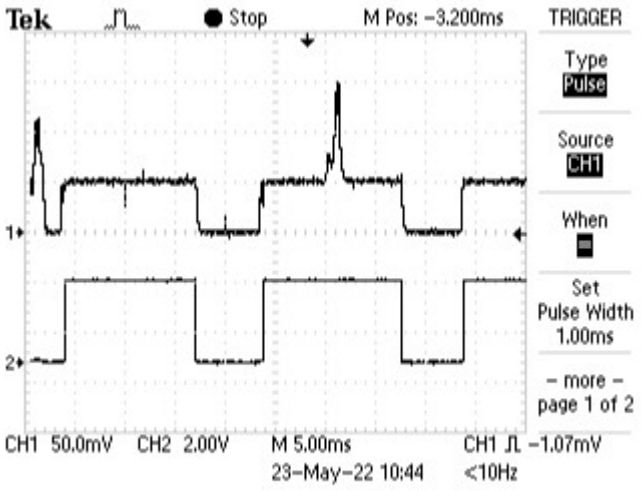

9a00 Encrypted

Microsoft Excel
97-2003 Worksheet

B Published Papers

This Appendix includes published papers by the Author and an industry framework for which the Author was an Editor.

B-1 Privacy and security of consumer IoT devices for the pervasive monitoring of vulnerable people

The following paper was published in the proceedings of the "Living in the Internet of Things: Cybersecurity of the IoT" at the IET, London, 28-29th March, 2018 (Poyner and Sherratt, 2018).
(Double-click to open in pdf reader).

Privacy and security of consumer IoT device
monitoring of vulnerable peo

I K Poyner, R S Sherratt**

** Biomedical Engineering, University of Reading, UK, i.poyner@pgr.reading.ac*

Keywords: IoT, healthcare, security, privacy, encryption.

Abstract

The Internet of Things (IoT) promises highly innovative solutions to a wide range of activities. However, simply being a technology company does not exempt an IoT company from needing to comply with the legislation applicable to their operating region that safeguards personal information. This will result in security and privacy requirements for healthcare solutions. There are several mature frameworks that address these issues, but they have been developed within the context of organised hospitals and care providers, where there is the expertise, processing power, communications and electrical power to support highly robust security. However, for IoT solutions aimed at vulnerable people, either at home or within their local environment, there are significant additional constraints that must be overcome. These include technical (low processing capability, power constrained, intermittent communications) organisational (how to enrol and revoke users and devices, distribution of cryptographic keys) and user constraints (how does a patient with physical and/or mental challenges configure and update their devices).

This paper considers at the legal frameworks and the security and privacy requirements for healthcare solutions. An overview of some of the primary frameworks is then provided followed by an assessment of how this is constrained within an IoT system.

Patients' concerns over
also a barrier to the ado

Failure to consider safe
project could result in
deployed devices, or i
fines (potentially £17M

Today, many of the sta
aimed towards patients
where physical, pers
contribute to the overall
monitoring systems i
maintained by comp
constrained in terms of
local processing capabil
the resource constraine
people interacting wit
home. Another compli
and apps aimed at digit
are re-purposed for m
interaction with techn
vulnerable people are
important protective be
and changing default p
to falling victim to soci

This paper assesses th
security models in the

B-2 Improving access to healthcare in rural communities — IoT as part of the solution

The following paper “Improving access to healthcare in rural communities - IoT as part of the solution”, was the Opening Presentation at the 3rd IET International Conference on Technologies for Active and Assisted Living (TechAAL 2019) (Poyner and Sherratt, 2019).

(Double-click to open in pdf reader).

Improving access to healthcare in rural comm of the solution

I. K. Poyner, R.S. Sherratt†*

** Department of Biomedical Engineering, University of Reading, RG6 6AY (i*

† Department of Biomedical Engineering, University of Reading, RG6 6AY (

Keywords: IoT, healthcare, remote, rural, constrained.

Abstract

AAL has benefitted tremendously from the near-ubiquity of powerful smartphones and very high data rates available over broadband and mobile networks. However, this is beyond the reach of many users. IoT systems offer the potential to extend some of the benefits to disadvantaged users. Such solutions will need to secure personal health information and provide a sufficient quality of service even when operating constrained user devices and communications.

1 Introduction

Technologies such as low-cost computers and broadband communications are enabling novel healthcare services that can help people, young and old, reduce the impact of chronic conditions, such as diabetes, cognitive challenges and dementia. In some cases, a user may share information in a controlled manner with other people who can provide support, such as family members, carers and organisations who have legal safeguarding obligations to user, such as social services.

However, communities in rural areas or developing countries may be inadvertently excluded from these transformative benefits because communications are not available, systems are unaffordable, or services have not been regionalised to make them easier for a user to understand in their native language.

where power is not a
Rwanda 92% of the po
internet should be affo
(based upon the assump
expenditure), yet only
internet. The Governm
computer literacy is a
Nation" (less than 9%
literate) [4].

IoT (Internet of Things
these issues. Low cost
simple (disappearing) u
than smartphones. Wi
range of 10-15km and
enable communities to
benefits being deployed

In this paper we look
services, requirements
by traditional technol
address some of the cor

2 Use cases

Four use cases are cons
1. Elderly person cop
a rural environmen
may occasionally b
surroundings or po
2. A teenager with Ty
expedition in a wil

B-3 IoT Security Assurance Framework

The IoTSF (IoT Security Foundation) Assurance Framework (IoT Security Foundation, 2021) is a framework aimed at promoting security in IoT products and services. Each section is tailored for a group of readers, whether they be consumers, software developers, hardware designers, cryptography custodians or business managers. It has been downloaded over 10,000 times since it was first issued in 2017. The author of this paper was heavily involved as an Editor during the Release 3.0, which involved significant changes as the focus moved from being a compliance document to an assurance framework. The publicly-available Framework document is included here, although the author was also heavily involved in the more detailed Questionnaire spreadsheet that provides assistance as to 'how' the requirements could be satisfied.
(Double-click to open in pdf reader).