SAPIENZA
UNIVERSITÀ DI ROMA

Privacy in characterizing and recruiting patients for IoHT-aided digital clinical trials

Sapienza - University of Rome

Ph.D. in Engineering in Computer Science – XXXI Cycle

Candidate

Fabio Angeletti

ID number 1244851

Thesis Advisor

Prof. Andrea Vitaletti

Co-Advisor

Prof. Leonardo Querzoni

January 2019

Thesis defended on 22nd February 2019
in front of a Board of Examiners composed by:

Prof. Riccardo Torlone (chairman)

Prof. Alessandro Farinelli

Prof. Paolo Prinetto

Privacy in characterizing and recruiting patients for IoHT-aided digital clinical trials

Ph.D. thesis. Sapienza – University of Rome

© 2019 Fabio Angeletti. All rights reserved

This thesis has been typeset by LaTeX and the Sapthesis class.

Version: February 23, 2019

Author's email: angeletti@diag.uniroma1.it

Dedicated to
my sister Carolina
my mother Lorella
my father Maurizio
my girlfriend Tatiana
Your unconditional love and support
made all of this possible.

Abstract

Nowadays there is a tremendous amount of smart and connected devices that produce data. The so-called IoT is so pervasive that its devices (in particular the ones that we take with us during all the day - wearables, smartphones...) often provide some insights on our lives to third parties. People habitually exchange some of their private data in order to obtain services, discounts and advantages. Sharing personal data is commonly accepted in contexts like social networks but individuals suddenly become more than concerned if a third party is interested in accessing personal health data. The healthcare systems worldwide, however, begun to take advantage of the data produced by eHealth solutions. It is clear that while on one hand the technology proved to be a great ally in the modern medicine and can lead to notable benefits, on the other hand these processes pose serious threats to our privacy.

The process of testing, validating and putting on the market a new drug or medical treatment is called *clinical trial*. These trials are deeply impacted by the technological advancements and greatly benefit from the use of eHealth solutions. The *clinical research institutes* are the entities in charge of leading the trials and need to access as much health data of the patients as possible. However, at any phase of a clinical trial, the personal information of the participants should be preserved and maintained private as long as possible.

During this thesis, we will introduce an architecture that protects the privacy of personal data during the first phases of digital clinical trials (namely the *characterization phase* and *the recruiting phase*), allowing potential participants to freely join trials without disclosing their personal health information without a proper reward and/or prior agreement. We will illustrate what is the *trusted* environment that is the most used approach in eHealth and, later, we will dig into the *untrusted* environment where the concept of privacy is more challenging to protect while maintaining usability of data. Our architecture maintains the individuals in full control over the flow of their personal health data. Moreover, the architecture allows the *clinical research institutes* to characterize the population of potentiant users without direct access to their personal data.

We validated our architecture with a proof of concept that includes all the involved entities from the low level hardware up to the end application. We designed and realized the hardware capable of sensing, processing and transmitting personal health data in a privacy preserving fashion that requires little to none maintenance.

Table of Contents

List of Tables

List of Figures

Achievements

- **Articles on journals**
  [Sensors&Transducers 2015] *"A Modular Design for Wireless Structural Health Monitoring Applications"* - **Fabio Angeletti**, Mario Paoli, Ugo Maria Colesanti, Andrea Vitaletti
  [Engineering Structures 2018] *"Ropeway Roller Batteries Dynamics: Modeling, Identification, and Full-Scale Validation"* - Andrea Arena, Biagio Carboni, **Fabio Angeletti**, Mathieu Babaz, Walter Lacarbonara
  [Sensors 2018] *"Towards an architecture to guarantee both data privacy and utility in the first phases of Digital Clinical Trials"* - **Fabio Angeletti**, Ioannis Chatzigiannakis, Andrea Vitaletti

- **Conferences**
  [The Ninth International Conference on Sensor Technologies and Applications - SENSORCOMM 2015] *"Wireless Sensor Networks in Structural Health Monitoring: a Modular Approach"* - **Fabio Angeletti**, Mario Paoli, Ugo Maria Colesanti, Andrea Vitaletti
  [The 25th International Conference on Software, Telecommunications and Computer Networks - SOFTCOM 2017] *"The role of blockchain and IoT in recruiting participants for digital clinical trials"* - **Fabio Angeletti**, Ioannis Chatzigiannakis, Andrea Vitaletti

- **Workshops**
  [The First International Workshop on Human-centered Sensing, Networking, and Systems - HUMANSYS 2017] *"Privacy preserving data management in recruiting participants for digital clinical trials"* - **Fabio Angeletti**, Ioannis Chatzigiannakis, Andrea Vitaletti
  [ACM 1st International Workshop on Knowledge Management for Healthcare (KMH) - KMH 2018] *"Random Projection to Preserve Patient Privacy"* - Aris Anagnostopoulos, **Fabio Angeletti**, Federico Arcangeli, Chris Schwiegelshohn, Andrea Vitaletti

- **Posters**
  [First International Conference on Network Medicine and Big Data: The Transformation of Medicine 2018] *"Random Projection to Preserve Patient Privacy"* - Aris Anagnostopoulos, **Fabio Angeletti**, Federico Arcangeli, Chris Schwiegelshohn, Andrea Vitaletti

Chapter 1

Introduction

During the last two decades we witnessed the tremendous progresses made in the field of wireless sensor networks, from the reduction in size to the enhancement in computation capabilities, from the coverage improvement to the reduced power consumption and the reduction of hardware costs. These advancements paved the way and facilitated the wider adoption of small electronic devices with interconnection capabilities. The so-called *Internet of Things* (IoT) is mainly composed by these devices. The IoT is taking a central role into new generation solutions that orchestrate myriads of devices, web services, business processes, people, companies and institutions.

The IoT is a highly dynamic and radically distributed networked system, composed of an incredible high number of objects [**?**]. The ubiquitous and pervasive devices composing the IoT are sensors, actuators, wireless display, intelligent appliaces, and, generally speaking, *connected* objects. They all generate a tremendous amount of data [**?**]. Robust, available and fast infrastructures that also offer storage solutions capable of handling this flow of data are needed. Moreover, the amount of this data builds the bases to very effective and powerful algorithms from the fields of machine learning and data mining [**?**]. Both IoT and the usage of its data are vastly considered as ones of the most expanding areas within future technologies and it is attracting attention in different industry applications [**?**], ranging from smart cities to home automation, from farming to precise agriculture, from manufacturing to healthcare and many more.

In this dissertation, we focused on a particular domain where the coexistence and cooperation of embedded systems (thus, the IoT) with our social life is unveiling a brand new era of exciting possibilities, the *eHealth* (e.g. Figure **??**).

Medicine has existed for thousands of years, during most of which it was an art frequently having connections to the religious and philosophical beliefs of local culture. In recent centuries, since the advent of modern science, most medicine has become a combination of art and science. While stitching technique for sutures is an art learned through practice, the knowledge of what happens at the cellular and molecular level in the tissues being stitched arises through science. From our more recent and more ortodox point of view, medicine is the science and practice of the diagnosis, treatment, and prevention of disease. It encompasses a variety of health care practices evolved to maintain and restore health by the prevention and

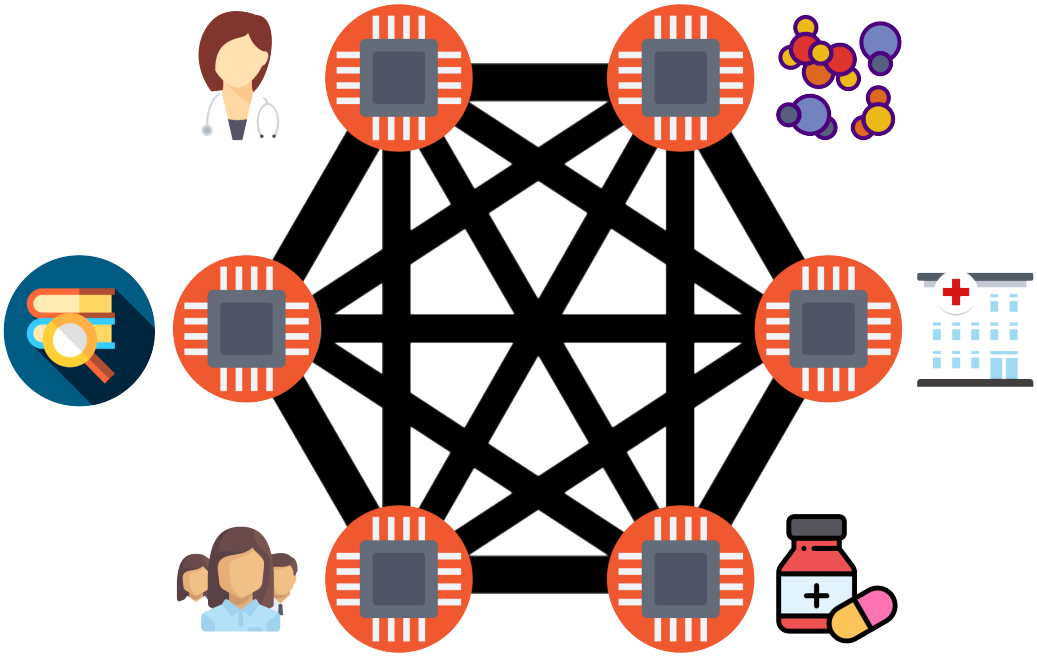

Figure 1.1. Schematic interconnection of entities involved in *eHealth*. Icons made by [prettycons, smashicons, prosymbols, freepik, vectors market] from www.flaticon.com

treatment of illness. Contemporary medicine applies biomedical sciences, biomedical research, genetics, and medical technology to diagnose, treat, and prevent injury and disease, typically through pharmaceuticals or surgery, but also through therapies as diverse as psychotherapy, external splints and traction, medical devices, biologics, and ionizing radiation, amongst others.

It was only a matter of time before the IoT entered the healthcare world, giving birth to eHealth. With the constantly increasing amount of data that is inherent to an IoT world dedicated to the healthcare, all parties involved can benefit. For instance, just to name a few, physicians can obtain better and refined update regarding a developing disease, hospitals can adjust their beds management in order to react faster to emergencies, drug developers can potentially access a wealth of real-world, participant-generated data that is enabling better insights and streamlined clinical trial processes.
Electronics healthcare solutions and, generally speaking, the *Internet of Health Things (IoHT)* follow the same trend as other IoT solutions: they are adopted with steadily more frequency. About 73% of healthcare executives say that IoHT is a disrupting techology for the next years and it is already becoming one of the most funded areas in IoT [**?**]. In the field of healthcare, common devices that are nowadays upgraded with connectivity and "smart features" include watches, weights, phones, clocks and general purpose instrumentation such as pressure monitors. Most of them are equipped with biometrical sensors or are capable of measuring biological parameter, such as heart rate, apnea, sleep cycles, body temperature, blood pressure and oxygenation, body weight, and many more. These devices that gather data

regarding a specific individual open the way to *precision medicine*. It aims to adapt the treatment to the specific individual, thus needs to access as much information as possible regarding the patient.
Pervasive IoHT enables cost savings for both the administrations and the individuals, but on the other hand it has some barriers, like privacy and security concerns, lack of skilled workers, poor interoperability and more [**?**]. The concepts of "privacy" and "protection of personal information" are constantly undermined in our society because of their value and malicious third parties are posing threats. Data breaches [**?**, **?**, **?**] are on the rise and we need techniques to control and protect the users' privacy in an effective and scalable way.

Given the sensitive nature of healthcare data, there is a strong need to protect the privacy of the patients from third parties. The European Commission, furthermore, recently enforced the General Data Protection Regulation (GDPR) in order to strengthen data protection and now this regulation must be applied to any organisation or individual that collects and processes information related to EU citizens, regardless where the data is physically stored or where they are based[**?**, **?**]. On the other hand and at the same time, analysis of such data are crucial for both the medical research and the drug industry. Consequently, there is a need to design approaches that allow data processing without exposing the personal underlying information, with a particular focus on the chance of re-identification of individuals from their data.

All the eHealth processes greatly benefit from accessing as much data as possible. The privacy preserving approaches require to guarantee privacy to the users and the techniques involved require an extra effort in terms of computation, energy and badwidth. This is an obstacle in wearable devices where there are hard constraints in processing power, communication coverage and bandwidth and energy reserve. Luckly we will see later how to overcome these limitations.

1.1 Clinical trials

Clinical trials are experiments or observations done in clinical research. Such prospective biomedical or behavioral research studies on human participants are designed to answer specific questions about biomedical or behavioral interventions, including new treatments (such as novel vaccines, drugs, dietary choices, dietary supplements, and medical devices) and known interventions that warrant further study and comparison. Clinical trials generate data on safety and efficacy. They are conducted only after they have received health authority/ethics committee approval in the country where approval of the therapy is sought. These authorities are responsible for vetting the risk/benefit ratio of the trial – their approval does not mean that the therapy is 'safe' or effective, only that the trial may be conducted [**?**].

Developing drugs is a challenging process. Only around one in 10 drugs in development actually make it through to the market [**?**, **?**]. This low rate to enter

the market is one factor contributing to the high costs of drug development and also the slow adoption of new drugs. A recent study indicates that developing a drug from bench to market costs an estimated $2.6 billion [**?**]. A large portion of those costs is related to:

- recruiting an adequate number of patients
- retaining the patients throughout the trials

Currently, more than $244,000$ studies are registered in the world out of which more than $42,000$ are currently recruiting [**?**]. Some of these studies require hundreds of individuals but others require thousands of participants, each of whom must meet precise criteria in order to be fit in a particular trial. Thus, it is not surprising that 80% of these important studies are delayed due to recruitment problems, according to the Center for Information and Study on Clinical Research Participation (CISCRP) [**?**]. Long recruitment phases prolong the execution of trials thus taking longer for innovative new medicines to be studied and approved, leaving patients to wait years for new treatment options.

Depending on product type and development stage, investigators initially enroll volunteers and/or patients into small pilot studies, and subsequently conduct progressively larger scale comparative studies. Clinical trials can vary in size and cost, and they can involve a single research center or multiple centers, in one country or in multiple countries. Clinical study design aims to ensure the scientific validity and reproducibility of the results. Trials can be quite costly, depending on a number of factors. The sponsor may be a governmental organization or a pharmaceutical, biotechnology or medical device company. Certain functions necessary to the trial, such as monitoring and lab work, may be managed by an outsourced partner, such as a contract research organization or a central laboratory.

The clinical trials are affected by a number of uncertainties driven by issues that could arise during the clinical trial itself. the number of patients required by each step of a clinical trial need to grow quite exponentially during the process. Enrolled patients could be less prone to participate if the trial is too cumbersome to sustain (it may require movements on daily basis to reach a specific installation, or following too much precautions and guide lines). While there exist trial in which data can be take remotely (avoiding then the physical movement of patients), often the *clinical research instutions* need to verify the reliability of this data. It can be costly and sometimes not feasible because of the accuracy of the instrumentation (or the cost of it).

Another aspect to consider is that currently the recruiting of patients is made by-hand by nurses or, generally, human operators. As in any human relation and work, this process is prone to errors. This need to access an appropriate pool of patients in order to execute clinical trials is well known to the broader public. It is observed on multiple occasions: a growing pool of people willing to participate. In a 2015 study of CISCRP [**?**] identified that 81% of responders consider clinical research studies “very important” to the discovery and development of new medicines and 80% of them would be willing to participate in a research study.

According to a 2012 on-line survey [?], 85% of the responders perceive privacy concerns as a barrier to share health information. It is clear that collected data may be used to extract or infer sensitive information about users' private lives, habits, activities and relations, which all refer to individuals' privacy [?, ?]. About half of the responders were either concerned or very concerned about the re-identification of their anonymized health and medical information. If data were irreversibly anonymized, 71% of respondents were willing to share data with researchers. During a clinical trial recruiting phase, when the benefits of a possible future enrollment have not been fully clarified, patients expect that their medical condition information are kept confidential.

Given the strict qualification criteria imposed by the researchers, only about 5% of patients eventually constitute the group participating in clinical trials. It is, therefore, imperative to introduce new methods for facilitating recruitment that respects the privacy and confidentiality of the patients in order to maximize the participation of people - particularly in rare diseases where the communities of patients are small.

During the execution of the trials, the collection of high-quality data is absolutely vital. For this reason, trial centres require regular tests and observations to be conducted at their premises in order to guarantee the accuracy of data collection. Interestingly, 70% of potential participants live more than two hours away from the nearest study centre [?]. It is, therefore, common for patients to travel to those centres for regular tests and observations, sometimes several times each week for the duration of the trial. Such complexities sometimes overcome the perceived benefits of participating in a trial, inevitably increasing the attrition rate of patients. Clearly, redoing patient recruitment further delays the execution of the trials.

Understanding the above issues and addressing them adequately is critical in developing successful digital health solutions. As technology becomes more accessible and affordable, the role of digital health data will become vital in clinical trials. It is well known that smartphones are already a ubiquitous technology – in 2015, almost two-thirds of people in the U.S. owned a smartphone and almost half owned a tablet [?]. During the same year, about 300 clinical trials were reported to involve a wearable technology [?].

The first ever reported clinical trial was conducted by James Lind in 1747 [?], later on the biggest change in clinical trial was the introduction of randomized experiments [?, ?, ?, ?]. By the late 20th century, RCTs were recognized as the standard method for "rational therapeutic" in medicine [?]. To improve the reporting of RCTs in the medical literature, an international group of scientists and editors published Consolidated Standards of Reporting Trials (CONSORT) Statements in 1996, 2001 and 2010, and these have become widely accepted [?, ?].

A randomized controlled trial is a type of scientific experiment which aims to reduce bias when testing a new treatment. The people participating in the trial are randomly allocated to either the group receiving the treatment under investigation or to a group receiving standard treatment (or placebo treatment) as the control. Randomization minimises selection bias and the different comparison groups allow the researchers to determine any effects of the treatment when compared with the no treatment (control) group, while other variables are kept constant. The RCT is

often considered the gold standard for a clinical trial. RCTs are often used to test the efficacy or effectiveness of various types of medical intervention and may provide information about adverse effects, such as drug reactions. Random assignment of intervention is done after subjects have been assessed for eligibility and recruited, but before the intervention to be studied begins.

Random allocation in real trials is complex, but conceptually the process is like tossing a coin. After randomization, the two (or more) groups of subjects are followed in exactly the same way and the only differences between them is the care they receive. For example, in terms of procedures, tests, outpatient visits, and follow-up calls, should be those intrinsic to the treatments being compared. The most important advantage of proper randomization is that it minimizes allocation bias, balancing both known and unknown prognostic factors, in the assignment of treatments [**?**].

Clinical trials need people to enroll. Participation of possible candidates is one of the prominent issues. Steve Cutler, chief operating officer of clinical research organisation ICON says: "Five percent of patients is probably the higher end of the range for participation in trials. I'd like to see that get to 20% within my lifetime, and I think we have the opportunity and the potential to get there." Kevin Julian, managing director of Accenture adds: "Another exciting development is the ability to bring specific patient profiles into the recruitment process – not just connecting patients to trials that they might be interested in but digitally prescreening patients for the inclusion-exclusion criteria on a trial and learning early on if a patient is a candidate, or perhaps in a particular population seeing whether there are even enough patients that meet the criteria. It's better to know ahead of time than work it out later by trial and error."

Many trials do not recruit sufficient participants and this can make it more difficult to use the results of the research in practice. Effective strategies for improving recruitment would be of great benefit to researchers designing and running trials [**?**]. Some groups of individuals are underserved by the medical health-care system in recruitment for clinical trials, thus it is difficult to include them [**?**, **?**, **?**, **?**]. Telemedicine helps reaching out more people, as explained in [**?**, **?**]. At the opposite of the spectrum, instead, there are trials in which the possible participants represent a very large number. Here we need a way to reduce this set to a smaller one that has the highest fitting to the trial. A solution was proposed in [**?**], the authors reduce the number of participants using electronic screening before recruiting in a clinical trial.

1.1.1 Phases of a clinical trial

According to [**?**], a clinical trial of experimental drug, treatment, device or behavioral intervention may proceed through four phases. To help better understand them, a graphic is provided in the following Figure **??**.

Going deeper, it is necessary to highlight the existence of another base step, the **Phase 0**. Finally, we can structurate a clinical trail as a path through five phases as reported below [**?**].

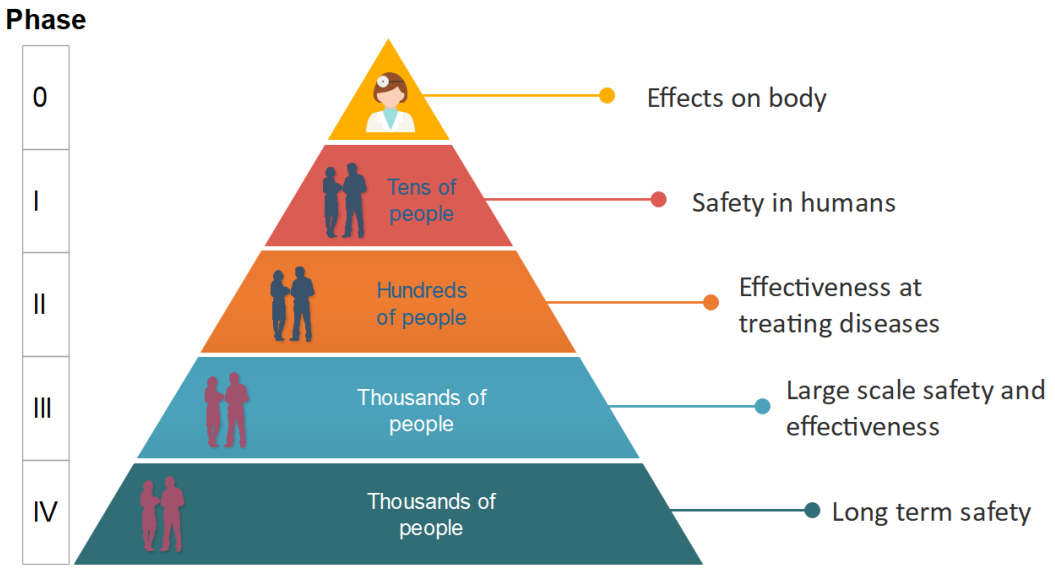

Figure 1.2. Flow diagram to show the different phases in a clinical trial. Icons made by [prettycons, smashicons, prosymbols, freepik, vectors market] from www.flaticon.com

- **Phase 0:** the object of the clinical trial (such as a treatment, a device, an experimental drug or a behavioral intervation) is thoroughly studied. This phase in particular could be years long.

- **Phase I:** *Clinical Pharmacology and Toxicity.* These first experiments are primarily concerned with drug safety, not efficacy, and hence are usually performed on human volunteers, often pharmaceutical company employees. The first objective is to determine an acceptable single drug dosage (i.e. how much drug can be given without causing serious side-effects). Such information is often obtained from dose-escalation experiments, whereby a volunteer is subjected to increasing doses of the drug according to a predetermined schedule. This phase will also involve studies of drug metabolism and bioavailability. After studies in normal volunteers, the initial trials in patients will take place. Typically, the studies in this phase might require a total of around 20-80 subjects and patients.

- **Phase II:** *Initial Clinical Investigation for Treatment Effect.* These are fairly small-scale investigations into the effectiveness and safety of a drug, and require close monitoring of each patient. The trials in this phase can sometimes be set up as a screening process to select out those relatively few drugs of genuine potential from the larger number of drugs which are inactive or over-toxic, so that the chosen drugs may proceed to the trials in the following phases. Seldom will this phase go beyond 100-200 patients on a drug.

- **Phase III:** *Full-scale Evaluation of Treatment.* After a drug is shown to be reasonably effective, it is essential to compare it with the current standard treatment(s) for the same condition in a large trial involving a substantial number of patients. To some people the term 'clinical trial' is synonymous with

such a full-scale trial, which is the most rigorous and extensive type of scientific clinical investigation of a new treatment. Studies investigate the efficacy of the biomedical or behavioral intervention in large groups of human subjects (from several hundred to several thousand) by comparing the intervention to other standard or experimental interventions as well as to monitor adverse effects, and to collect information that will allow the intervention to be used safely.

- **Phase IV:** *Postmarketing Surveillance.* After the research programme leading to a drug being approved for marketing, there remain substantial enquiries still to be undertaken as regards monitoring for adverse effects and additional large-scale, long-term studies of morbidity and mortality. This phase is sometimes used to describe promotion exercises aimed at bringing a new drug to the attention of a large number of clinicians, typically in general practice. This latter type of enquiry has limited scientific value and hence should not be considered part of clinical trial research.

1.2 Healthcare and IoT, the birth of IoHT

The term *eHealth* encompasses a broad range of meanings. It is used by many individuals, institutions and organizations both in academy and outside. Still, it is a neologism that lacks of a clear definition [**?**]. Generally speaking, it is a relatively recent term for healthcare practice supported by electronic processes and communication [**?**], we support this definition through our works. Cleared out what the term means by itself, we can focus on describing what is the *eHealth*, and what it changed in the medical field.

Since the beginning of the XXI century [**?**], different branches of medicine joined the use of telecommunication and information technology to provide the so-called *telemedicine* [**?**, **?**]. With the explosion of the IoT in the market [**?**], a large number of devices with internet connectivity started to work among us.
Millions of these devices are sold every year [**?**], following a trend that does not seem to diminish. Thus, a vast subset of the world population has access to low cost devices capable of collecting personal health data, like heart rate, steps, blood pressure, pulse oximetry, sleep cycles, body weight and composition and much much more [**?**].
The tremendous amount of data generated by these devices could be exploited to help in various field of healthcare, from early diagnosis of illness to more effective treatment options, from individual-targeted therapies to behavioral analysis. The advancements in machine learning and data mining techniques can really enhance the performances of current health solutions as well as exploit the data to provide cutting edge new paths for treatment and diagnosis. The interest in using advances in technology in conjuction with medicine is clear [**?**, **?**].

Each aspect of our lives is impacted (or will be impacted in the nearest future), in particular, regarding the health processes: from the way we wake up each morning to counting how many steps we made each day, from recording our calory intake meal after meal to storing our body composition and blood pressure. We were used to wake up using an alarm clock, then it was the time of the radio alarm clocks,

now some devices can wake up us at a variable time depending on our sleep stage [**?**]. The sphygmomanometer, symbol of modern medicine (with, of course, the stethoscope) is now often replaced with electronic alternatives that also measure heart rate in few seconds. A large number of commonly used devices in medicine are now replaced with more technological alternativse. These newly designed devices try to minimize reading errors, human mistakes, reliability and offer newer functions and connectivity. Thinking about the old fashion thermometer that required more than 5 minutes to obtain a reading, is it now normal to image the modern infrared-based or thermocouple-based alternatives that are extremely faster and easier to read thanks to the LCD (Liquid Crystal Display). Modern activity trackers can calculate an incredible range of human body parameters: heart rate, calory burn, pulse oximetry, steps and also blood pressure.

Remote monitoring of human behaviour and condition is well known in literature [**?**, **?**, **?**, **?**, **?**]. Telemedicine can be seen as a branch of information technology that tries to address the issues of medicine coming from distances between patients and specialized structures, such as costs, time to intervention, diagnosis and more. Telemedicine helps all the individuals with no distinction between gender, age or social status.

We live in a society with an increasing number of elder people [**?**, **?**, **?**]. As reported before, the elderly have high need of assistance. This high demand is not fullfilled completely by the available physician and specialized operators [**?**, **?**, **?**, **?**]. It is well understood that the elderly are the most vulnerable to illnesses [**?**, **?**] and studying them throughly could really help the research. The elderly are more prone to disabilities, so that they could be primary beneficiaries from telemedicine [**?**]. "Today, the number of elderly people and patients with reduced autonomy or with chronic diseases are steadily increasing. In addition, a stay in hospital or nursing home is very expensive. Thus, in recent years we have witnessed the development of projects to keep these people at home while providing them the needed care and assistance" [**?**]. However, it is not new that the interaction between elderly people and modern technology must be taken into account [**?**, **?**, **?**, **?**]. Nonetheless, the research community has made important progresses in that direction [**?**, **?**].

The concept of medical tourism [**?**] taken a lot of attention from the community [**?**, **?**, **?**]. With the use of eHealth, this phenomenon can be greatly reduced.Moreover, social status deeply impact the accessibility of individuals to the healthcare system (both public and private) [**?**, **?**] and it is demonstrated that different educations lead to the same result [**?**, **?**]. Telemedicine can help on all in all of this matters, addressing multiple issues together [**?**, **?**, **?**, **?**].

eHealth technologies are now underway internationally, often with an incredible high amount of money used as investment. Governments believed in this technology. For example, England has invested at least 12.8 billion pounds in a National Programme for Information Technology (NPfIT) for the National Health Service, while the the Obama administration in the United States (US) has used more than 37 billion dollars in health care with eHealth[**?**]. "The eHealth IMPACT study provides

empirical evidence on the benefits of eHealth systems and services. It demonstrates the potential of eHealth as enabling tool for meeting the 'grand challenges' of European health delivery systems" [**?**]. Among other solutions, the same concept of clinical trial could be perfectioned using all the data gathered by the health-centric IoT devices and the modern machine learning techniques.

1.3 Digital clinical trial

Now it should be clear that the collected data from IoHT devices could improve the clinical trial process. It can provide advantages in all the steps that involve the participants, from the recruiting to the enrollment in the program, from the monitoring during the trial to the share of personal data to third parties. The whole data, as soon as it is digitalized, become easily accessible and the internet could help its sharing even further. The strong connection now possible between the clinical trials and the information technology field paved the way to *digital clinical trials*. In this revisited clinical trials, the whole process is supported by recent technologies that allows to reduce the gap between physicians and patients and between clinical research institutions. Moreover, it can particularly help in all the situations in which the physical distances between the *clinical research institutes* and the patients are considered strong barriers. In these cases, telemedice and eHealth solutions can be the only feasible options for a person to participate in a trial (e.g. if he/she is mobility impaired or reaching the clinical research institution requires taking multiple transportation systems). The wide adoption of technology changes the way each phase of a clinical trail is accomplished.

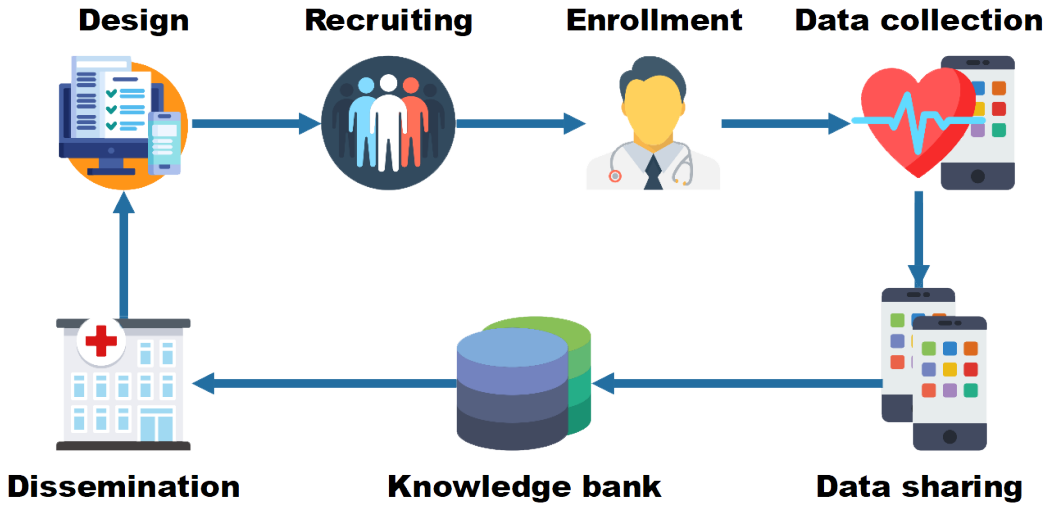

Figure 1.3. The generic iterative phase suggested in digital clinical trials. Icons made by [prettycons, smashicons, prosymbols, freepik, vectors market] from www.flaticon.com

At first, all but the **Phase 0** depicted in Figure **??** are deeply affected by the use of health-related IoT devices. Since we are now in a faster communicating

environment, it is possible to adjust requirements, steps to follow, control parameter and much more at any time. Thus, these phases can be made iterative. This particular approach poses the bases for the digital clinical trial, of which each phase could follow the iteration presented in Figure **??**.

The model is based on the ability of modern technologies to communicate over the Internet in order to reach nearly an unlimited number of potential participants. The model also leverages the capabilities of IoT to collect relevant data for a specific study employing suitable devices deployed outside the "study centres", possibly at participants' homes. Initially, a digital clinical trial is **designed**, possibly taking into account insights from previous trials. In the sequel, during the **recruiting phase**, patients matching the needs for the designed trial are recruited. These patients are required to **enrol** upon accepting the informed consent. Enrolled patients **remotely collect useful data** and **share those data** to build a **knowledge bank**. The outcomes of the trial are **disseminated** and can contribute to the design of a new trial. The cycle then can repeat itself several times, possibly improving each phase. The advantages are obvious, it is possible to better target a specific group of people as well as reach previously unreachable individuals, also it is possible to change parameter on-the-fly and have access to much more information then before, also in a real-time fashion. Moreover, all the data is digitized. This implies that is easier to share. Also the results accomplished by different organizations are fastly available, and this makes the refining iterative process simpler and faster.

1.4 eHealth and privacy

Digitalization of clinical trials poses some interrogatives to answer to. How is my data managed? Who is in charge of controlling the flow of information and how he can guarantee no data breaches? **Privacy** must be really taken into account by both possible candidates and third parties. This matter is well-known in the literature [**?**, **?**, **?**, **?**, **?**, **?**]. The motives are different but they require strong answers to address the matter and make the potential candidates comfortable to participate without fears. Privacy of information collected during healthcare processes is necessary because of significant economic, psychologic, and social harm that can come to individuals when personal health information is disclosed [**?**, **?**].

Privacy of information collected during healthcare processes is necessary because of significant economic, psychologic, and social harm that can come to individuals when personal health information is disclosed [**?**, **?**]. Multiple definitions of privacy exist, each one focused on different declinations of the same principle: *"the ones right to manage valuable personal information"*. Certain studies account some critical points regarding privacy: improper access, unauthorized use (both direct or secondary), errors and collection of personal information [**?**, **?**, **?**, **?**].

Consider the following three different definitions of *privacy* that help convey the implication of the concept of privacy within IT applications and services [**?**]:

- "Privacy is the claim of individuals, groups or institutions to determine for

themselves when, how, and to what extent information about them is communicated to others." [?]

- "[Privacy is] the appropriate use of personal information under the circumstances. What is appropriate will depend on context, law, and the individual's expectations; also, [privacy is] the right of an individual to control the collection, use, and disclosure of personal information." [?]

- "Privacy involves the policies, procedures, and other controls that determine which personal information is collected, how it is used, with whom it is shared, and how individuals who are the subject of that information are informed and involved in this process." [?]

The above definitions indicate that privacy includes diverse concepts however they all accept the emerging need to manage and protect the personal information adequately. Essentially the goal of information privacy raises issues of access control (user authentication and authorization) and the need for data authentication.

Accordingly, the American Civil Liberties Union (ACLU) believes that a privacy policy for health information should be based on the following principles [?, ?]:

1. Strict limits on access and disclosure must apply to all personally identifiable health data, regardless of the form in which the information is maintained.

2. All personally identifiable health records must be under an individual's control. No personal information may be disclosed without an individual's uncoerced, informed consent.

3. Health-record information systems must be required to build in security measures to protect personal information against both unauthorized access and misuse by authorized users.

4. Employers must be denied access to personally identifiable health information on their employees and prospective employees.

5. Patients must be given notice of all uses of their health information.

6. Individuals must have a right of access to their own medical and financial records, including rights to copy and correct any and all information contained in those records.

7. Both a private right of action and a governmental enforcement mechanism must be established to prevent or remedy wrongful disclosures or other misuses of information.

8. A federal oversight system must be established to ensure compliance with privacy laws and regulations.

1.4.1 Privacy and security

In a digital health system, all information is converted into a digital form. Therefore data protection and privacy protection are very closely connected. This is due to the need of transmitting, processing and utilizing the information in fast, simple and reliable ways, thus in digital ways. It is, therefore, crucial to agree on a common definition of privacy, how it is related to security and how to identify and address privacy risks.

In terms of a healthcare information security system, the goals of the system in terms of security can be simply stated as follows [**?**, **?**, **?**].

1. To ensure the privacy of patients and the confidentiality of health care data (prevention of unauthorized disclosure of information).
2. To ensure the integrity of healthcare data (prevention of unauthorized modification of information).
3. To ensure the availability of health data for authorized persons (prevention of unauthorized or unintended withholding of information or resources).

In this sense, the goal of security is the application of cryptographic protocols for data transmission and storage.

1.4.2 The role of trust

Online users show privacy concerns about the usage, the disclosure and the protection of their personal health information [**?**]. Moreover, they are also sensitive to the possible further dissemination of their personal health information. The research is well aware of the concern of privacy about the health information of individuals [**?**, **?**, **?**, **?**, **?**]. Among all the different types of personal and sensible information, a prominent place is occupied by the medical and clinical data. Among online users, is understood that exist a concrete concern regarding the usage, the disclosure, and the control that they have on their personal health information [**?**]. These pool of users are sensible also to the fact that it is possible that undesirable social and economic consequences can happen following a misuse of such data [**?**]. As part of the Health Insurance Portability and Accountability Act (HIPAA), introduced in 1996, a huge step in the handling and protection of sensible health information was made and also it brought to the forefront some privacy concerns [**?**].

These studies indicate that the lack of trust in ICTs and digital health care affects very seriously any effort to migrate from the conventional healthcare procedures to an electronic system. Trust has been the fundamental pre-requisite for the progress of commerce and prosperity in human societies [**?**]. The “trust” can be explained as the acceptation to depend on a third party (another person, institution, company, or other) based only on the belief of its integrity and/or benevolence [**?**, **?**, **?**]. Since trust determines to which extent an individual wants to depend on others, it is also a primary factor in changes in behaviour [**?**] and helps understand others’ motives and social interactions [**?**, **?**]. The central role of trust as a major type of social capital in online activities is well established [**?**, **?**, **?**]. Research on traditional business interactions has demonstrated the significant role of trust [**?**, **?**, **?**, **?**], moreover

the research on online purchasing activities has shown its the critical role in online business transactions [?, ?, ?, ?, ?, ?, ?, ?, ?]. The central role of trust in electronic commerce has been demonstrated multiple times [?]. The challenge of trust in a digital environment is the strong separation between how people perceive a potential danger and how people understand it [?, ?].

According to the above, any successful digital healthcare system should target at increasing citizen's trust. It is clear that both *trust* and *security* play central and fundamental roles: "The more people trust others, the less concern they have for misuse of personal information" [?]. Trust, however, is difficult to establish in the digital health domain since it requires interactions between computers, between humans, and between humans and computers.

There is much ongoing the development of new trust management models for complex and dependable computer systems. As privacy is connected to security, a similar relationship is also observed between trust and security [?]. Researchers propose the application of automated trust mechanisms in distributed systems [?]. Various schemes have been proposed for the design of secure information systems have been proposed which are based on automated trust management protocols [?, ?, ?]. The composition and propagation of trust information between elements of information systems are also of pivotal concern and a number of research works are devoted to them [?, ?, ?, ?].

1.4.3 Data protection regulations

The need for protecting individuals' privacy has been recognized by the law enforcement agencies leading to the creation of laws for data protection. A new European Union-wide framework known as the General Data Protection Regulation[1] (GDPR) has been introduced that provides a more uniform interpretation and application of data protection standards across the EU. Essentially it constitutes a fundamental change in the management of data privacy designed to protect and empower all EU citizens data privacy and with severe implications in the way organizations across EU approach data privacy. While the purpose of GDPR is to protect personal data at large, namely "any information relating to an identified or identifiable natural person", in this dissertation the focus is on clinical trial data as they necessarily require the collection and analysis of sensitive personal data (e.g. health data).

The regulatory framework defines three main roles: **The Subject**, namely the resident or individual providing his/her data to the organization for the purpose of the clinical trial. **The Data Controller**, namely the *clinical research institute* that determines the purpose and meaning of the processing (i.e., the clinical trial) of personal data provided by the subjects. **The Data Processor** that processes the personal data on behalf of the Data Controller. Note that in many cases the *clinical research institute* has the double role of Data controller and Data Processor. The following is a short summary of the main requirements defined in the law enforcement directive:

- **Explicit Consent.** Clear and definite conditions for acquiring consent from

[1] `https://ec.europa.eu/commission/priorities/justice-and-fundamental-rights/data-protection/2018-reform-eu-data-protection-rules_en`

data subjects (citizens) to process data;

- **Data Protection Officer.** A person is appointed to handle the necessary internal recordkeeping requirements;
- **Sanctions.** Non-compliance can result in serious penalties;
- **Territorial Scope.** The directive applies to all organizations processing data from data subjects (citizens) residing in the EU, not only EU based organizations;
- **Right to Access.** The data subject shall have the right to obtain from the controller confirmation as to whether or not personal data concerning him or her are being processed, and, where that is the case, access to the personal data and some other information;
- **Right to rectification.** Incorrect data has to be rectified;
- **Right to be Forgotten.** Data subjects have the right to request data controllers to erase their data;
- **Data Portability.** Data subjects have the right to request their data in a portable format, which allows one to transfer its data to another data controller;
- **Data Protection by Design and by Default.** Develop default privacy protection mechanisms and implement monitoring processes;
- **Notification Requirements.** Data breaches must be reported without undue delay.

Clinical trial data is, however, a "special" data category whereby processing is necessary for scientific or research purposes. This special data category negates the subject's right to erasure or portability. This is due to the fact that clinical data cannot just be removed or transferred from a dataset, without affecting the audit trail or the statistical outcome. Subjects can, however, leave a trial to prevent additional data collection.

In this context, the right of data portability means that clinical trial subjects have the right to receive their personal data in a commonly used and machine-readable format, and transmit such data to another organization.

Clinical trial providers must identify the data that is being processed, where it is transferred to, who processes the data, what it used for, any risks and processes, and ensure all employees are trained. Furthermore, they have to provide all these information to potential participants in a trial and keep records to show what individuals have consented to, what they were told, and when and how they consented. Note that, a clinical trial provider is a processor from a customer perspective but also a controller of data in terms of personnel, sales, and sub-contractors. As a consequence, clinical trial companies have obligations to make sure that rules are in place and followed.

Particular interest is article 32 of the directive that states that *"the controller and the processor shall implement appropriate technical and organisational measures to*

ensure a level of security appropriate to the risk". To this purpose, a crucial component of data collection in clinical trials is the distinction between pseudonymization and anonymization. Any pseudonymized data that can still be tied to an individual patient with the help of other information will still be considered personally identifying information (PII). Only fully anonymized data will lose the PII label, so trials must make the distinction between these two data types in trial protocols.

During this dissertation we will focus also on the technical implications of the Article 32 in the context of clinical trials when IoT devices are employed at home to collect relevant health and personal data for the trial.

1.5 State of the art

From the study in [**?**], it is clear that IoT will soon revolutionize the healthcare system that is transforming itself exploiting the potentialities of IoT. While on one hand the Internet of Healthcare Things (IoHT) technology offers ridiculous benefits, on the other poses new challenges. The work in [**?**] validates the effectiveness in using IoT devices in the following of diabetes outpatients. In [**?**] it is recognize that current personal health data is prone to hacking because of security vulnerabilities. In [**?**, **?**] the authors propose social media based approaches to make the people aware of clinical trials. These solutions greatly reduce the cost of advertising (often done through other media, like newspapers, television and radio) but do not help in characterizing the population before the actual recruitment. Anyway, exploiting social network allow to reach a really wide audience with ease. To overcome the difficulties in recruiting adolescents (in particular girls), the authors in [**?**] propose a mixed approach of social media use (Facebook) with the traditional paper mailing. This approach is interesting because firstly tries to attract interested adolescents using a technology well-known to them (the social media), then, in a second phase, it uses the traditional paper mailing to communicate further. The inclusion of people over 90 years old is of particular interest in clinical research, in fact they compose the fastest growing segment of the population [**?**]. In the *90+ study* analyzed in the same work [**?**], one of the inclusion criteria for the study was that the individuals must be within one hour drive from the study location. This limit can be (partially) overcome using remote monitoring system or similar solutions offered by eHealth.

The authors in [**?**] applied the MARKIT (Marketing and Information Technology) model to the SMART study (a clinical trial of weight loss for college students). Some of the subjects in the SMART study were monitored using IoHT devices. In the work [**?**] all the data is collected within a single system, where study staff can monitor, for example, the completion of questionnaires and more. The authors in [**?**] propose a high level view of their architecture to efficiently use wearables IoT in healthcare. From their study also emerge the need for standards and regulations. The work in [**?**] propose the submission of online personality questionnaires in order to increase the efficiency of recruitment. This additional step helps to identify the potential participants who will meet the key criteria. The authors in [**?**] propose an architecture to collect and process health data produced by specialized IoT devices. Their approach uses a centralized structure with an added privacy preserving and security enforcement module at the edges. They do not dive, however, into the

recruiting process of digital clinical trials. The article in [?] highlights the importance of wearables devices in both recruiting participants for digital clinical trial and the successive follow up (in following strict treatments but also in the long run).

The introduction of blockchain technology to healthcare is not completely new in the field, for example in [?] the authors propose a platform to conduct trial and better support precision medicine. It is crucial to keep in mind that smart wearables integrate sensing, computation and wireless communication in small, low-power devices that in many cases may operate in uncontrolled environments. Such low-sized embedded devices have limited sensing, signal processing, and communication capabilities and are usually battery operated. Due to this resource-constrained environment of operation, applying standard security and privacy requirements is extremely challenging [?]. As an example consider that some smart devices have limited computing and storage capabilities, thus cryptographic algorithms and protocols that require intensive computation, communication, or storage are simply not applicable. It is too costly (in terms of computation) to authenticate using a public key and too costly (in terms of memory and computation) to store one-way chains of keys. Also consider that some smart devices may be battery operated, forcing security mechanisms to reduce their energy consumption. These constraints greatly increase the difficulty of securing IoT-enabled systems and make them more vulnerable to security threats [?, ?, ?].

1.6 Data management in trusted space

A common approach to address the requirements and technical challenges outlined in the previous sections is to introduce a centralised, cloud-based trusted authority that is responsible for the control and processing of the data of the clinical trials. The data collected by the IoT devices at home during the different phases of a clinical trial depicted in Figure ?? are shared with this trusted authority that controls and stores them on a cloud-based infrastructure. In this way, existing commercial platforms (e.g., such as AWS IoT[2]) can be used to accelerate the development process and integrate existing business processes and IT enterprise infrastructures to enrich the delivered services.

Three main spaces are identified: the private space, the trusted space and the public one. During all the phases of the clinical trial, the data in the *private space* generated by the IoT devices, is enriched by other relevant data possibly residing in the *public space*, such as gender, sex, age etc. and delivered to the *trusted space*. The *Clinical Research Institute* is operating within the trusted space in order to analyze the data as an integral part of the research. It is evident that the user has severely limited control over the personal data residing within the trusted space. It is therefore critical that the trusted space conforms to all regulations relevant to data protection. For this reason, the Data Controller and the Data Processor take care to anonymize or pseudonymize the data residing within the trusted space in order to be compliant with Article 32 of the GDPR. For a graphical representation see Figure ??.

[2]https://aws.amazon.com/iot/

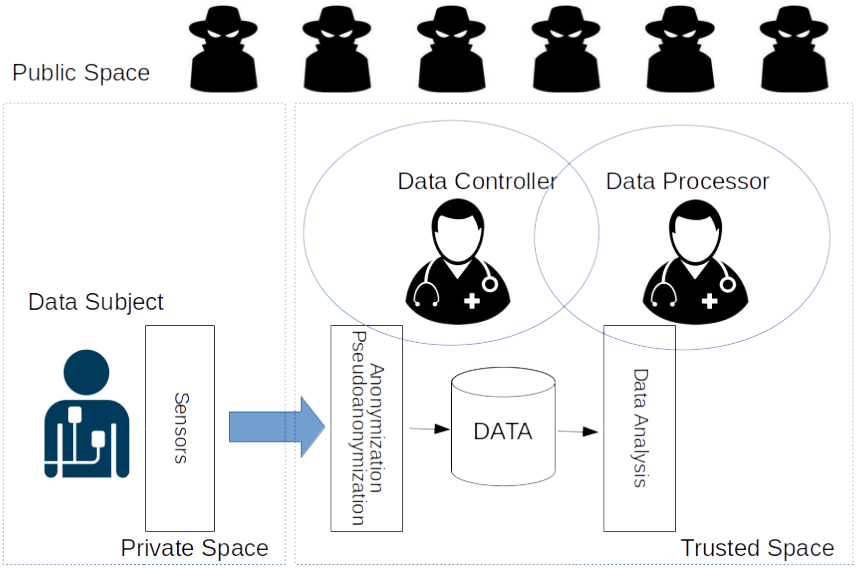

Figure 1.4. A simplified representation of the common approach in which according to Article 32 of GDPR anonymization and/or pseudoanonymization are in charge of the Data Processor and Data Controller.

In the following, three techniques that can be employed to guarantee the anonymization or pseudonymization of personal data collected will be analyzed. Remark that there is an evident trade-off between the privacy that can be guaranteed to the participant and the usefulness of the data for the purposes of the clinical trial; as much the data are handled to preserve the privacy of the participant, as more difficult it becomes to extract relevant information to be used in the clinical trial. Note that the data in the public space are among the *"other information"* that can be used to tie pseudo-anonymized data to an individual patient, thus voiding the anonymization process.

1.6.1 K-anonymity

K-anonymity is a practical approach for data anonymization. The key idea of k-anonimity is that the features that would allow the identification of users are handled such that always exist at least k records in the dataset with the same set of features, so in principle is difficult for an adversary to distinguish among those k records. To clarify this concept, consider the example in Table **??**. In this example, the adversary knows some information about the users, namely name, age and ZIP. If the adversary can have access to the medical records, despite they do not contain the name of the users, he can immediately correlate this information with the ones in his availability, to infer sensitive information about the disease of the users. Through k-anonimity (2-anonimity in this case) this risk can be reduced. Indeed with his information the adversary can claim that either Joe or Nic have disease A (or B), but not exactly who.

There is a clear trade-off between the re-identification probability that can be tolerated and the utility of data; while higher values of k imply a lower probability

Name	Age	ZIP
Joe	15	1
Nic	18	2
Lou	35	3
Mary	40	4

Age	ZIP	Disease
15	1	A
18	2	B
35	3	C
40	4	D

Age	ZIP	Disease
[15,18]	[1,2]	A
[15,18]	[1,2]	B
[35,40]	[3,4]	C
[35,40]	[3,4]	D

Table 1.1. On the right the adversary's knowledge, in the center the original medical records which, on the left a 2-anonymous table.

of re-identification, they also introduce more distortion to the data, in some cases reducing significantly the usefulness of the data. [**?**] provides a survey on k-anonymity in data mining while [**?**] explores the applicability of k-anonimity to health records. In the latter, the authors suggest that a hypothesis testing approach can be effectively used to control over re-identification risk and to reduce the extent of information loss compared to baseline k-anonymity.

1.6.2 l-diversity

In [**?**] the authors show two simple attacks to a k-anonymized dataset that can lead to severe privacy problems. The *homogeneity attack*, exploits the limited diversity in some sensitive attributes. In particular, if the value for a sensitive attribute within a group of k records is the same, that value can be predicted exactly even in a k-anonymized dataset. As an example of homogeneity attack, consider the case in table **??** where the Disease is C for both the records in the group with Age [35,40] and ZIP [3,4]. The *background knowledge attack* relies on some background knowledge that is possibly not encoded in the dataset, but allows the attacker to infer the most likely values for some attributes. To overcome the homogeneity attack, the most simple definition of l-diversity [**?**] requires that the records in a group shows at least l distinct values.

1.6.3 Differential Privacy

The main purpose of differential-privacy [**?**] is to make indistinguishable the output of an algorithm that analyzes a dataset and computes statistics, when a record in the dataset is either present or absent. In other words, looking at the output of the algorithm, one cannot tell whether any individual's data was included in the original dataset or not. This implies, that an adversary cannot learn anything (w.h.p.) about the presence or absence of that particular user, irrespectively from the peculiar characteristics of that user.

More formally, given a randomized algorithm A and two datasets $D1$ and $D2$ that differs in exactly one record (i.e., the data of one person), A is ϵ-differential private if for any $S \subseteq Range(A)$

$$Pr[A(D1) \in S] \leq e^{\epsilon} Pr[A(D2) \in S]$$

The architecture of a differential privacy system [**?**] is represented in Figure **??**. The analyst submit a query to the privacy guard, a software that assesses the privacy impact of the query "using a special algorithm". The query is delivered to the

database that responds. The guard adds some "noise" according to the evaluation of the privacy impact and the noisy rensponse is finally delivered to the analyst.

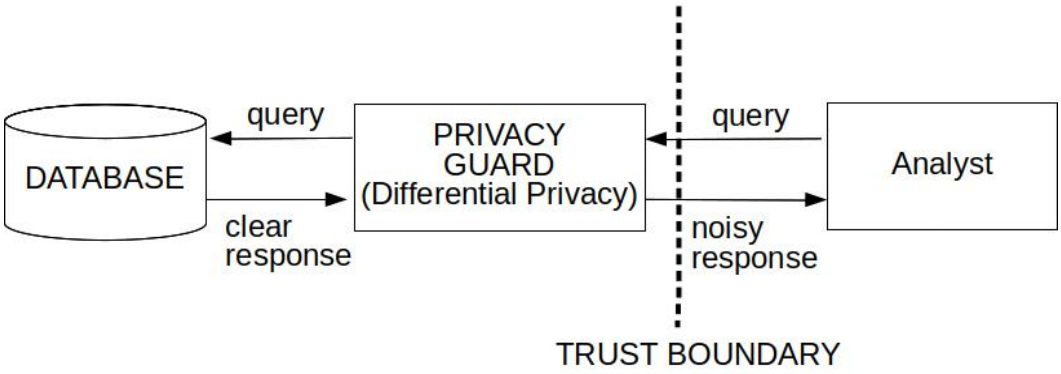

Figure 1.5. A simplified reference architecture

The correct evaluation of the privacy impact is crucial in this process and it is primarely related to the selection of the parameter ϵ, which is the parameter controlling the tradeoff between privacy and accuracy. While this is one among the most critical aspect for the applicability of differential privacy in practical cases, to the best of our knowledge, there is still no rigorous method to evaluate it in the literature. Dwork [**?**] indicates that the value of ϵ is a "social question", leaving de-facto open the problem, while in [**?**] the authors discuss the challenge of setting the proper value of ϵ given the goal of protecting individuals in the database with some fixed probability; they show that the clues about the fact that a specific individual is in the database or not can change depending on the query, the values in the data, and even on values not in the data. More recently, the authors of [**?**] proposed a model that expresses the balance between privacy and accuracy and they used such model to choose ϵ on a series of simple statistical studies. Despite such efforts, still a satisfactory evaluation of ϵ is a challenge and make difficult the applicability of differential privacy in practice. Indeed, in the literature the value of ϵ can range from 0.01 to more than 5. Finally, [**?**] analyses the main criticism about differential privacy.

The paper [**?**] discusses the applicability of differential privacy to the health care domain. While the motivations supporting the use of differential privacy and the corresponding challenges are very well explored, unfortunately the actual application of this techinque in real world helath care is very limited. Very recently, the paper [**?**] has promised a step forward towards practical differential privacy for SQL queries. The authors implemented FLEX, a tool to enforce differential privacy for real-world SQL queries on any existing database with negligible performance overhead. Remarkably, the approach has been recently adopted by Uber to enforce differential privacy for their internal data analytics [3].

[3] `https://medium.com/uber-security-privacy/differential-privacy-open-source-7892c82c42b6`

1.7 Challenges of IoHT-generated data

Today, multiple IoT devices for healthcare (namely *IoHT*) are available, from the professional infusion pumps and ventilators to the smartwatches, from portable blood analyzers to mobile EKG units and much more. The *clinical research institutes* expect that these devices will meet certain requirements, thus they provide useful data. From the work in [**?**] it is clear that IoT-generated data must retain some properties such as accuracy, freshness, availability and so on. The successful integration of IoT technologies in healthcare requires to address very specific technical challenges. For example, a *clinical research institute* that conducts a digital clinical trial needs to guarantee certain technical requirements throughout the complete process (see Figure **??**), from specifying the desired patient profiles and the digital screening process up to the collection of data and conducting the analysis.

Our architecture, that will be introduced starting from Chapter **??**, retains some of the following data properties in digital clinical trials. It also guarantees privacy protection to the individual and it assures quality of data to the *clinical research institute*, at least during the first phases (see Figure **??**) of a digital clinical trial where both the users and the *clinical research institutes* are most vulnerable. We want to protect the interests of business parties (like the *clinical research institutes*) while guaranteeing as much privacy as possible to the users. Out architecture provides a method that allows the CRIs to track the data generated, validating the *authenticity* of the devices and retrieveing their *accuracy*. Since the same architecture uses encrypted communication channels and storage solutions, it can assure the *confidentiality* of data to the users. The hashing of data is used, instead, to validate the *integrity* of gathered data. During the first phases of the trial, users' data remain into their respective private spaces, thus our architecture should not leverage online cloud storage solutions to increase data *availability*. Still, during the first phases, data is not provided directly to the CRIs thus we can not guarantee *freshness* of data. Anyway, we timestamp all the information in a secure and reliable way (leveraging the blockchain technology) so that the CRIs can check for it in the next phases when, instead, we use IoHT devices to send data possibly guaranteeing also *freshness*.

In the sequel, the most critical technical challenges related to the acquisition, storage and processing of data are presented, with a focus on eHealth applications. Despite in our architecture we relax the concepts of *freshness* and *availability* of data, we report them in the following for sake of completeness.

1.7.1 Accuracy

The term *accuracy* itself may refer to the closeness between an accepted reference value and the test result [**?**]. A measurement, instead, is described by its *trueness* and *precision*. To better explain the concepts, *trueness* represents the closeness between an accepted reference value and the average of a large number of tests [**?**]. The *trueness* itself is affected by the *systematic error* that is a value that shifts all measurements in a systematic way [**?**]. Contrary with respect to the *systematic error*, there is the *random error* that affects mainly precision and can vary in an unpredictable way. This kind of error occurs for a plethora of reasons, like noise in the measurements or a too small sample size. Finally, *precision* is not related to the

true absolute value but refers to the closeness between test results [**?**].

At an abstract level, a sensor is an electronic component that converts a physical value (such as O_2 in the air, light and temperature within a room, etc.) to an electronic form, digital or analogic. More advanced sensors, instead, apply particular algorithms on top of raw data to obtain data. Examples of these sensors are the optical heart rate and pulse oximetry sensors that rely on an optical sensor varying its value, or various blood analysis that use chromatography to retrieve components' concentrations.

Heart rate, the glucose level in the blood, number of steps made per hour, daily calories intake and much more are examples of information that can be gathered from IoT devices. Like any other information, they have a certain accuracy that characterizes the data [**?**]. For example, it was studied [**?**] that heart rate monitoring made by common activity trackers and smartwatches have accuracies that range from 99.9% to 92.8% and, thus, in certain scenarios they can be treated as accurate. In [**?**, **?**], the authors measured the performances of a very common activity band with respect to professional calorimeters. While during activities made on plane surfaces the accuracy was relatively high ($> 80\%$), a slight inclination change produced a notable worsening in performances, achieving a poor evaluation of burnt calories with accuracy degrading by more than 40%.

Another interesting aspect is that accuracies of different parameters can vary deeply depending on sensor positioning. As an example, sensors positioned on the wrist, on the chest or on the hip achieve different accuracy [**?**].

1.7.2 Authenticity

Data may be collected from high-quality, high-accuracy equipment targeting professional uses or from low-cost, low-accuracy devices suitable for recreational activities. Remote patient monitoring requires the use of equipment that provides a constant and measurable performance so that the experts can evaluate correctly the health status of the patient. It is therefore required to establish an authentication method in order to verify the source of data and avoid fake or data of poor accuracy. A typical example could be the assumption of opiates drugs, that cause dependence [**?**]. A patient is willing to fake data in order to receive stronger medications or additional doses of the same drug but for a longer period of time. Another example could be the faking of data coming from a smart scale in order to benefit from incentives [**?**]. These examples demonstrate the need for an authentication method for all the IoT devices that participate in a digital clinical trial.

A consistent number of works in the literature provide primitives for authentication in low power systems and, in general, in IoT devices. For example a two-way authentication system based on Datagram Transport Layer Security (DTLS) protocol could serve [**?**], or an authentication and access control framework based on CoAP [**?**]. Certificate-based authentication mechanism for WSNs can also be used in order to allow the sensor nodes and the end-users to authenticate each other and initiate secure connections [**?**].

It is equally important to take into account the reputation of the devices in order to detect promptly attacks, malicious behaviours or tamperings. Towards this end a trust management model based on fuzzy reputation for IoT can address this issue [**?**].

The problem is well studied in a broader context, for example the case of a Social IoT (SIoT) [?].

1.7.3 Confidentiality

Personal health information is perceived as one of the most important data to protect from third parties. As already stated in [?], 85% of the respondents perceive privacy concerns as a barrier to share health information and it is clear that the individuals request a high standard of confidentiality. While some years ago IoT devices were usually constrained in terms of processing power and memory [?], new devices take advantage of the technological advancements in the silicon industry that offer high processing power with little energy consumption. The new capabilities of embedded processors and microcontrollrs allow advanced algorithms to be executed within the IoT device [?, ?, ?, ?].

Data confidentiality is mostly achieved through encryption, using an algorithm such as AES, DES or RSA [?]. These algorithms are highly optimized and represent a mature technology, but often they require some processing power (it depends also on the parameters for encryption and the strongness willing to achieve) [?]. It is important to highlight that, in order to assure confidentiality, some encryption algorithms require to successfully realize a key exchange before opening a secure communication channel [?, ?].

One of the most used silicon architecture in small devices is the ARM Cortex-M. Within these Integrated Circuits (IC), it is common to find hardware accelerators for a security application, most notably the AES accelerator [?]. Running encryption algorithms in hardware allows very constrained devices, like activity trackers, smart wearables but also RFID tags [?] to communicate confidentially, guaranteeing high levels of security.

1.7.4 Freshness

Depending on the particular clinical trial scenario, data need to be processed and evaluated within certain delay requirements. In trials targetting wellness and safety, delay is a critical requirement. For example for heart diseases such as arrythmia, identifying and generating early warnings require very short response times [?, ?]. In other applications the freshness requirement can be relaxed regarding slowly changing parameters. For example, the glucose level in the blood can be delayed by minutes since it changes relatively slowly [?].

A critical aspect of existing IoT devices targeting healthcare is their *lack of computational power to locally process the ECG recordings and detect abnormal behaviour*. Therefore recorded signals need to be transfered to cloud services where advanced analysis algorithms are executed for processing and Integrated [?, ?]. In a typical IoT architectures, data flows from the IoT devices are transmitted to a nearby gateway device and then to cloud services for further processing, analysis and integration [?]. In such a typical IoT deployment, there are several technical issues that need to be addressed in order to guarantee any freshness requirements imposed by the specific clinical trial. Depending on the sensor used, in certain cases even a trace of a short period can require large amounts of data to be transmitted

over the wireless network and eventually to the core network. Considering that data flows can follow different paths, in particular beyond the local gateway and within the core network infrastructure, it is natural to encounter delays, during disassembly and reassembly, as well as jitter in the communication [**?**, **?**].

Another technical issue relates to dense deployments of IoT devices within small areas where a very large number of devices communicate continuously with the gateway. Despite that each sensor may transmit a potentially small quantity of data, the total number of data flows deeply stress current communication protocols that rely on medium contention (e.g., CSMA - Carrier Sense Multiple Access) such as Wi-Fi 802.11 and ZigBee 802.15.4 [**?**] to cite the most common ones. This issue could cause loss of data due to buffer overloading (particularly true on constrained devices) and to delays [**?**].

It is also important to consider that wireless comunication are also susceptible to cyber attacks like jamming [**?**, **?**]. In such cases, the frequencies used to enable the communications are saturated with noise, making the medium unusable. Various techniques were proposed to limit this issue, between the most notable one, there is the *spread spectrum* with its declinations [**?**, **?**]. It is evident that minimal possible latency, network bandwidth preservation, and efficient data storage resource utilization are elements of paramount importance to address the freshness requirement [**?**].

Recently, researchers have examined the possibility of allowing IoT devices to become capable of executing advanced alerting algorithms locally. In this alternative IoT architecture, segments of sensor data are classified wether they are important or not for the clinical trial and thus if they should be stored and transmitted to the cloud infrastructure. In this way the overall data traffic is minimized while additionally the IoT device would conserve battery power and minimize memory requirements. This concept of combining the resource-bound last-mile sensors of any IoT-related application with computational capabilities is receiving increasing attention from researchers and practictioners. The so-called *Fog computing* approach extends the cloud computing paradigm by migrating data processing closer to production site, accelerates system responsiveness to events along with its overall awareness, by eliminating the data round-trip to the cloud. Offloading large datasets to the core network is no longer a necessity, consequently leading to improved resource utilization and quality of experience (QoE) [**?**]. Another very important benefit of the Fog computing approach is the increased level of data control. Since the data collected from the wearable devices is not forwarded to the cloud, we are in the position to allow the user to maintain control of all collected data. In other words, Fog computing reinforces our goal of guaranteeing the privacy of confidential data, which is of paramount importance when designing smart healthcare systems.

Apart from the particular approach to address the freshness requirement, it is critical that data generated by the IoT device is timestamped. In this way, potentially dangerous situations can be avoided [**?**]. Sensor that assess the sleeping patterns of individuals that do not incorporate accurate timestamps can easily lead to naive miscalculations of the time spent in the bed. Addressing this requirement requires that IoT devices are equipped with accurate clocks that are periodically synchronized.

1.7.5 Availability

The role of IoT devices is to collect data that will be used in the clinical trial. As soon as the physical phenomenon is sampled by the sensor, it is stored locally. Moreover, under the *Fog Computing* paradigm the data is analyzed locally. Therefore *data availability* depends on both the connection, storage and processing technical approach followed.

Data storage follows two main paradigms: *centralized* and *decentralized*. In *centralized* solutions, a single entity, such as a server or generally a device, stores all the data. While this can be convenient in terms of cost, resources need and more, it also makes the system less robust against hardware failures and power outages. In fact, a centralized system has a *single point of failure*. On the contrary, decentralized solutions allow for better scalability, do not suffer from the single point of failure and is robust also against large-scale power outages.

In IoT applications it is common to find locally centralized systems [**?**, **?**, **?**, **?**, **?**] that send data to the cloud periodically, where the storage solutions are mostly decentralized [**?**, **?**, **?**, **?**]. This hybrid approach present some strengths like the simplicity of installation, maintenance and connection since a single device that acts as a gateway should be configured and connected to the internet but at the same time suffers from the single point of failure and power outages.

During recent years, we witness the increasing usage and development of low power, long range and low-cost radios that paved the way for local-infrastructure-independent devices [**?**, **?**]. In most cases, communication infrastructures such as LoRaWAN [**?**, **?**, **?**, **?**] and SigFox [**?**, **?**] allows the usage of IoT devices without the need for a gateway. The main limits of such architectures are the offered bandwidth [**?**] that is limited and their cost. In [**?**, **?**], the authors investigate the limits of the PHY and MAC layer of SigFox and LoRa protocols. The real-time capabilities of such protocols are limited so that applications that require tight real-time communication or require high bandwidth are very limited. At last, it is important to highlight that the solution proposed by SigFox is strongly limited on the downlink communication (from the cloud to the devices) allowing only a few packets per day to be delivered.

Availability is also influenced by the amount of energy available on the specific IoT device. Clearly, for devices connected to the main supply, this is not an issue, but the vast majority of IoT devices used in healthcare (such as smart bands, activity trackers, smart watches but also smart scales, blood pressure monitors, pulse oximetry meters and so on) are powered through batteries. Despite the very fast improving in silicon technologies, batteries are not following the same trend and improve their capacity of about 5% - 8% every year [**?**]. Moreover, lithium-based batteries are not considered strongly safe, as demonstrated by the very flourish research activity [**?**, **?**, **?**] that try to avoid dangerous explosions and thermal runaway typical of this kind of batteries.

Battery depletion causes unavailability, thus the research is trying to address the issue with multiple approaches, from designing low power algorithms to changing in network topologies to implementing newer and more energy efficient hardware. For example, in [**?**], the authors present a low-power system for acquiring and classify biosignal coming from body sensors, while in [**?**] a mechanism is presented to adapt the radio power in order to decrease the overall energy consumption. In [**?**] the

authors present a lightweight multicast forwarding for service discovery and in [?] the other authors describe how to optimize power consumption using BLE (Bluetooth Low Energy) technology.

1.7.6 Integrity

Wireless communication technologies used in IoT, such as ZigBee [?, ?], LoRa [?, ?] and SigFox [?, ?] do not offer packet fragmentation. This impacts all the IoT applications where data that need to be transfered cannot fit within a single packet size. For example, LoRa allows a maximum packet size of 256 bytes, ZigBee of 133 bytes, while SigFox allows only 12 bytes for upstream packets and 8 bytes for downstream packets. These sizes do not reflect the actual payload of the packets, they often include the header and sometimes also the preamble. It is evident that IoT devices must implement data fragmentation mechanism. Such mechanisms must guarantee the integrity of information during data reassembling.

Interestingly, data fragmentation is not the only threat to integrity. *Data integrity* is strongly connected to the concept of *protection of information* from possibly malicious third parties, cybercriminals or any external interference from the initial transmission to the final reception of data. Thus, the system must be aware of the threat whenever it tries to tamper the data [?]. Malicious third parties could be interested in making revenue for their false outsourced data. A solution for such problem is, for example, investigated in [?], where the authors provided an analysis on data integrity verification based on authenticator suitable for both the cloud and the IoT. Also in [?] the authors present their solutions in order to achieve privacy preservation during the communications between all the components of an IoT system. In [?] the authors' present public-key cryptosystems as desirable solutions whenever there is a need for data integrity and authenticity.

1.8 Structure of the thesis

We shown how IoHT solutions changed the healthcare system and the clinical trials, giving birth to the digital clinical trials. The need for privacy is intrinsic in personal health data and we are interested in guaranteeing as much privacy as possible, especially during all the phases that compose a digital clinical trial. We now introduce the arguments covered in this dissertation, in Figure **??** we grouped them for ease of explaination.

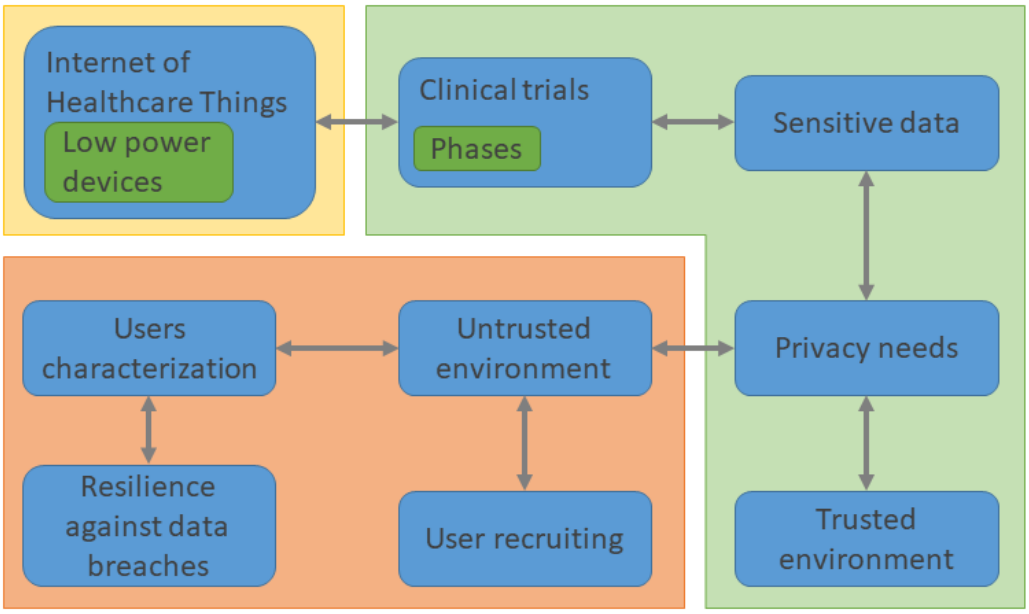

Figure 1.6. Dissertation arguments

In this introductory chapter we introduced the clinical trials and their phases and we gave some insight regarding digital clinical trials. Then we analyzed the need for privacy that is intrisic in any healthcare related application and the how the enforcement of GDPR regulation changed the way to handle sensitive data. The common centralized approach used nowadays in personal data-focused application were illustrated, among with some possible techniques to enhance privacy. Challenges posed by the usage of data generated by IoHT devices were introduced. All these concepts are highlighted within the green box in Figure **??**. The rest of the dissertation is organized as follows.

In Chapter **??** we present the core of our work, diving into an alternative approach of conducting digital clinical trials. We suppose to work in an *untrusted environment* where the only chance to effectively protect the personal data is containing it within the private space of the users as much as possible and only very limited (and anonymized and privatized) information is exchanged outside. The interested third parties, namely the clinical research institutions, are interested in accessing the personal data of the users. They want to accomplish two mainly tasks: *characterize* the population of users and *recruit* suitable patients for a specific digital clinical trial. The recruiting of patients relies on a blockchain technology (the *smart contracts*)

and on the *recruiting function*. The characterization phase can be used instead even before starting a digital clinical trial. We show two different approaches: one based on a distributed privacy preserving clustering algorithm and another that relies upon the properties of *random projections*. The latter approach guarantees user privacy on one hand while allowing the execution of interesting machine learning algorithms for clustering patients on the other. We demonstrate that this approach offer strong privacy also in case of data breaches. In Figure **??** we highlight the covered arguments within an orange box.

In Chapter **??** we show the embedded device that we designed and realized to validate our architecture. In the specific we dig into the hardware and software solutions that were necessary to obtain a device capable of long lasting operation with little to none maintenance. The data produced by this device is secure and it is possible to track its authenticity. In Figure **??**, this chapter is highlighted within the yellow box.

Finally, in Chapter **??** (not shown in Figure **??**) we briefly recap the work done and its limitations. Hints for future works are presented at the end.

Chapter 2

Data management in private space

In this chapter, we illustrate the advancements in enhancing the privacy of the users during the first phases of digital clinical trials and our concept of private space where the private data is located [**?**, **?**, **?**]. We took into account the privacy risks posed by the use of IoT devices aimed to healthcare, and we introduce some solutions able to contain them. At the end of the chapter, we present a practical attack that could endanger our solution and we show how this solution can contain the damage subsequent to the data breach [**?**, **?**].

It is evident that when data management is done with a trusted authority, the user has to completely trust the *Clininal Research Institute* for the protection of its sensitive data. Unfortunately, the uncontrolled growth of internet-centred services has led us to accept many compromises about how data are shared. Interestingly, for certain phases of a clinical trial (namely users characterization and patients recruiting) it is possible to avoid the transfer of most of the private data from the private space to the trusted space while allowing the *Clinical Research Institute* to conduct these steps without accessing clear confidential data. In fact, during the recruiting phase the users can choose to participate without disclosing any information, while during the characterizing phase, all the data shared are anonymized or shared in privacy-preserving fashion. The central idea is to take advantage of the increased computational capabilities of the IoT devices to reinforce the privacy of confidential data while still conforming to all requirements relevant to data protection, including the GDPR. Therefore the user may retain complete control of the private data and be free to decide if the data will be accessible by the researchers or not. Moreover, in [**?**], the authors propose a distributed solution that targets diabetes clinical trails. Similarly to our solution, they leverage the blockchain technology to safely store information and uses multi party computation in order to process data. Their approach however is not GDPR compliant since it does not give individuals the ability to revoke the access to their personal data.

In particular during the *initial design*, *recruting* and *enrollment phases* of a clinical trial (see Figure **??**) an alternative approach is feasible, in which the personal user data is not stored in any trusted space, instead, this data always remain in the

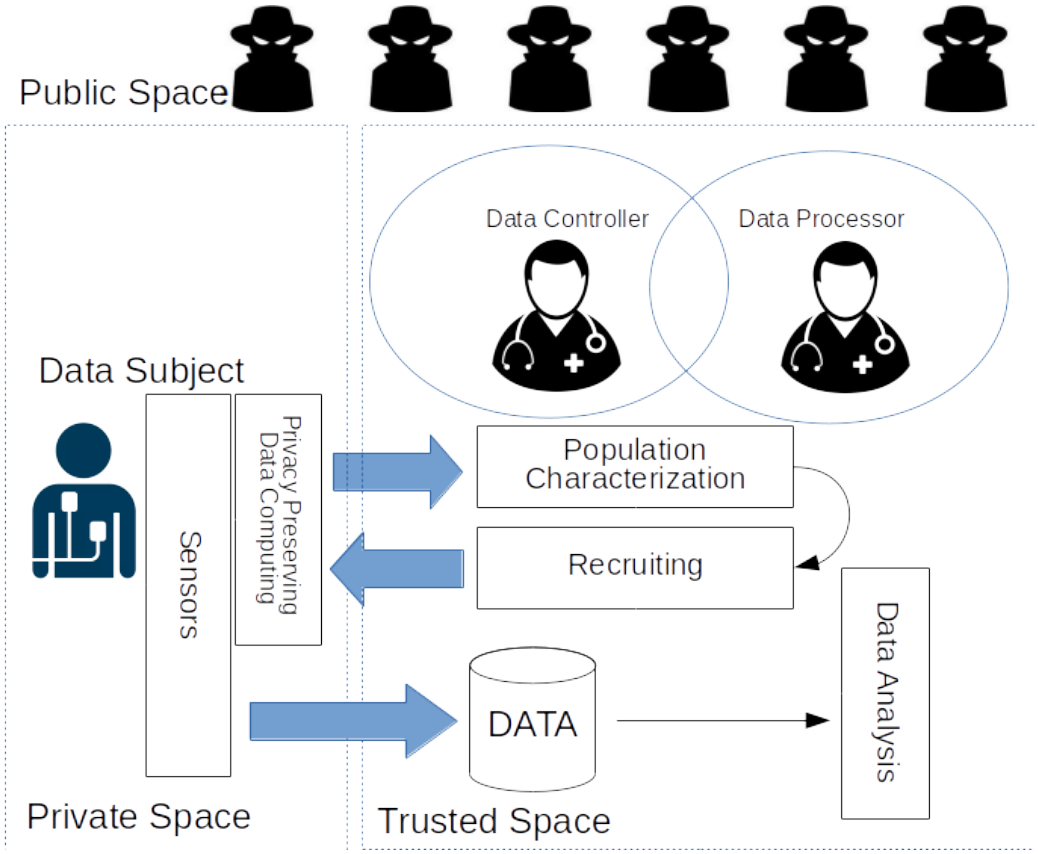

Figure 2.1. A simplified representation of the first phases af a digital clinical trial where the privacy of the users is retained. Once the user accept to enroll, data ar threated as in the common approach depicted in figure **??**

private spaces of the users. In the subsequent *enrollment phase*, when the users are enrolled on a clinical trial, the data needed by the *Clinical Research Institute* to conduct the actual research are transferred to the trusted space with all the necessary anonymization or pseudonymization steps described in the previous section. Such an alternative approach allows users to retain the maximum control over their data until the point when they are sure that the *Clinical Research Institute* can generate value from the use of their data. At this point the users will receive their fair part of the value created. Therefore, this approach of moving data management from the trusted space to the private space follows the arguments of the *My Data is Mine* declaration[1].

In more details, during the *recruiting phase*, the *clinical research institutes* has two main needs, namely a) to characterize the community of potential users so that it can design an effective clinical trial and b) to recruit suitable users. Both these tasks can be done without having access to the data collected by the users and thus without any potential violation of the privacy of the users. A simplified representation of the proposed approach is shown in Figure **??**. Initially, the user participates in the characterization process that allows the clinical research institution to better design the clinical trial. Then the *clinical research institute* starts the recruiting process contacting all the users. Only those users that match the specific criteria are invited to enrol in the trial. The matching is computed in the private space of the users thus no personal information is disclosed. As a consequence only users actually enrolled in the clinical trial will deliver their personal data to the trusted space, while all the others will not reveal any relevant information except the one necessary to characterize the population, process that however has been designed to preserve the privacy of the user.

The critical assumption in this alternative approach is that IoT devices and commonly available data processing equipment residing within the private space of the user will be capable of guaranteeing the security and confidentiality of the sensitive data. These processing elements interact with the Data Controller and Data Processor residing within the trusted space. It is evident that the resulting distributed system is composed of a variety of different subsystems that operate with a plethora of privacy and trust requirements. The design of such a system-of-systems is particularly complex due to the interactions of IoT systems with real-world processes [**?**]. Inevitably, delivering robust applications requires testing and performance evaluation of individual system components as well as compositions of the system on real hardware in large-scale deployments. Given the increased production of IoT hardware nodes and the introduction of new tools for managing IoT testbeds (e.g., see [**?**, **?**]) it is important to incorporate in the design process a real-world testbed deployment, such as the IoT-Lab facility[2].

In the following sections, we present the design and the realization of three phases of a digital clinical trial (namely the recruiting, the enrollment and the users characterization) that follow the distributed data management approach introduced before. The design is accompanied by a proof-of-concept (PoC) implementation that allows to evaluate the proposed solutions in real-world hardware and to get a first

[1] `http://www.mydataismine.com/manifest`

[2] IoT-LAB: a very large scale open testbed, `https://www.iot-lab.info/`

feedback on the feasibility of this alternative approach.

The PoC uses embedded devices with different processor architectures (such as Microchip ATMega128RFA1, ST Microelectronics STM32L476 and STM32F103) and different wireless technologies (both standard - like 802.15.4 - and non standard protocols communicating over the ISM band, at 868 MHz or 2.4 GHz) to emulate the wearable devices for the generation of biometric data.

A Raspberry Pi3 Model B+ produced by Raspberry Pi is used to represent a commonly available data processing unit residing in the *private space* that on one hand interacts with the IoHT devices of the user, and on the other hand it interacts with the parties outside the private space, such as the blockchain and the *Clinical Research Institute*. It has a powerful ARM Cortex-A53 microprocessor with 1GB of LPDDR2 RAM on-board (see Figure **??**). The gateway supports the same wireless technologies used by the wearables devices and, thus, can communicate with them. This device is also referred to as the *"gateway" device* since it enables the interaction between the private and trusted spaces.

Finally, a standard PC (Intel i7-6500U, RAM 8GB DDR3) is used to represent the *Clinical Research Institute* and implement the necessary functionality residing in the *trusted space*.

Figure 2.2. The Raspberry Pi 3 Model B+ produced by Raspberry Pi. Copyright CC BY-SA. Credits `www.raspberrypi.org`

The blockchain is a distributed database that permanently stores data. It is often referred as a *distributed ledger*. This database is accessible from everywhere there is an connection to the pool of participants, commonly through a simple internet connection. Every entity can write into the database but it is costly, instead reading data from it is free. This so-called *cost* depends on the blockchain used. It can be,

for example, processing power, virtual or real money and other. Every change made in the database is recorded, so that also a small modification leaves a permanent trace. Some implementations of the blockchain, like *Ethereum* that we used, allow to embed small portions of code inside the chain, allowing for example to release an amount of virtual coins after an event happened. This portions of code are called *smart contracts*. We rely heavily on smart contracts. For example, we use them to store the hash of the private data and the public keys of the smart devices used, but also to identify nodes capabilities and specifications (like accuracy, reliability, etc.).

2.1 Architecture

We now introduce the general architecture in Figure **??** with a specific emphasis on private and untrusted spaces. This architecture allows to conduct recruitment in digital clinical trials using a decentralized approach. We consider individual users that are using a collection of (a) wearable devices attached on them and (b) smart devices positioned within their home environment. These devices compose the IoHT of each user in Figure **??**. The *gateway* is in charge to communicate with the IoHT and to manage the storage in the private space and it also handles all the communications happening between the private and the untrusted space. The *blockchain*, the *internet* and also the *clinical research institutes* are considered part of the untrusted space since the single user has little to none control over them.

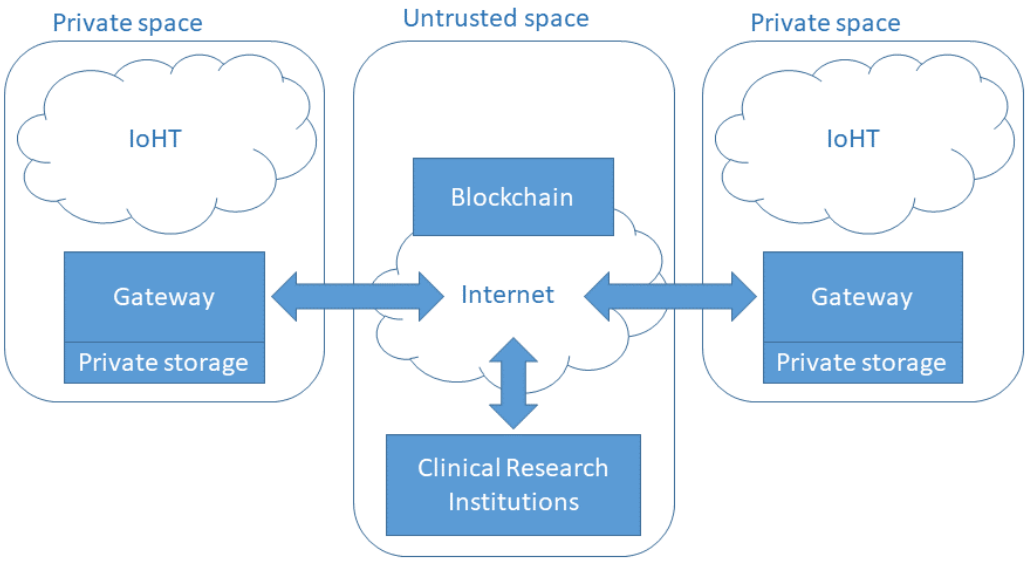

Figure 2.3. The macro blocks composing our general architecture, with a specific emphasis on private and untrusted spaces

Move one step ahead, we now show in detail (Figure **??**) all the steps required to recruit and enroll a patient in a digital clinical trial. Into the private space of the user (namely, the home environment) there are IoT devices dedicated to gather data. They collect information about the health status of the individual, his habits and more. Currently, all wearable devices (and smart devices) are equipped with an

identifier that encodes the manufacturer, model number and production series. The proposed approach requires that they are also equipped with a tamper proof memory space where the manufacturer places a private key. The corresponding public key is stored on the blockchain and thus publicly accessible by any entity. This allows a research institute to identify the devices available to the individual and to decide if the data collected are within the accuracy requirements.

Within the home environment there is also a so-called *gateway* to which periodically all available IoT devices communicates with. These communications are private and digitally signed using an asymmetric key scheme. The data received by the gateway (step 1 in Figure **??**) is stored locally into a local storage (the so-called *private database* in Figure **??**). Furthermore, whenever the wearable devices communicate the collected data to the *gateway*, all packets are digitally signed and therefore the *gateway* can check their integrity and authenticity, and discard those that are arriving from untrusted devices. Periodically, the *gateway* applies a hash function over the collected (trusted) data with its signature (stored into the private database) and the result of this hash function is safely stored into the *blockchain* (step 2 in Figure **??**) (thus, into the untrusted space).

Saving the hashing into a blockchain solution will allow later the *Clinical Research Institute* to check that the dataset provided are indeed the ones collected over the period of time claimed. Moreover, the clinical research institute can check the authenticity of the retrieved data since the hash take into account the signature of the devices generating data and the corresponding public keys of these devices are stored into the blockchain too.

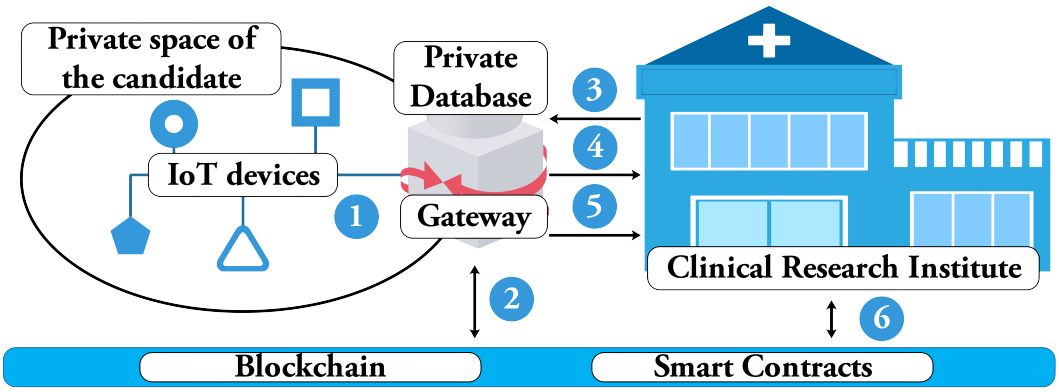

Figure 2.4. Recruit and enroll a patient in a digital clinical trial

Up to now, no data is brought externally from the private space except for the hash. A clinical research institute interested in recruit for a digital clinical trial should now design a *recruiting test*. The recruiting test is a function that analyses a historic dataset provided by the individual users and decides if the candidate should be included or excluded. The analysis is encoded into code that generates only constant-size output (i.e., O(1)) and thus cannot include any personal data that can violate the confidentiality of the candidates. The *Clinical Research Institute* sends the code to all interested parties (step 3 in Figure **??**). Each *gateway* locally execute it on the locally stored historic data. If the user agrees to participate in the

trial, and the result of this function states that this user is fit for the trial, only then the output of the recruiting test is transmitted to the *Clinical Research Institute* without revealing any of the private data (step 4 in Figure **??**). If the candidate is eventually enrolled for the actual trial, he provides his actual historic data (step 5 in Figure **??**) and the *Clinical Research Institute* can finally check the data integrity employing the hash function published in the chain (step 6 in Figure **??**).

As already observed, in the enrollment phase (that started in step 5 in Figure **??**), participants have to trust the data management of the research institute. In view of this, instead of the blockchain, a more traditional trusted and certified database could be used to store all relevant information. However, we decided to employ the blockchain in our solution, because its P2P (and trustless) nature is in line with the principles of participation, personal data ownership and accountability that inspired our design.

2.2 Guaranteeing originality and authenticity of collected data

Initially, the *Clinical Research Institute* conducts certain accuracy evaluation tests over various off-the-shelf wearable devices and smart devices to compile a list of "trusted" devices that it considers accurate enough so that data collected from these devices can be used through the digital screening phase. Based on an adequate set of genuine historic values collected from one or more of these "trusted" devices, the clinical research institute analyses the data (e.g., using a combination of logistic regression models, principal component analysis, etc.) and decides if the patient will be included or excluded. This is translated into the following requirements.

- The quality of data provided by participants is guaranteed (only consider data collected from approved devices).
- Data are collected by certified and trusted devices.
- Fake data cannot be introduced into the private space.

One way to address the above requirements is for the wearable devices to digitally sign all the packets containing genuine measurements before transmitting them to the *gateway*. The digital signing is done using an identifier that encodes the manufacturer, model number and production series along with a private key that is installed in a tamper-proof memory space. The corresponding public key is stored on a blockchain residing in the untrusted space to make it publicly accessible by both the *gateway*s and the *Clinical Research Institutes*. Therefore the *Clinical Research Institute* is capable of identifying the devices available to each individual user and deciding if the data collected are within the accuracy requirements. Moreover, the *gateway* can check their integrity and authenticity, and discard those that are arriving from untrusted devices.

A proof-of-concept implementation is done using the open-source `uECC` library[3] available on `RiotOS`[4]. This implementation supports the recommended elliptic curve

[3] `https://github.com/kmackay/micro-ecc`
[4] `https://riot-os.org/`

[**?**] over binary fields with equation $y^2 + xy = x^3 + x^2 + 1$ along with the irreducible polynomial $f(x) = x^{163} + x^7 + x^6 + x^3 + 1$. The curve's order (the number of points on it, r) and the base point is $G(x, y)$ are listed on Table **??**. The testing blockchain residing in the trusted space is implemented based on the `Ethereum` Blockchain App Platform[5].

Parameter	Value
r	$0x400000000000000000020108a2e0cc0d99f8a5ef$
x	$0x2fe13c0537bbc11acaa07d793de4e6d5e5c94eee8$
y	$0x289070fb05d38ff58321f2e800536d538ccdaa3d9$

Table 2.1. Parameters for the Basic Elliptic Curve Operations.

The evaluation indicates that the signature function is by far the most time-consuming, while the hashing is always relatively fast. Figure **??** shows the time necessary to sign 64000 bytes of data on the resource constrained *M3 Open nodes* dividing them in chunks of different size (1000×64 bytes, 100×640 bytes, 10×6400 bytes, 1×64000 bytes).

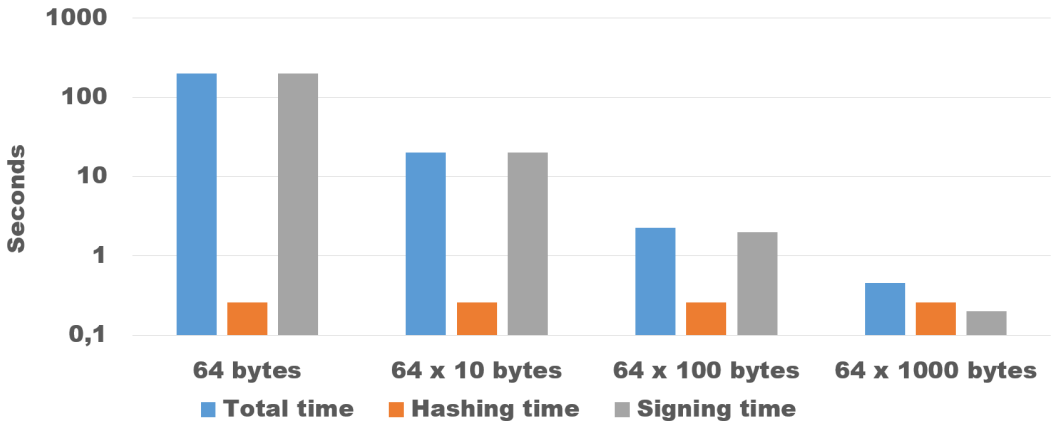

Figure 2.5. Time necessary to sign chunks of data of different sizes for an overall of 64000 bytes. Please note the logarithmic scale.

The total time necessary to sign 1000 chunks of 64 bytes were more than 200sec, namely more than 400 times bigger than the time necessary to sign the chunk of 64 Kbytes (about half a second). Furthermore, when we consider chunks of 64 bytes, the signature occupies about one-third of the payload, while in the 64000 bytes case it is less than 0.1%. This result suggests that some form of aggregation is always necessary to implement a practical solution.

The *M3 Open Node* (see Figure **??**) is the first device that we used in our PoC. It is provided by IoT-Lab and its microcontroller is an ARM Cortex M3 with 32 bit architecture running at 72 Mhz. It has 64kB of RAM.

Despite we have no performance measurement on our node (based on an ARM Cortex M4 with 32 bit architecture running at 80 Mhz. It has 128kB of RAM), its performances are comparable with the ones shown in Figure **??**.

[5] `https://www.ethereum.org/`

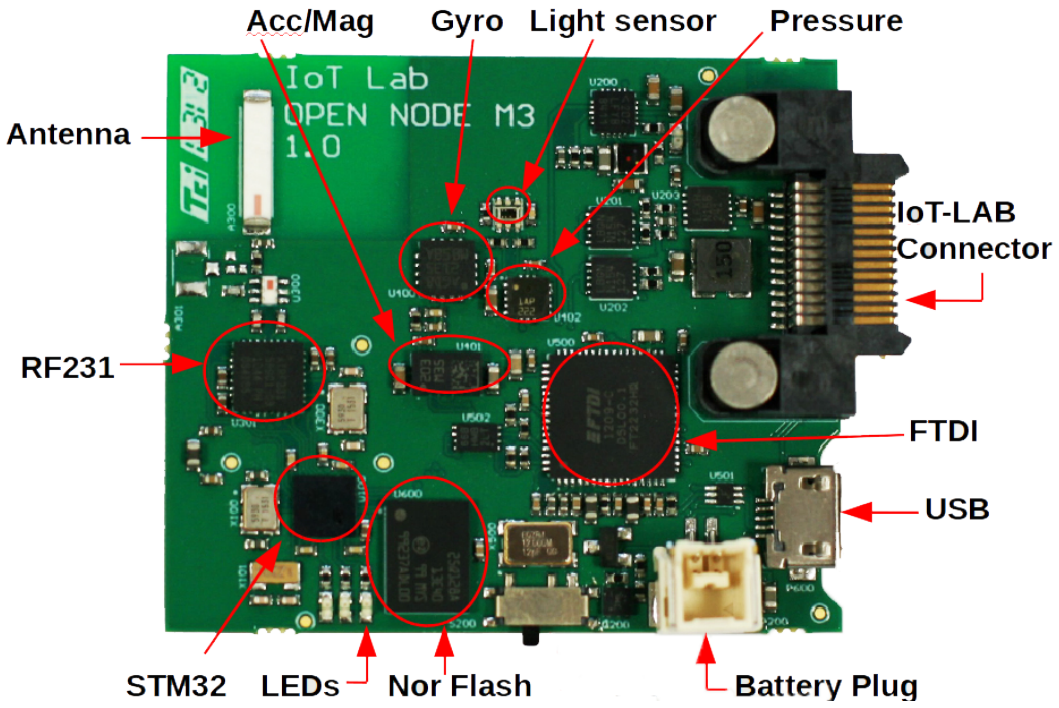

Figure 2.6. M3 Open Node from IoT-Lab. Copyright © FIT IoT-LAB

2.2.1 Interacting with the blockchain

Data acquired in the private space are periodically hashed and stored in the blockchain. This allows the *Clinical Research Institute* to validate the authenticity, integrity and conformance of collected data. Since the interaction with the choosed Ethereum blockchain technology is expensive in terms of computing power, it was convenient to delegate it to the gateway.
The blockchain is also used to store the public key associated with the healthcare devices and published by the manufacturers of the IoHT devices on the blockchain. In this way the clinical research institutions can check which device generated each piece of data. Moreover, since it is possible to understand the class of devices generating data, it is possible to discriminate only the devices that are accurate enough for the specific application. Providing guarantees on the performance of validating transactions in the blockchain is not possible by definition, but usually is in the order of minutes (using Ethereum blockchain technology). On the contrary, the time necessary to read the blockchain is negligible ($O(1)$) because the entire blockchain is stored locally.
The very structure of the blockchain provides a mean of timestamping. So that, clinical research institution (or any interested third parties) can retrieve the hash of the stored data and considering the associated timestamp can check if the data was actually retrieved in the considered timestamp. Thus, malicious gateways cannot inject data after the recruiting phase stating it was generated before, since the blockchain does not allow it.

2.3 Privacy preserving digital screening process

The *Clinical Research Institute* starts the recruitment phase based on the patient profiles specified in the previous steps. The screening relies on an adequate set of genuine historic values collected from one or more of these "trusted" devices, the clinical research institute analyses the data (e.g., using a combination of logistic regression models, principal component analysis, etc.) and decides if the patient will be included or excluded. We translate these steps into the following requirements.

- The inclusion/exclusion of a candidate is based on the result returned by a well-defined *recruiting test* automatically executed over a provided data set of historic values.
- The candidate must provide proof that the historical data set is real and collected over the stated period of time.
- An individual that wishes to be considered for a specific clinical trial expects that the privacy of his/her personal data will be respected and the confidentiality of the private data will be guaranteed. If during the digital screening phase the individual is excluded, then the digital health system should guarantee that no personal data will be retained by the *Clinical Research Institute*.

It is assumed that a recruiting test is used by the *Clinical Research Institute* to discriminate whether a candidate is suitable to enrol for a given clinical trial. Different recruiting tests are assumed based on machine learning algorithms within the following two main scenarios.

S1 The data in the private space, possibly pre-processed to extract relevant features, are used to learn some model. The parameters of such a model are given as inputs to the recruiting test.

S2 A given pre-computed model is embedded by the *Clinical Research Institute* into the recruiting test, and some features of the data in the private space are given in input to such model in order to classify the candidate.

In **S1**, both features extraction and machine learning are performed by the *gateway*, while in **S2**, the model for the machine learning phase is computed by the *Clinical Research Institute* and provided to the *gateway* to execute the classification based on the features extracted.

2.3.1 The recruiting test

The *recruiting test* is a function provided by the *clinical research institute*. Its role is to analyze the historic dataset of the individual and to decide if the candidate should be included or excluded. It is, again, important to highlight that at this point the personal data of an individual is still in his/her full control. The evalutation of this function is processed completely within the *private space* of an individual, so that none of his/her personal data is disclosed (see Figure **??**).

The analysis is encoded into code that generates only constant-size output (i.e., O(1)) and thus cannot include any personal data that can violate the confidentiality

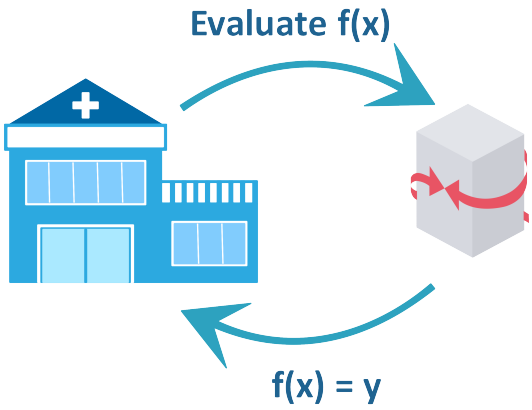

Figure 2.7. Recruiting function

of the candidates. It is mandatory to highlight that trivial recruiting functions must not be allowed. In fact, a sequence of trivial recruiting functions can execute a binary search. This is a known limitation at the moment, so that we consider that there is a mechanism that rejects such functions. The *Clinical Research Institute* sends the code to all interested parties (step 3 in Figure **??**). Again, the *gateway*s locally execute it on the locally stored historic data. The user can choose (see Figure **??**) to send the result of this recruiting function to the *Clinical Research Institute* without revealing any of the private data (step 4 in Figure **??**). The user is also free to refuse to reply to the recruiting function, exiting the recruiting process immediately. If the participant is eventually enrolled for the actual trial, he/she provides her actual historic data (step 5 in Figure **??**) and the *Clinical Research Institute* can finally check the data integrity employing the hash function published in the chain (step 6 in Figure **??**).

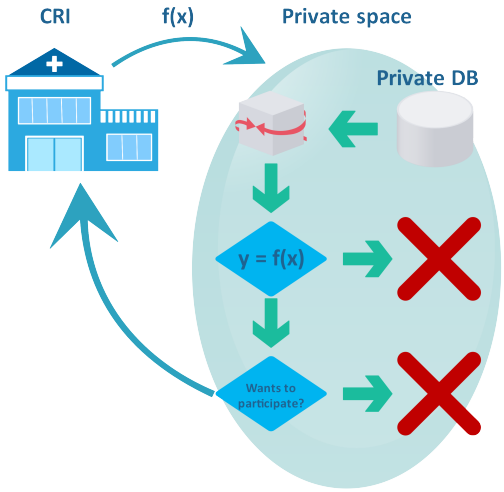

Figure 2.8. Recruiting function with user interaction

At a preliminary investigation with physicians, in many practical cases, the recruiting test is relatively simple and mostly based on threshold functions. However,

in our experiments, we performed some tests to evaluate the ability of our PoC to allow more complex and general recruiting tests based on machine learning algorithms.

2.3.2 Performances

In order to evaluate the performance of the two scenarios introduced in Section **??**, we choose four real-world datasets that are provided by the UCI Machine Learning Repository [**?**]. They vary both in the number of samples and in the number of features, accordingly to table **??**. Two of them target health issues such as cancer and heart diseases (*Arcene* and *Heart Diseases - Cleveland*). *EEG Eye State*, instead, correlates the EEG with open or closed eyes while *Gisette* has been selected to stress the performance analysis.

Table 2.2. Datasets used to evalute performances

#	**Dataset Name**	**Number of Samples**	**Number of Features**
1	Arcene	100	10000
2	EEG Eye State	13444	14
3	Heart Disease	270	13
4	Gisette	6000	5000

Machine Learning. The main goal of this experiment is to evaluate the ability of our architecture (in the specific, the ability of our gateway) to support the computation of complex recruiting functions on realistic data-sets. Six well-known machine-learning algorithms are used, available in the *Python* package *scikit-learn* [**?**]. Their execution times are shown in table **??**.

Remark that the objective of the evaluation conducted here is to measure the performance in terms of execution time. Since both processing units are executing the same algorithms and with exactly the same implementations, the accuracy is the same. The fine-tuning of such algorithms, in order to improve the effectiveness of the results, is not within the scope of this work.

In all cases the *gateway* is an order of magnitude slower than the PC but it can execute the algorithms in a reasonable time. The algorithms that are more demanding in terms of memory requirements can easily exhaust the memory generating significant delays (e.g., due to memory swaps) and in some cases create *Memory Errors* on the resource-constrained *gateway*.

Features Extraction. Raw data generated by the IoT devices in the private space are elaborated in order to extract relevant features provided as input to the machine learning algorithms. In this experiment, the ECG data are analyzed following the methodology of [**?**]. First, a Butterworth Band Pass Filter of the 5th order is applied. Then the Fast Fourier Transform (FFT) is calculated on the signal and the associated Power Spectral Density (PSD) are computed. Finally, the mean, the standard deviation, the variance and the maximum peak of the signal are obtained for statistical purposes.

Table 2.3. Execution times of common machine learning algorithms. Each row corresponds to the equivalent dataset in table **??**

	#	PC	Gateway
Support Vector Machines	1	0.147322s	0.695988s
	2	45.561s	565.288s
	3	0.005111s	0.061623s
	4	250.205s	1180.824s
Logistic Regression	1	0.133114s	1.639809s
	2	0.177369s	2.124157s
	3	0.001498s	0.020984s
	4	1.707945s	22.940s
k Nearest Neighbors	1	0.010120s	0.086751s
	2	0.009555s	0.125092s
	3	0.000355s	0.002438s
	4	1.315992s	18.851s
Gaussian Mixture Models	1	0.245765s	2.103226s
	2	0.667314s	7.934804s
	3	0.011209s	0.087351s
	4	30.658s	Memory Error
k-Means	1	0.116964s	0.859489s
	2	0.029887s	0.293893s
	3	0.003146s	0.035679s
	4	8.616173s	Memory Error
PCA	1	0.463008s	3.250940s
	2	0.055881s	0.662027s
	3	0.000539s	0.003660s
	4	40.488s	Memory Error

Table 2.4. Execution times to elaborate ECG

	PC	Gateway
Filtering	0.114648s	1.483026s
FFT and PSD	0.621372s	4.410873s
Statistics	0.016048s	0.115040s
Total	0.752068s	6.008939s

Table **??** reports the execution times required to analyze one hour of real ECG data sampled at 360Hz [**?**]. Again, the findings indicate that the *gateway* is roughly an order of magnitude slower than the *PC*, but it can extract important and complex features in absolutely reasonable times, with just a bit more 6 seconds of time required.

2.4 Distributed privacy-preserving characterization

Relying upon regulation to obtain guarantees on privacy and safety requires trustness, both with respect to the enforcement of the regulamentation by the government and on its effective application by the companies. In some digital clinical trial applications, however, it is important to infer something about an individual instead of accessing its whole personal data. Technically speaking, it is possible, under some circumstances, to infer something about an individual not from exactly its personal data but from a "privatized" version of it. The mathematical method of *random projection* could be exploited in order to mask the clear data of an individual. Moreover, this approach also provides some guarantees of resilience against a data breach.

In this section we show a distributed approach to characterize a group of users in a privacy-preserving fashion, while the next Section **??** is dedicated to our work on random projections applied to personal health data. We concentrate on the results given by common machine learning algorithms both on the clear health data and on the random projected health data. It will be clear that, thanks to random projection methods, it is possible to share "privatized" health data without disclosing personal information while making the same data useful for different kind of algorithms. In the end of this chapter, we will focus instead upon the possible de-privatization of projected data following a data breach. Under real-world usage scenarios, the chances to deprivatize data will be very low.

The *Clinical Research Institute* is interested in grouping candidates according to their characteristics in order to a) understand the community of potential patients, b) better design the clinical trial. For this purpose, a *privacy preserving data clustering* is conducted by carefully specifying the desired features relevant for the purpose of the trial. The features include the evaluation of specific biometric attributes (e.g., body composition, heart operation, daily activity, etc.) collected by the patient using wearable technologies over a given period of time (e.g., blood pressure for past week, etc.). During this process, the user's privacy is preserved since none of the private data is disclosed. The only information received back from the institution are the ones necessary to identify the clusters without disclosing data on the single users.

For the sake of simplicity, horizontally distributed data are considered in which the personal data of each party are disjoint and the parties want to jointly cluster their records without revealing their personal data. In the experiments reported here, the library presented in [**?**] is used to implement the proposed algorithm of Samet and Miri [**?**]. The performance of the privacy-preserving clustering algorithm is evaluated in comparison to that achieved by a standard PC available in the *trusted*

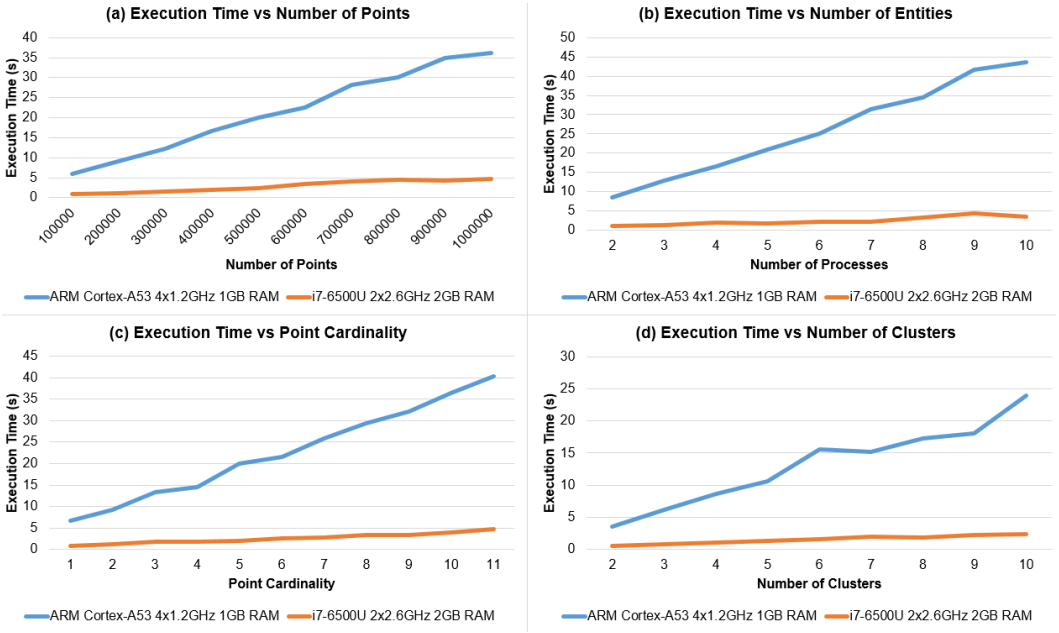

Figure 2.9. Execution time (in seconds) of the privacy preserving clustering algorithm on PC (orange) and *gateway* (blue)

space and the low-capabilities PC commonly available within the *private space*. The results of the performance evaluation are shown in Figure **??**.

Four parameters are examined independently: the total number of points used, the number of candidates (entities) participating in the privacy-preserving computation, the point cardinality, and the number of clusters (clusteroids) used as input (k) to the algorithm. The effect of each of these parameters is evaluated independently, by keeping fixed the others. However, it is clear that all these parameters affect the overall performance of the algorithm.

For the purposes of these analyses, synthetic datasets were used with specified bounds and distributions that are within biometric data reflecting heart rate, body mass, steps, etc.. Each parameter-set was run and evaluated over 10 iterations– in realistic situations, the number of iterations may vary with convergence time depending on the data and random initialization of the centroids. Depending on the parameters examined, the clustering algorithm may require a different number of rounds to reach convergence. In order to provide comparable results, the execution time is being reported as the mean time required to conduct 1 round of the algorithm.

2.5 Random projections and privacy

With the availability of accessible and widely used cloud services, it is natural that large components of healthcare systems migrate to them; for example, patient databases can be stored and processed in the cloud. Such cloud services provide enhanced flexibility and additional gains, such as availability, ease of data share, and so on. This trend poses serious threats regarding the privacy of the patients and the trust that an individual must put into the healthcare system itself. Thus,

there is a strong need of privacy preservation, achieved through a variety of different approaches. We study the application of a *random projection*-based approach to patient data as a means to achieve two goals: (1) provably mask the identity of users under some adversarial-attack settings, (2) preserve enough information to allow for aggregate data analysis and application of machine-learning techniques. As far as we know, such approaches have not been applied and tested on medical data. We analyze the trade-off between the loss of accuracy on the outcome of machine-learning algorithms and the resilience against an adversary. We show that random projections proved to be strong against known input/output attacks while offering high quality data, as long as the projected space is smaller than the original space, and as long as the amount of leaked data available to the adversary is limited.

Given the sensitive nature of healthcare data, there is a strong need to protect the information of the patients. Furthermore, the recent adoption of General Data Protection Regulation (GDPR) strengthens data protection and now it must be applied to any organisation or individual that collects and processes information related to EU citizens, regardless where the data is physically stored or where they are based[**?**, **?**]. At the same time, analysis of such data are crucial for medical research and the drug industry. Consequently, there is a need to design approaches that allow data processing without exposing the personal underlying information.

For this reason, there has been a series of techniques for perturbing data such that information on individual data points cannot be leaked, while aggregate information is preserved. Examples of such approaches are k-anonymity [**?**] and differential privacy [**?**]. The various approaches put different importance on the privacy requirements; for instance, differential privacy attempts to alter the data such as to provide very strong privacy guarantees, typically, without specifying the usefulness of the resulting data. for general-purpose data analysis.

In this chapter we apply a method, which can be found in [**?**], where the explicit goal is to obtain a dataset that remains useful after the perturbation (still providing some privacy guarantees). More specifically, our approach is based on *random projections* (RP), a technique that is typically applied primarily for efficiency reasons. It is based on a fundamental result from the work of Johnson and Lindenstrauss [**?**] and the idea is the following: Assume that we have large amounts of data (n datapoints), lying on a high-dimensional space. Then, if we project each point to a random subspace of dimension $O(\log n/\epsilon^2)$, with high probability all pairwise distances between the data points are preserved within a factor of $1 \pm \epsilon$. This technique has found multiple applications in streaming algorihtms, in finding nearest neighbors in high-dimensional spaces, in reducing the dimension of databases, and so on.

Our idea of applying a random projection approach on privacy-preserving data mining is the following. Assume that we have the records of multiple patients. Then we can consider a random projection of these data. The result of Johnson and Lindenstrauss guarantees that if we execute algorithms that depend on the pairwise distances on the data (e.g., several clustering or classification algorithms), then the results obtained are with high probability similar to those obtained on the real data set (and the error can be quantified). Furthermore, because the projections are random, one cannot use the projected data to obtain the real data: each datapoint

appears to be random. This last statement holds also in the case where a third party is interested in estimated data and even if he knows the projection matrix. In fact, estimating real data starting from the random projected one remains hard even if the attacker has some significant computing power. This trade off is studied in previous works (e.g. [**?**, **?**]).

It is not clear a priori that this approach could work in the application on medical dataset that interests us. For instance, the lemma by Johnson and Lindenstrauss is typically applied on settings where the original data lie on a very high-dimensional space. However, in practice, the original dimension may be low (for instance in our dataset it is about 50). In this chapter, we look at this and other issues by applying the random projection to a dataset containing information about 70K cases of diabetes [**?**]. We show that it is possible to reduce the dimensionality of the data and still obtain accuracy scores that are comparable with the ones obtained from the original non-projected data. At the same time, we also show how sensible and private patient information such their age or gender are safe against attacks that try to reconstruct the mapping between the original data and the projected data after applying random projections.

2.5.1 State of the art

Data leakages are very common [**?**]. In this work, we are more interested in reducing the ability of an attacker to reconstruct non-yet-leaked data from the leaked one. Within medical premises, there are multiple individuals who could obtain access to protected information from the rightous doctor to untrustful workers. This could lead to multiple entities knowing protected personal health information.

Before 2003, with the enforcing of HIPAA rules, some private medical information were regularly shared among professional [**?**]. Following the guidelines from Health Insurance Portability and Accountability Act (HIPAA) [**?**], the US government made the first concrete attempt to mitigate the chance of re-identification of patients. In 2009, it was clear that HIPAA is not sufficient to protect the privacy of an individual. In fact, the HIPAA was not able to protect the user personal information after the anonimization process that substituted HIPAA parameters with IDs. In a famous case [**?**], some researchers were able to re-identify users and also their sexual orientation and other information. Moreover, the availability of correlated data (coming from the same source or other sources) could greatly help to identify a patient. Data breaches continue to increase year after year, between 2005 and 2014, only in the US, more than 26 million of people had some form of personal health information leaked [**?**].

Therefore, more elaborate tecniques, which add noise to the data, have been developed in the last years to protect users' privacy and still maintain a good level of accuracy when exploring and analyzing the data. One of these is differential privacy [**?**]. This approach focuses on providing statistically coherent responses querying a database, i.e. third parties are interested to query for information about a sample of a population, not a single individual. Instead, we are interested also in providing data about a specific individual, for example investigating if he or she is suitable for a clinical trial.

In [?] the authors proved that the Johnson-Lindenstrauss transform can be used as an alternative approach to achieve differential privacy. The method is then compared against other techniques, such as adding Gaussian noise to data or randomized response. The proposed approach has superior accuracy bounds than the others, while still keeping secure the privacy of the records. The authors also criticize the work of Liu et al. about releasing data to third parties after applying random projections in order to protect sensitive information while still preserving accuracy of different data mining algorithms: an adversary that has some background knowledge can infer approximations of the original data. We address this issue in the scenario of known input-output data (section **??**) and show how in real world scenario regarding medical data, under reasonable assumptions for the power of the adversary, it is difficult for an attacker to discover private information from projected records.

In the literature there exist a very large number of works regarding the re-identification of person starting from various data, within some degree it is called "breaking the k-anonimity". For example in [?] the authors presents a method to re-identify a user from its preferences.

In this chapter, we aim at investigating to what extent RPs can provide useful data for machine learning algorithms (e.g. classification) on a group of potential patients while preserving at the same time the privacy of individuals. RPs have been employed in a number of healthcare applications, for example to segment tumor areas [?], to enhance tomography [?], to cluster DNA microarrays[?] or to classify cancer [?].

In [?], RPs were used to mask clear data projecting them in smaller spaces, while in [?] and [?], similarly to our work, the authors discuss how to exploit RPs to enhance data privacy. The authors in [?] also discuss the utility of the RP in reducing complexity of problems while maintaining the usefulness of the projected data for algorithms. It is anticipated that by 2020 there will be more than 26 billion devices involved in IoT related applications [?]. Surely, not all of them will be part of the healthcare field, however we expect a very large amount of information to process. The usefulness of RP in reducing problem complexity (or resource requirements) is well understood and exploited as useful resource in the literature [?, ?, ?, ?, ?]. For example, in [?], the authors explore some ways to reduce high dimensional data for clustering while, in [?], is presented a work on classification of small patches of images from a very large database that takes advantage of the properties offered by RPs.

During the last two decades, the contribution of machine learning and data mining algorithms in healthcare applications became more frequent year after year. This is well demonstrated in the literature, for example in [?, ?, ?, ?]. One last aspect to consider is the chance to link together multiple datasets. For example in [?], the authors presented the infrastructure of a databank in order to enable record-linkage research studies. This linkage on one hand could deeply help the development of newer treatments or drugs, but on the other hand poses threats to the privacy of the individuals.

2.5.2 Problem formulation

We consider a reference scenario in which a group of users, characterized by private features, are potentially suitable for a clinical trial. Only a limited number of users in the group will be actually enrolled in the trial. For the enrolled users, namely the patients, the private features will be eventually made public to participate to the clinical trial in the most effective way. Some knowledge on the group is of primary importance for the researchers to understand the size and the characteristics of potential patients. In general, users are well disposed to support this need of the researchers provided that their privacy is preserved. The main problem we want to address in this chapter is:

Can we learn something on the group of users as a whole, while preserving the privacy of the individuals who will not participate in the trial?

More formally, we consider a group of n users, where each user u is characterized by m features. We represent the corresponding dataset as a matrix $X \in \mathbb{R}^{m \times n}$, with m rows (the features) and n columns (the users). As already observed, in the era of big data, m and n can be particularly big.

Giannella et al. [**?**] show how it is possible to break the privacy in some contexts of distance preserving mappings. Liu [**?**] instead, highlights how mappings that do preserve distances within certain bounds like random projections can boost the privacy guarantees. We will apply these techniques in order to prove that users' privacy can be kept safe against malicious attackers.

We are interested in understanding to what extent the random projection technique, which has been originally conceived to reduce the dimensionality of a dataset, can also be used to preserve the privacy of the users. In particular, we apply a random projection to X, such that if $R \in \mathbb{R}^{k \times m}$ is the random-projection matrix $Y = RX$ is the transformed matrix after applying the random projection, with $Y \in \mathbb{R}^{k \times n}$. We denote by x_i^u the column in X associated to user u_i, and with y_i^u the corresponding column in Y. In the scenario we are describing the projected matrix Y is known to the public, it is indeed the dataset on which researchers try to distill information on the group; the transformation matrix R and the original data X are private. Some columns of X may become public once the corresponding users will eventually decide to participate to a clinical trial, in other words some pairs (x_i^u, y_i^u) will become public.

We can now better describe the problem, splitting it into two sub-problems:

Accuracy. *Can we learn something on the group exploiting Y?* Here we want to understand if the results of some machine-learning algorithms on Y are a good approximation of the ones obtained on X. If we answer positively to this question, we can at least conclude that what can be learned from the original data can be also learned from the projected data.

Privacy. *Can we preserve the privacy of the individuals that will not participate in the trial?* As already observed, Y is public whereas only some columns of X will eventually become public when the corresponding users will decide to participate in a clinical trial. Consequently, some pairs (x_i^u, y_i^u) will become

public. Here we want to understand if an attacker knowing Y and the some pairs (x_i^u, y_i^u) can possibly know something about the other users that do not participate in the trial.

We now elaborate on these two dimensions.

2.5.3 Accuracy

Lemma **??** provides a technique to generate a low-dimensional representation of the original data maintaining the pairwise distance within an error ϵ. Since the pairwise distance is the key ingredient for many classification tasks performed by machine learning algorithms, this property allow us to have some guarantees that the solution found in the low-dimensional space is a good approximation of the solution in the original and higher dimensional space. Furthermore, reducing the size of the input data speeds-up the execution time of the algorithms and limits the amount of resources needed. It can be proved that a random projection, is a mapping f that fulfills the previous lemma with positive probability. This is often referred as *JL-embeddings*.

Johnson and Lindenstrauss lemma

(Johnson and Lindenstrauss) Given $\epsilon > 0$ and an integer n let k be a positive integer such that $k \geq k_0 = O(\frac{log(n)}{\epsilon^2})$. For every set P of n points in R^m there exists a mapping $f : R^m \to R^k$ such that for all $u, v \in P$

$$(1 - \epsilon) \parallel u - v \parallel^2 \leq \parallel f(u) - f(v) \parallel^2 \leq (1 + \epsilon) \parallel u - v \parallel^2$$

2.5.4 Privacy: Known Input–Output Attack

We now try to answer one of the questions we raised in the previous section: Can a malicious third party who knows some pairs (x_i^u, y_i^u) (i.e. that a particular record x_i^u is associated to y_i^u after its projection) learn information about other records?

Liu in his Ph.D. thesis [**?**] describes a *Bayes privacy model* to measure the privacy offered by a perturbation technique. The model considers the attacker's apriori and a posteriori beliefs about the data and uses Bayesian inference to evaluate the privacy. For completeness, we repeat his framework here.

Let x be the unknown private data, y the perturbed data and θ the attacker's additional knowledge about the data. Then the MAP estimate of x given y and θ is

$$\hat{x}_{MAP}(y, \theta) = \arg\max_x f_{x|y,\theta}(x|y, \theta)$$

with $f_{x|y,\theta}$ the conditional probability density of x given y and θ.

Let X_p denote the first p columns of X and X_{n-p} the remaining columns. We define similarly Y_p and Y_{n-p}. We further assume that the columns of X_p are linearly independent and that X_p is known to the attacker (i.e., the attacker has full knowledge of p patients). Y is entirely known to the attacker, because as we stated before, it is publicly available to conduct experiments on the projected data.

For the next reasoning the following hypothesis must be verified:

- The original data arose from as a sample from a matrix variate distribution.
- The projection matrix R is a $k \times m$ random matrix with each entry indipendent and identically distributed with 0 mean and unit variance. R has a matrix variate Gaussian distribution with mean matrix $M = 0$ and covariance matrix $\Sigma = I_k \otimes I_n$.[6]
- Y has a matrix variate Gaussian distribution with mean matrix $M = 0$ and covariance matrix $\Sigma = I_k \otimes \frac{1}{k} X^T X$

The attacker will try to produce $\hat{x_i}$, with $1 \leq i \leq m - p$, such that $\hat{x_i}$ is a good estimate of the undisclosed private record x_i. In other words the attacker's target is to try to give an estimation of one of the records contained in X_{n-p}, given that he knows the records in X_p and their randomly projected counterpart in Y_p.

We now derive the MAP estimate of x given $y = Rx$ and the known matrices X_p and Y_p

$$\hat{x}_{MAP}(y, \theta) = \arg\max_x f_{x|y,\theta}(\mathbf{x} = x|\frac{1}{\sqrt{k}}Rx = y, \frac{1}{\sqrt{k}}RX_p = Y_p)$$

which can be simplified in

$$\arg\max_x f_{x,y,\theta}(\frac{1}{\sqrt{k}}R\overline{X} = \overline{Y})$$

where $\overline{X} = [xX_p]$ and $\overline{Y} = [yY_p]$.

We further suppose that the attacker has no other background knowledge about the private data, so we can assume that $\theta = 0$.

The previous result can be written as

$$\arg\max_x f_{x,y}(\frac{1}{\sqrt{k}}R\overline{X} = \overline{Y}) =$$
$$\arg\max_x f_{\frac{1}{\sqrt{k}}RZ|Z}(\frac{1}{\sqrt{k}}RZ = \overline{Y}|Z = \overline{X}) f_Z(Z = \overline{X})$$

If we assume that f_Z is distributed uniformly over an interval, we finally get

$$\hat{x}_{MAP}(y) = \arg\max_x f_{\frac{1}{\sqrt{k}}RZ|Z}(\frac{1}{\sqrt{k}}RZ = \overline{Y}|Z = \overline{X})$$

In [**?**, Theorem 5.3.8] is shown that the probability density function we obtained has the following form

$$(2\pi)^{-\frac{1}{2}k(p+1)} det(\frac{1}{k}\overline{X}^T\overline{X})^{-\frac{1}{2}k} etr\{-\frac{1}{2}\overline{Y}(\frac{1}{k}\overline{X}^T\overline{X})^{-1}\overline{Y}^T\}$$

We want to maximize this function in order to solve the problem of finding the best estimate of x given the observation of X_p.

Liu proposes an algorithm to estimate the nondisclosed records of a certain dataset. Experimental results have shown that while decreasing the number of column records known to the attacker (denoted by p) the relative error of the estimation

[6] $\otimes$ indicates the Kroenecker product of two matrices [**?**]

increases. The error in the estimation increases also decreasing the dimensionality of the projected subspace (denoted by k). In particular the algorithm uses the Nelder–Mead simplex algorithm to find the optimal solution of the maximization problem.

2.5.5 Experimental results

In this section, we present experimental results obtained on a dataset containing information about 70000 cases of diabetes diagnosticated in 130 US hospitals during the decade 1999-2008 [**?**] [7]. From now on we will refer to this dataset as the *diabetes dataset*.

We focus on the *classification* of patients based on their privatized data. Following the work in [**?**, **?**], we choose to use *random forest classifier* in our dataset to classify the users. Moreover, from the work in [**?**], we know that random forest classifiers works really well with random projections. In Figure **??** we report the effectiveness in terms of accuracy running the random forest classifier [**?**] on the original data and on the projected data in multiple lower dimensional spaces. To run and validate the classification algorithm, we divided the whole dataset into two parts: *train* and *test*. In the dataset we decided to predict the range of glucose level in the blood. So that, the algorithm was firstly trained with the records within the *train* part of the *diabetes dataset*, providing all the target values. Thus, we made the random forest classifier algorithm predicts the target values in the *test* part giving its features as input. Moreover, we tested the effectiveness of RPs also with *k-nearest neighbors (k-NN)* classifier, the results were reported in Figure **??**. Our approach was inspired by [**?**]. The results are quite different because in the first experiment we taken a feature of the dataset (the range of glucose level in the blood) as the value to predict, instead with the second experiment we choose to run firstly a *kMeans* clustering algorithm (on the whole dataset) to obtain labeled groups and then, with the *k-nearest neighbors (k-NN)* classifier we predicted the values.

The blue line represents the accuracy of the machine learning algorithm on the original data. The orange line, instead, represents the accuracy of the same algorithm on the projected (obfuscated) data. We tested the classification algorithm on projected spaces in different sizes, starting from only 2 components up to 10 components.

The lines plotted in Figure **??** presents the average values for each projection space, while the vertical wiskers represent the confidence interval corresponding to a specific projection space. For the baseline (classification on the clear data) we ran the classification algorithm 50 times, in each round starting from a random state of the random forest classifier. Since the original data is not projected into any space, we have only a baseline with the associated mean value and confidence interval. Thus, we reported the confidence interval only at the lefties part of line using wisker again. Instead, for the accuracy of classification on the projected data, we ran the algorithm more than 100 times. In each round the algorithm generated a value for each projected space. The results were obtained using the *scikit-learn* package on *Python 3.6*.

[7] The dataset is called "Diabetes 130-US hospitals for years 1999–2008 Data Set" and is available at this page

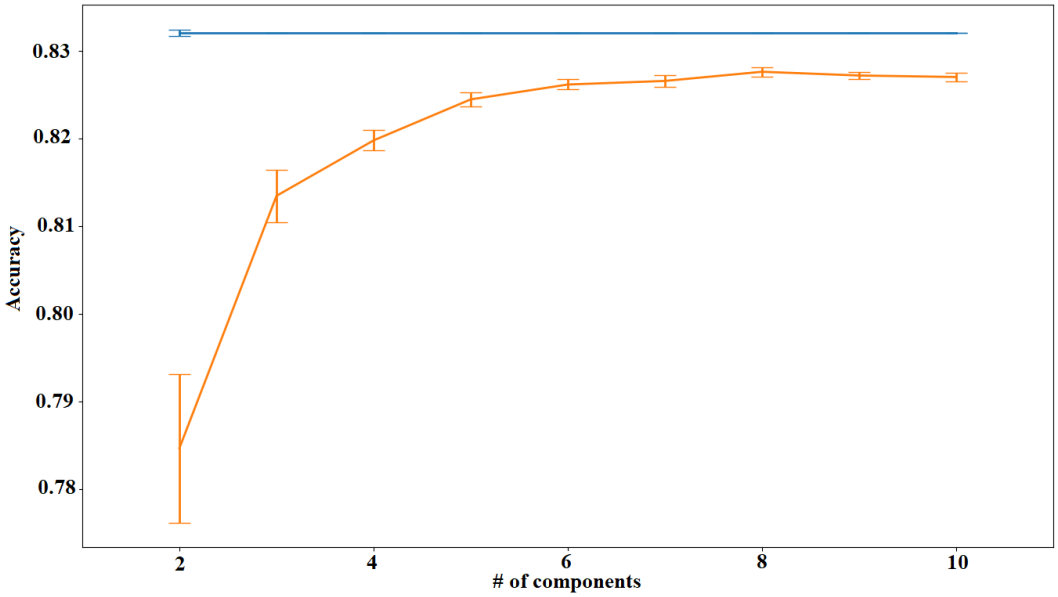

Figure 2.10. Accuracy of the *random forest classifier* algorithm on the original data (blue line) and on the projected data (orange line), varying the projection space (# of components). Mean values are reported as lines and 95% confidence intervals are reported as vertical lines.

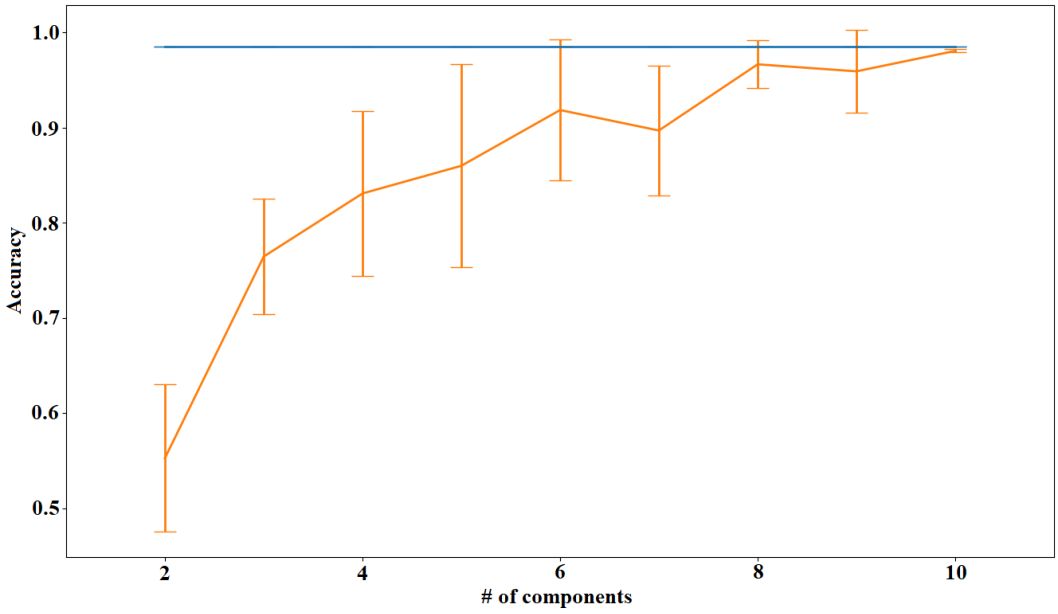

Figure 2.11. Accuracy of the *k-nearest neighbors (k-NN)* algorithm on the original data (blue line) and on the projected data (orange line), varying the projection space (# of components). Mean values are reported as lines and 95% confidence intervals are reported as vertical lines.

In [**?**, **?**] the authors explore the security of such techniques: they show how it is possible to use data dimensionality reduction techniques to lower the complexity of data mining algorithms while preserving their accuracy and how those techniques preserve the privacy of users.

The authors start from the same privacy hypothesis we have presented in **??** and study how an attacker in possession of a collection of linearly independent private data records and their corresponding transformed part can gather some insight about other records.

We present the results we got running the algorithm of [**?**] on this dataset. After choosing a number p of record pairs (x_p, y_p) we select a record x for which we do not know the mapping; the algorithm we are using will try to give an estimation $\hat{x}$ of the original record x.

We used two techniques to evaluate how similar to the original records the algorithm's estimations were. We measured the distance between the estimation $\hat{x}$ provided by the algorithm and the original record x. We compute the relative error between the two vectors with the following:

$$E(x, \hat{x}) = \frac{||x - \hat{x}||_2}{||x||_2}$$

The error E increases with the Euclidean distance between the two. Notice that with this notation it may happen that the error is greater than one: this could verify in the case that the distance between x and its estimations $\hat{x}$ is high and the norm of x is a small value. This could happen if the algorithm's estimation is very far off from the original record.

This measuring has the drawback to lack an upper bound for the dissimilarity. Neither the cosine similarity helps, since in our case we are not interested only in the direction of vectors but also in their magnitude.

A solution is provided in [**?**], where a radial basis function kernel can be used for representing similarities: we are going to use $1 - \frac{1}{e^{dist(x,\hat{x})}}$ as a similarity function between x and its estimation $\hat{x}$, where $dist(x, \hat{x}) = ||x - \hat{x}||^2$. The bigger the Euclidean distance between two vectors, the bigger the error $e^{dist(x,\hat{x})}$ will be. In this way we have a $[0, 1)$ bound for the similarity of the estimations. By applying the inverse we get a value in the range [0,1): if x and $\hat{x}$ are the same vector (perfect reconstruction performed by the algorithm) then $\frac{1}{e^{dist(x,\hat{x})}} = 1$.

Our workplan is the following: for every subspace of dimensionality k we apply the algorithm with different knowledge about the number of pairs (x_p, y_p) the attacker knows. We go from $p = k - 1$ to $p = 1$. In the next figures we display the results of our experiments, with the two different measuring techniques we used to quantify the similarity between the original records and the estimated ones. We report the mean of the errors for every pair (k, p) and the variance. On the X axis are placed the tuples (k, p) for which we have conducted the experiments, on the Y axis we placed the reconstruction errors.

On low-dimensionality subspaces we get a high relative error, meaning that it is not possible to give an effective approximation of the original (private) data records. In higher dimensions the approximation is closer to the original data. We

ran our experiments with 10 features of the dataset, since with vectors of higher dimensionality it becomes more difficult to run the reconstruction algorithm in reasonable times; also with higher dimensionalities the algorithm we are using outputs vector reconstructions that are very dissimilar from the original ones.

We applied the random projection to reduce the feature space in different dimensions, from 10 to 3. Notice, however, that even when the projected space has the same dimension of the original space, we already get a significant relative error, meaning that on the average it is not possible for the attacker to extrapolate any useful information about the patients' records. So for records of higher dimensionality there is already a safe privacy bound when applying random projection to them, at least against this kind of attacks.

We assigned an increasing numerical value to nominal features, that is, we assigned 0 to the text *male* and 1 to text *female* in the *gender* feature.

We applied random projection to this records, from $k = 10$ (no dimensionality reduction) to $k = 3$; the number p of pairs (original record, projected record) known to the attacker is in the range $k - 1 \leq p \leq 1$.

With $k = 2$ we obviously have only $p = 1$: we omit this result since it is not meaningful with respect the other results we get for higher k and p, because it does not show how knowing less (or more) information about the original data changes the reliability of the reconsturction we get.

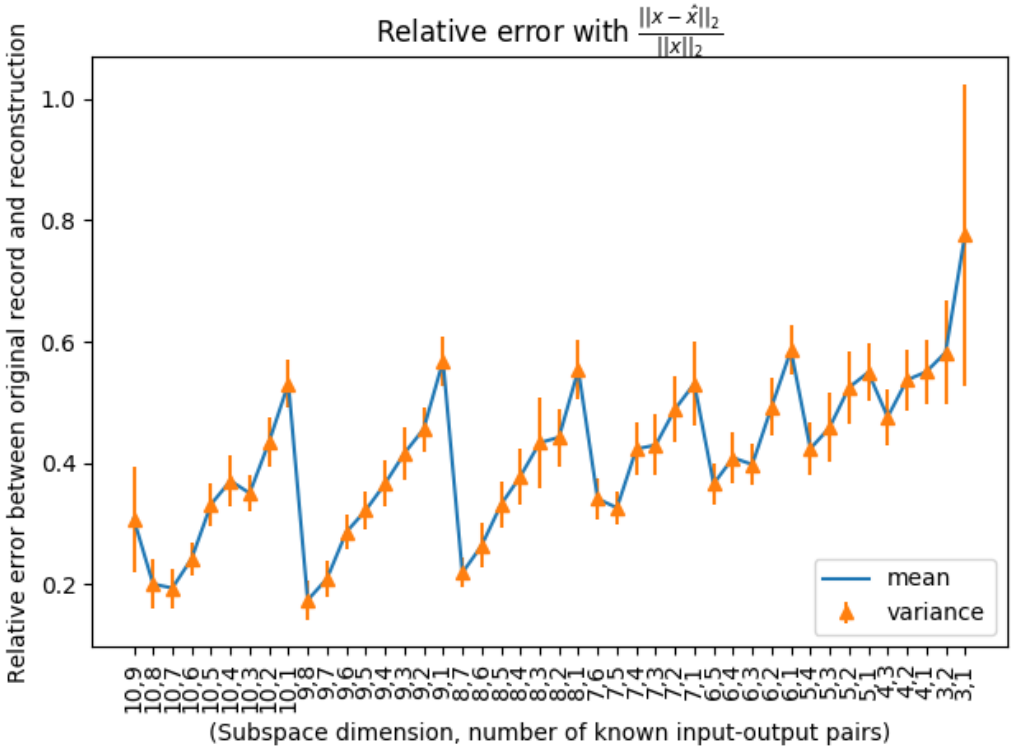

Figure 2.12. Mean and variance of the relative error while using the formula $\frac{||x-\hat{x}||_2}{||x||_2}$

In the next figures we show the mean and variances of the errors for every tuple

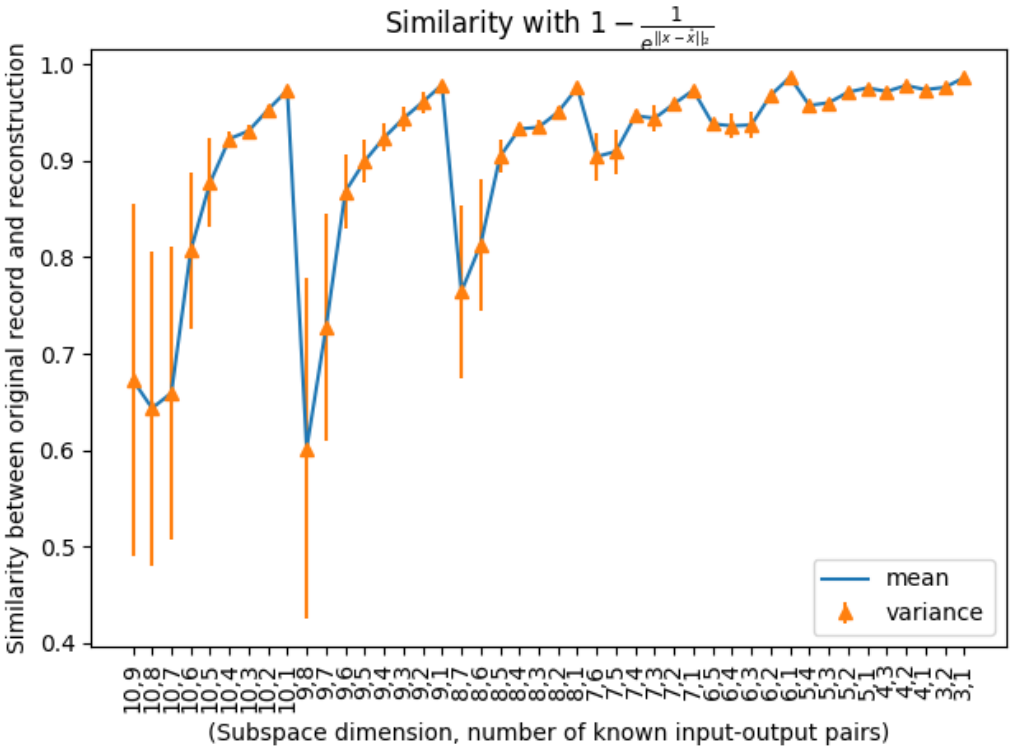

Figure 2.13. Mean and variance of the similarity between original records and their reconstruction while using the similarity function $1 - \frac{1}{e^{||x-\hat{x}||_2}}$

(k, p) for which we have conducted the experiment. It can be seen from the charts that as the number of known input-output pairs p decreases, the reconstruction error increases. Together with the dimensionality reduction, disclosing a scarce number of known input-output pairs can help with the task of preserving the privacy of users involved in clinical trials.

In this case we are projecting low dimensionality vectors ($k = 10$) but we still get high reconstruction errors when applying the techniques we have explained. This is another confirmation of the thesis that random projections help keep the privacy of users when their information is shared among research institutes.

Chapter 3

Sensing and processing in low power and wireless devices

In this chapter, we present the work done on refining an embedded solution that senses physical values, processes them and sends them to a sink (a gateway), achieving years of lifespan on batteries. Our first solution, the so-called *MMS* [**?**, **?**] was refined multiple times and posed the basis to the realization of the *RealTMS* [**?**]. The RealTMS was used to validate the architecture presented in Chapter **??** and proved to be a capable device for IoHT applications.

One of the propellant that is fueling the vast deployment of the Internet of Things (IoT) is the availability of inexpensive and inter-operative hardware. The technology made impressive progresses during the last two decades and now it reached the required level of miniaturization and low prices to take over the wider audience. Among all, most of the devices composing the IoT can operate with limited processing power, transmission bandwidth and energy. Thus, they can be deployed in different environment and sometimes they will never require maintenance during their entire life cycle (they will not need, for example, the changing of batteries, reprogramming and so on).

We will particularly focus on this matter: the optimization of current state of the art solutions to achieve long lasting operational state without maintenance. This will allow the deployment of devices in the most challenging environment, where the human interaction is strongly expensive or impractical. In the specific, we investigated throughly the capabilities of current state of the art solutions to efficiently use low energy power sources in order to achieve multiple functions, such as sensing an environmental value, its local processing and the transmission of data through multi-hop network up to a gateway.

The *"sensing"* action implies the interaction of an electronic device with the environment, in order to transform a physical value to a digital one. Since real world parameters are mostly analog values and modern systems are digitally based electronic devices, they need to convert the environmental parameters to digital values. Most of these conversions start from of a chemical or a physical reaction, but at the end they are mostly evaluated as the changing of a capacitance or of

a resistance. The Analog to Digital Converters (ADC) are specific circuits used to achieve these conversions. The very action of converting an analog value to a digital one requires a certain amount of energy, moreover there as some sensors that requires the heating a small plate to sense a particular value (for example some relative humidity sensors) or some other requires the forcing of air flow to sense the amount of dust in the air, just to cite a few of them. All these operations require energy, draining the batteries. One step further, an IoT device needs to *"process"* the gathered values (for example to retrieve the dew point from temperature and relative humidity) and it also uses processing power to handle the communication protocol. This actions, again, consume energy. In the end, the device will send the important data through a *"wireless"* link up to other devices. The proper action of partecipating in communications requires energy, again, and communicating through a wireless medium usually requires more energy.

We need to take care of all these actions in order to mantain the energy consumption as low as possible, using possibly sophisticated solutions on both hardware and software.

In this chapter, we present our work on designing, realizing and testing real devices capable to operate for years without the need of battery replacement, reprogramming or any kind of human interaction. Despite some of the works illustrated in the next sections aim to the Structural Health Monitoring (SHM) and, more generally, the engineering structure monitoring, they posed the basis for our proof-of-concept device used in the previous chapters. In fact, we took the design solutions presented in this chapter and we used them in our dedicated device for IoHT. We had to add some sensing components useful in clinical trials, like temperature and relative humidity sensor and shock/vibration sensor. The hardware used in the previous chapter (the *RealTMS*) is also more powerful than the one introduced here, and can accomplish asymmetric key encryption.

IoHT devices often need to use low power techniques. For example, the widely used *"holter"* devices run for more than a whole day measuring heartrate, pressure or other biometrical values. Wearables devices capable of sensing heart rate, sleep cycles and so on uses low power techniques in order to operate for days without recharging or changing batteries. While the very action of "changing batteries" seems in contrast to the effort described above ("totally removing maintenance"), we need to take into account the functions that a specific devices must do. For example it is possible to make an activity band that senses heart rate to last months, but instead of measuring the heart rate every ten seconds, it will sense heart rate every two hours. In the following, the concept called *"duty cycle"* will be explained in details.

3.1 A Modular Design for Wireless Structural Health Monitoring Applications

In this section we present the Modular Monitoring System (MMS), a low-power wireless architecture dedicated to Structural Health Monitoring (SHM) applications. Our solution features an easily customizable modular architecture, fulfilling the needs of many SHM applications. The MMS supports mesh network topology

and offers excellent coverage and reliability, taking advantage of Wireless Sensor Networks (WSN) technology. Later in this same section we present the research and development activity made on the MMS.

Structural Health Monitoring (SHM) deals with the detection of damages to which civil and industrial structures, such as roads, bridges, canals, buildings and aero-space vehicles are subjected to. It can prevent collapses and breaks, avoiding permanent damages to structures, thus simplifying and improving the effectiveness of their maintenance. Depending on application scenarios, SHM requires many different types of sensors, including pressure sensors, vibrating-wire strain gauges, inclinometers, crackmeters, etc.

Nowadays, most SHM systems in the market are wired. However, the deployment of a wired system in a wide area or in a harsh environment, can pose both economical and practical limitations. For this reasons, WSNs have been proposed as an ad-hoc network infrastructure to support SHM, avoiding prohibitive costs of wired systems and easing the on-field deployment. Nowadays, SHM supported by WSN is an active and well-established research field and some wireless SHM systems are now entering the market. The work presented in this chapter introduces the MMS, a wireless, low-power, scalable hardware architecture dedicated to SHM. A key feature of the MMS is its high modularity that allows easy customization of the platform depending on the number and the type of sensors required by the specific application scenario. In addition, MMS fully supports multi-hop wireless communication paradigm and mesh networks.

The remainder of this chapter is organized as follows: Section **??** presents the state-of-the-art on both wired and wireless SHM systems. In Section **??** we explain the motivations supporting the design choices made during the development of the MMS, while in Section **??** we introduce the features characterizing our system. In Section **??**, we present the hardware prototype while in section **??** we report tests and validation results.

3.2 State of the art

Nowadays, most SHM systems available on the market, such as the Geomonitor by Solexperts [**?**], are wired. However, deploying those systems can be cumbersome: besides the installation costs, a detailed deployment plan is required to face evolving needs of different construction phases. Moreover, cables are subjected to accidental cuts and damages and, in some scenarios, their installation is infeasible or inappropriate (e.g., historical buildings). Along with wired systems, some standalone data loggers dedicated to SHM are commercially available, among the others: Solexperts SDL [**?**], Geokon 8002-16-1 [**?**] and Keller GSM-2 [**?**]. These systems are simpler to install, but they do not allow real-time remote monitoring and require frequent in-situ access by qualified personnel to collect data.
The introduction of wireless communications in SHM gives immediate advantages in terms of easier deployments and reduced maintenance and personnel costs. However,

when monitoring devices are battery-powered, the use of wireless communications is among the most energy demanding feature that can significantly limit the devices lifetime. Prominent examples of wireless monitoring systems are: National Instruments Wireless Data Acquisition (WiFi-DAQ) [**?**] and MicroStrain's wireless sensor network [**?**]. While the former supports IEEE 802.11 standard, the latter is compliant with IEEE 802.15.4. Both products adopt a conservative approach by limiting wireless communication to 1-hop. Supporting multi-hop wireless communications was investigated in several research papers in the last decade [**?**, **?**, **?**]. Multi-hop networking provides a number of advantages: scalability (larger areas can be covered), reliability (failure and multiple routing paths without single point of failure) and ease of deployments (the presence of multi-hop relay nodes allows to bypass obstacles like walls and metal structures). Recently, National Instruments presented the NI WSN [**?**]: a multi-hop battery-powered WSN supporting up to 36 nodes configured in a mesh network and up to 3 years node lifetime.

3.3 Motivations

As seen in the previous section, WSNs are slowly entering the market of SHM applications. An attempt to develop a robust solution and to test it in realistic application scenarios was made in the GENESI Project [**?**]. The main goal of that project was to design and implement a "novel generation of green wireless sensor networks which can be embedded in buildings and infrastructures at the time of construction and be able to provide a monitoring and control intelligence over the whole structure lifetime". The project involved the monitoring of a bridge construction site in Fribourg [**?**] and the construction of a tunnel for the B1 underground line in Rome [**?**], by means of WSNs. The outcomes of these experimental activities, highlighted the advantages of WSN technology compared to old monitoring techniques. According to the application requirements provided by the SHM experts, GENESI nodes can support a number of heterogeneous sensors. However, a GENESI node can manage only a single sensor per type, while there are some applications in which multiple instances of the same sensor are needed. As an example, in a 3-axis deformation analysis, a single wireless node may need to interface with three instances of a vibrating-wire strain gauge while to monitor a concrete junction the node may need to interface a current-loop inclinometer sensor and a resistive displacement sensor. Other commercial solutions, such as the NI WSN described in Section **??**, can handle multiple instances of the same sensor but can not support different sensor families at the same time. In general, the development of new ad-hoc devices addressing the specific requirements of an SHM deployment is impractical, while the adoption of commercial solutions in such contexts is not optimal in terms of flexibility, size and costs. This brought us to propose a novel low-power slotted modular system made by a set of modules connected through an internal communication bus. This solution features one wireless master module managing a group of extension modules, each one designed to interface a specific sensor set. The flexibility of the proposed architecture, named MMS, allows us to support a vast number of SHM applications by simply plugging into each node the required extension modules. By changing the master module, it is also possible to easily modify the wireless technology, thus

effectively adapting to the heterogeneous needs of indoor and outdoor communication requirements.

3.4 System architecture

The original MMS architecture presented in [**?**] was based on a master/slave communication abstraction where "a *master module*, provided with radio capability and responsible for most of computational tasks, communicates through a low-power shared bus with a maximum of four *extension modules* (slaves)". The new architecture embraces the same principle but requires the design of a new low-power shared bus.

The low-power shared bus in [**?**] was based on the Serial Peripheral Interface (SPI) Protocol with dedicated Chip Select (CS) and interrupt lines for each module. The rationale for this choice was to minimize the power consumption of the whole platform by allowing the master module to selectively activate each extension module acting on the corresponding chip select line, thus, without affecting the power consumption of the other modules installed on the platform. Despite the effectiveness of such solution from the power consumption perspective, the elevated number of dedicated lines (2 shared lines for the power supply, 3 shared lines for the SPI and 2 dedicated lines, CS and interrupt, for each extension module) forced us to develop a backplane board with a pre-defined number of slots at design time. In the original release, shown in figure **??**, we chose to support one master module and 4 extension modules, as a good trade-off between size and modularity.

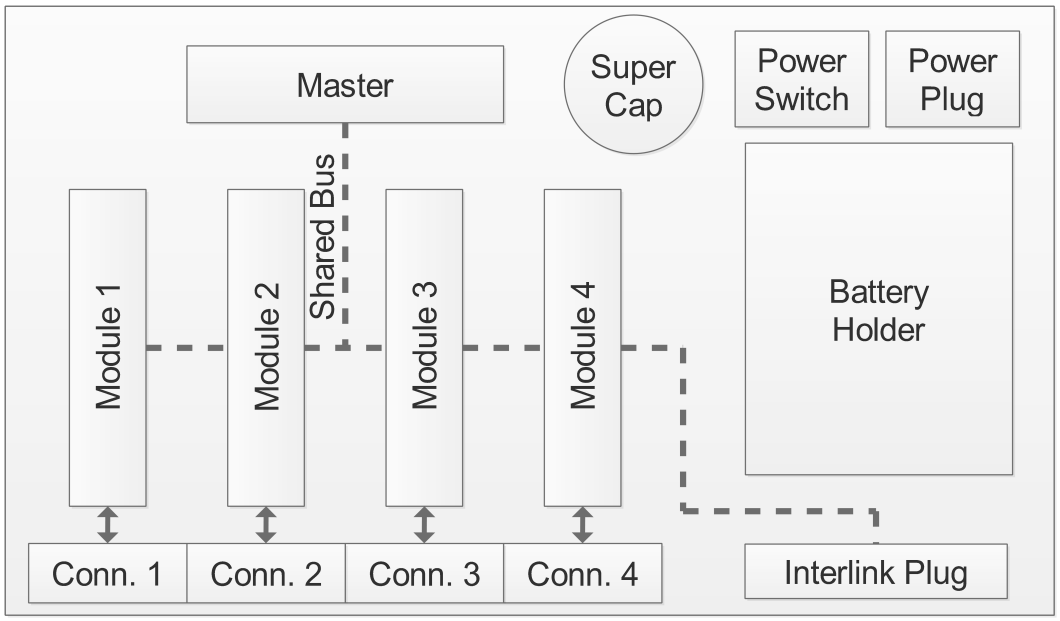

Figure 3.1. First version of the Modular Architecture.

This choice quickly revealed its limits: the additional backplane board increased the overall size of the platform and its costs for the electronics, and it also required extra costs for the ad-hoc housing. To overcome these limits we took hints from

existing modular housing solutions available on the market and we updated our design as described below. We found out that most of the considered modular products are based on O type DIN rail standard [**?**], a widely used solution for mounting circuit breakers and industrial control equipment inside racks. To ease the communication between modules, such solutions commonly provide a 5 lines shared bus. Hence, we implemented our system in a modular DIN housing and we redesigned the low-power shared bus to be compliant with the reduced number of lines offered by DIN bus. This led us to the new MMS architecture described in the following sections.

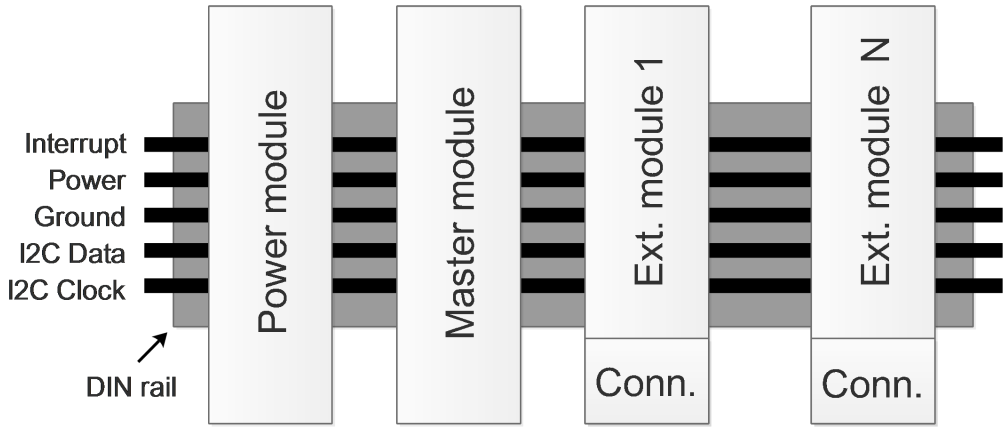

Figure 3.2. New Modular Architecture.

3.4.1 Low-Power Shared Bus

The availability of only 5 shared lines in the bus drastically reduces the implementation options for a master/multi-slave communication system. Excluding power and ground lines, the low-power shared bus has to be implemented on the remaining 3 shared lines. To overcome this constraint, we used the Inter Integrated Circuit protocol (I^2C) which is a multi-master/multi-slave serial communication protocol based on two open-drain lines bus [**?**]. The communication protocol, which is master initiated, is based on the transmission of a 7 bit address, i.e., the slave address to which the master wants to communicate, followed by a read/write bit and a set of protocol dependent commands. The third open-drain shared line has been used by the extension modules to trigger the master to start a communication. The master module polls to slaves to identify which extension module triggered the interrupt on the third line. Figure **??** shows the new architecture with a detailed view of the new low-power shared bus implementation.

This new solution easily scales with the number of extension modules (the standard implementation of the I^2C bus is designed to support up to 127 devices), thus increasing the flexibility and optimizing the size of the MMS. However, from the power consumption perspective, we need to pay attention to several drawbacks:

- **Wake-up.** When the master module issues an I^2C start command, all the extension modules wake up to perform the address matching, even if the communication is addressed to one module only.

- **Polling.** When an extension module starts a communication, it notifies the master using the interrupt line. Subsequently the master polls all the devices to find the initiator.
- **Latency.** The I^2C protocol needs more data to be sent over the bus with respect to the SPI one, this increases the latency and the power consumption of the MMS.

In section **??** we analyzed each of these aspects measuring the power consumption overhead related to this new implementation with respect to the original low-power shared bus. The results show that the overall consumption of the MMS is marginally affected by this choice.

3.4.2 Master Module

As in the first version of the MMS, the master module is responsible for managing the wireless connectivity within the WSN and manages the extension modules. When switched on, the master performs a discovery routine for dynamically assign the I^2C addresses to each extension modules and to retrieve the configuration information from them. Then, the master computes the sensing schedule and switches to the operational mode. Below a typical master-initiated interaction with the extension module is described:

1. The master issues an I^2C start command over the bus and sends the address of the extension module followed by a write bit and a data request command.
2. The extension module wakes-up and starts the conversion while the master enters a low-power state, waiting for an interrupt.
3. The master wakes up when an extension module pulls down the interrupt line.
4. The master polls all the modules searching for the initiator.
5. When found, the master sends a *Data Read* command to retrieve the new data.
6. When the read finishes the extension module puts the interrupt line in high impedance and enters sleep state again.
7. The master module checks the interrupt line level for other extension module that are willing to communicate. If the line is still low, it jumps to step 4 otherwise it goes back to sleep.

3.4.3 Extension Module

The extension modules provide the hardware interface to external sensors. Each module communicates over the low-power shared bus with the master by means of an I^2C-capable microcontroller. When powered-on, the extension modules participate to the I^2C address assignment managed by the master module. Once an extension module gets an address, it switches to a sleep state and wakes up only upon the detection of an I^2C start command on the bus. The extension module can initiate a

communication by pulling down the shared interrupt line of the low-power shared bus, which activates the polling procedure by the master.

Each extension module provides a register configuration area that enables the interaction with the master module. The register area is divided in 5 subsections as follows:

- The **Information area** [**read-only**] stores information such as device type and revision. This area allows the master to identify the extension module, e.g., sensors supported, channels available, etc.
- The **Command area** [**write-only**] is used by the master to trigger commands to the extension module such as a read command.
- The **Settings area** [**read/write**] stores settings for each sensor channel such as the sensor type connected, the periodic sampling value, etc.
- The **Channel flags area** [**read-only**] is used by the extension module to notify that a new data is available.
- The **Channels data area** [**read-only**] is used by the extension module to store the last sensor value read.

An example of extension module register implemented on our test board is shown in table **??**.

The hardware interface provided by each module depends on the sensors it supports: it can be fully digital, e.g., to interface RS-232 or RS-485 peripherals, or analog, to connect sensors such as current-loop, vibrating wire strain gauges, resistive etc. Each extension module can support different kinds of sensors or several sensors of the same type.

3.4.4 Power Module

The modular design of the new MMS allows us to support a number of interchangeable power modules which provide the power supply to the whole MMS through the power line on the bus. The modules support any kind of battery type providing a voltage between 3 and 5.5V as well as the 240V AC and energy harvesting solutions. Notice that this flexibility was not possible in the previous release of the MMS, since the power unit had to be integrated into the backplane at design time.

3.4.5 Additional extensions

We introduced the interlink bus which enables the MMS to support a virtually unlimited number of extension modules. In the specific, it connects the master modules of different MMS creating a tree structure where the parent MMS acts as master for its children.

To do so, the parent MMS uses the extension slots to make a connection through the interlink plugs of its children (see Figure **??**). Hence, the parent sees its children as enhanced extension modules. This mechanism can be replicated at each level of the tree to support an unlimited number of sensors. Each leaf node advertises the

Table 3.1. Extension module register area example.

Address	Register
0x00	Device type
0x02	Device description
0x03	Cmd channel 0
0x04	Cmd channel 1
0x05	Cmd channel 2
0x06	Cmd channel 3
0x07	Configuration channel 0
0x0e	Configuration channel 1
0x17	Configuration channel 2
0x1f	Configuration channel 3
0x27	Data flag channel 0
0x28	Data flag channel 1
0x29	Data flag channel 2
0x2a	Data flag channel 3
0x2b	Data channel 0
0x3b	Data channel 1
0x4b	Data channel 2
0x5b	Data channel 3

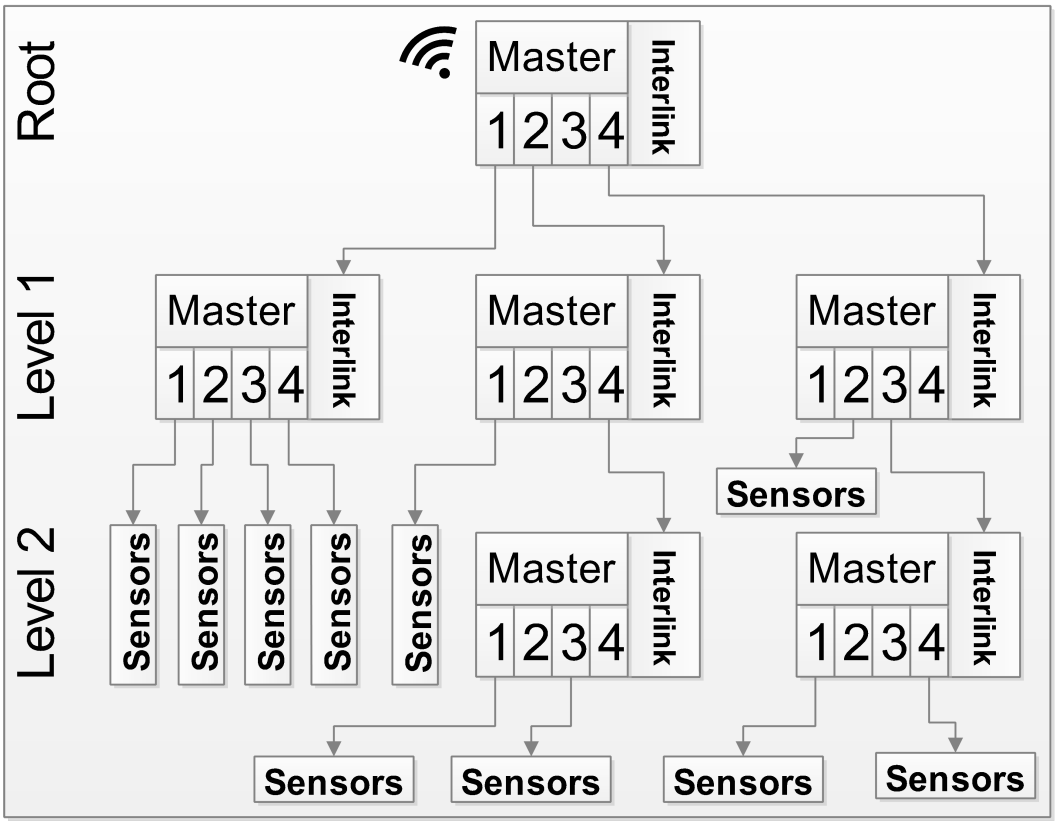

Figure 3.3. Interlink connections

connected sensors together with its identification number to its parent which aggregates the received information and forwards it to the upper level. This procedure is iterated until the root MMS is reached. The root MMS is the only one which manages wireless communications. From a WSN perspective, this makes the whole tree resulting as a single wireless node.

Thanks to the discovery procedure performed by the master during the boot phase, extension modules attached to the backplane exchange configuration information. Hence, an operator can add an extension module at any time by just power-cycling the MMS, without the need of an in-situ reprogramming. This feature eases the effort required by operators and qualified personnel. Based on our on-field experience, this functionality is often required in evolving environments, such as (but not limited to) SHM scenarios.

3.5 Hardware and firmware development

We designed, developed and assembled all the versions of the MMS at the Department of Computer, Control and Management Engineering of the University of Rome "La Sapienza". The the first version (shown in figure **??**) included a master module, a Resistance Temperature Detector (RTD) extension module and a backplane. The modules are plugged into the backplane by a standard Peripheral Component Interconnect Express (PCIe) edge connector.

Instead, the new version includes one master module, two demo extension modules and two power modules. As opposed to the first prototype, this version fits in a complete modular housing solution based on pluggable slots and snap-in connections for external sensors (figure **??**). The whole system was housed in a steel IP66 enclosure with a mounted DIN rail, an external antenna and IP66 cable glands.

Due to scarce of time, the MMS has some limitations. We suggest the development of several additional extension modules able to support most of the sensors used in SHM, such as strain wire gauges, crack meters, inclinometers, displacement sensors, etc. and additional master modules capable to support multiple wireless frequencies (e.g., 433Mhz, 868Mhz, 915Mhz) and certified industrial wireless protocols like Wireless Hart.

3.5.1 Master Module

The master module is responsible for radio communication and for managing the extension modules. It is based on the MagoNode OEM [**?**], a wireless hardware platform specifically designed for WSN applications. The MagoNode is based on the Atmel Atmega128RFA1 (RFA1) System-On-Chip microcontroller (MCU) equipped with 128KB of ROM (Read Only Memory), 16KB of RAM (Random Access Memory) and an embedded 2.4Ghz radio transceiver fully compliant with the 802.15.4 standard. The radio range is extended through a power efficient RF (Radio Frequency) front-end which enhances radio performance while keeping the power consumption low. The main figures in terms of power consumption are: radio

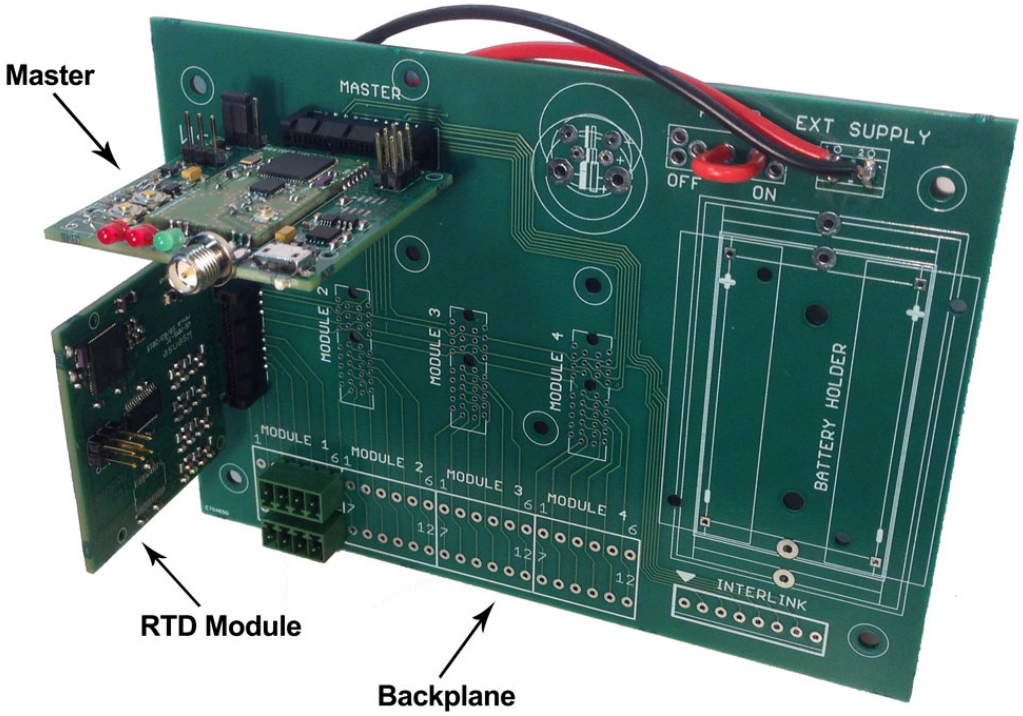

Figure 3.4. The old MMS with a master module, a RTD module and the interconnecting backplane.

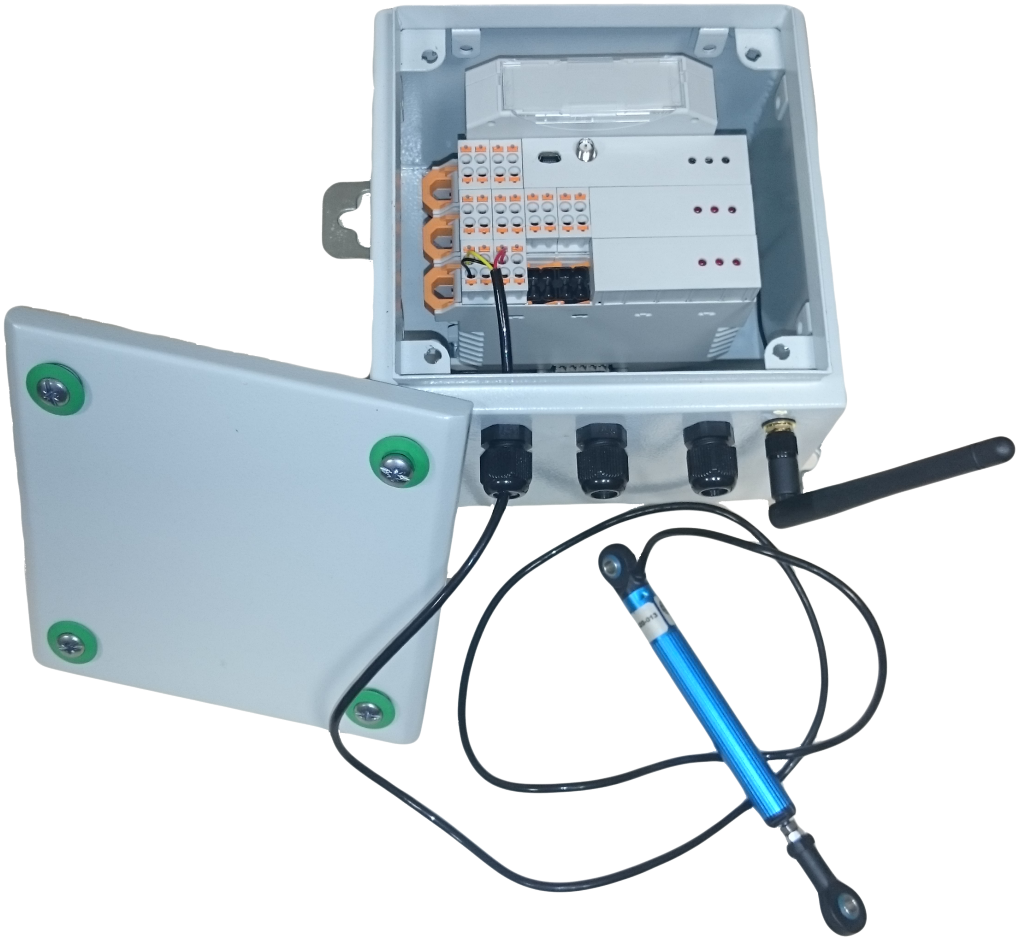

Figure 3.5. The newest version of MMS with a power module, a master module and two extension modules.

transmission 27.7mA @+10dBm, radio reception 14.5mA and <2μA in sleep. The master module is equipped with a NOR flash for persistent data storage, three status leds, a mini-usb plug and an RP-SMA (Reverse Polarity SubMiniature version A) antenna connector. An optional wireless HART [**?**] communication module can be installed based on application requirements.

The firmware is written using TinyOS [**?**], an event-driven, open-source operating system (OS) dedicated to Wireless Sensor Networks. TinyOS is designed to cope with typical constraints imposed by Wireless Sensor Networks: low computational capabilities, limited memory and scarce energy resources. TinyOS also provides a set of libraries implementing low-power wireless protocols for medium access control and routing. The implementation of the MMS firmware on TinyOS allows an easy and effective integration of our system in large-scale WSNs deployments.

3.5.2 Demo Extension Module

We developed a general-purpose extension module for demo applications which provides 4 24-bit analog channels, 2 excitation current outputs and 2 external voltage references. The device logic is handled by an efficient ARM Cortex M0+ MCU. The MCU provides the I^2C hardware interface and the interrupt channel required by the low-power shared bus communication, an SPI interface for the 24-bit ADC and a set of General Purpose Input/Output (GPIO) pins for driving the 3 status leds of the module. As opposed to the first MMS prototype which could switch off the power of the extension modules, in this case the low-power shared bus has an always-on power line. This forces all the extension modules to remain (by default) in a sleep mode when inactive: they are enabled only when an address match event on the I^2C peripheral occurs. The demo extension module supports a sleep mode with a current consumption of 1.5μA[1]. As the master module, the firmware was implemented using the TinyOS operating system which provides the primitives for I^2C and SPI communication and for handling the sleep mode of the device.

3.5.3 Power Module

We currently support two basic power modules: an AA and D type 3.6V thionyl chloride battery. A module to power the MMS via the 240V AC is in an advanced development stage and we are also working on the integration of an energy harvesting module supporting solar cells.

3.6 Experiments

In our experiments of the new MMS design, we first measured the power consumption during a local data acquisition (i.e., with wireless radio disabled), then we evaluated the system in a Wireless Sensor Network testbed. The former measurement validates the effectiveness of the low-power design in our architecture, while the latter demonstrates how the MMS features energy consumption levels suited for WSN

[1] An errata specific to the MCU revision used in the extension module avoided to reach this value which refers to the newest chip revisions.

applications. Finally, we performed a set of measurements aimed to demonstrate how the energy consumption overhead introduced by the new design of the MMS (see section **??**) is negligible if compared to the overall system consumption.

3.6.1 Local Data Measurement

All the measurements were done powering the MMS system with a Rigol DP1308A programmable DC power supply, providing $3.0V$. The MMS was connected in series to a Rigol DM3068 digital multimeter that sampled the current consumption at $10kHz$. The measurements were taken from an MMS made of a master module and a single demo extension module connected to a displacement sensor, i.e., a potentiometer with $4k\Omega$ series resistance. We performed a single master-initiated sampling request which consists in the following steps:

1. The master module sends the address of the extension module and sets a bit in the command area register to trigger one sample from a single channel on the selected extension module.

2. The extension module starts the conversion while the master goes back to sleep waiting for an interrupt.

3. The extension module generates an interrupt to the master notifying that it has a new data sample available.

4. The master polls all the modules, reading the data flag area. Whenever a flag notifies new data available, the master retrieves the new data.

In a second experiment, we repeated the same procedure using an additional demo extension module connected to an identical transducer where the master performs a sample request from each module. The current consumption of both tests are shown in figure **??**. The data conversion time, which is equal to 50ms, is determined by the sampling rate of the ADC which was configured to 20 samples per second. As expected, the current consumption during data conversion is doubled in the two extension modules configuration since each module performs the conversion at almost the same time. In addition, it is clearly visible how the sample request and data retrieval phases last twice in the two extension modules configuration since the master module needs to perform these activities sequentially for both modules. In both MMS setups, the system automatically switches back to sleep mode as soon as the data retrieval phase is completed.

Table **??** summarizes the average current and energy consumption of each phase measured on both configurations. Note that for the two extension modules configuration there are overlapping effects between phases: when the master module sends the start conversion command to the second extension module, the first module is already converting. Similarly, during the data retrieval phase, when the master module starts to download the new data from the first module, the second one is still converting. This explains the differences in the average current consumption during the sample request and data retrieval phases of the two modules configuration with respect to the single one.

Table 3.2. Local measurement - consumption details.

State	Single ext. module		Double ext. modules	
	Avg current (mA)	Energy (μJ)	Avg current (mA)	Energy (μJ)
Sleep	0.0035	N/A	0.005	N/A
Sample request	3.65	28	5.23	67
Data conversion	1.69	281	3.36	543
Data retrieval	2.79	104	3.57	178
Total	1.95	413	3.51	788

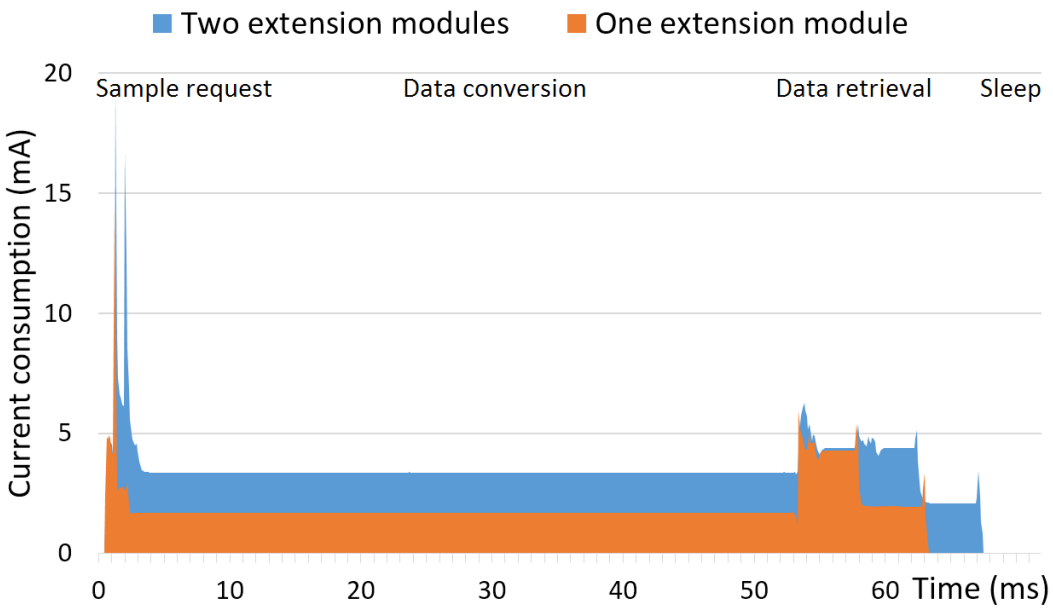

Figure 3.6. MMS current consumption during local data acquisition.

3.6.2 WSN testbed

We tested the two extension module configuration, described in the previous section, in an indoor wireless sensor network deployment located at the basement of our department. The testbed setup (figure **??**), consisted in 10 IRIS wireless devices [**?**] building a multi-hop wireless sensor network together with the MMS. Each IRIS mote was configured to generate one packet of 88 bytes every minute and to transmit it over the network. The MMS was set to locally sample one channel from each extension module generating 2 sensor readings every minute in order to transmit the sampled data over the wireless medium at the same rate as the IRIS motes. All data packets are routed through the multi-hop network toward a gateway, which is responsible for data aggregation and storing. We used the Collection Tree Protocol [**?**] as routing algorithm and the BoX-MAC medium access control protocol [**?**], both provided by the TinyOS operating system. To save energy, the MAC protocol was set with a radio duty cycle of 2% which represents a common value when long lasting wireless sensor networks are deployed. The testbed ran for 5 hours and collected statistical information embedded in the packets transmitted every minute by each node.
Table **??** reports detailed power consumption statistics, averaged over all the IRIS motes and compared with the one of the MMS. We observe that the power consumption of the MMS is similar to the one of the IRIS motes. Despite these measurements strongly depend on the position of the MMS in the network topology, this similarity shows that our solution is comparable in terms of energy requirements to a common WSN platform. Furthermore, the energy consumption of the sampling activity is less than 1% of the overall consumption that is largely dominated by the radio activity. We are aware that different transducers (e.g., the current-loop family) can consume more energy during conversion. This would significantly impact the fraction of energy consumed by the sampling activity. However, this consumption is exclusively related to the adopted transducer and does not depend on our architecture.

Table 3.3. Testbed - consumption details.

State	MMS Platform	IRIS Platform (AVG)
Energy Sleep (J)	0.26	0.13
Energy Radio Tx (J)	4.56	4.79
Energy Radio Rx (J)	1.23	1.28
Energy Radio Idle (J)	19.84	25.65
Energy Sampling (J)	0.23	0
Total	26.13	31.85

3.6.3 Design Overhead

As discussed in section **??**, we adopted a new design for the low-power shared bus, increasing the flexibility and reducing the cost of the MMS. However, these advantages come at the expense of a higher management complexity of the system that caused an increase of the power consumption. Obviously, the higher is the number of extension modules, the higher is the complexity, thus the power consumption overhead increases

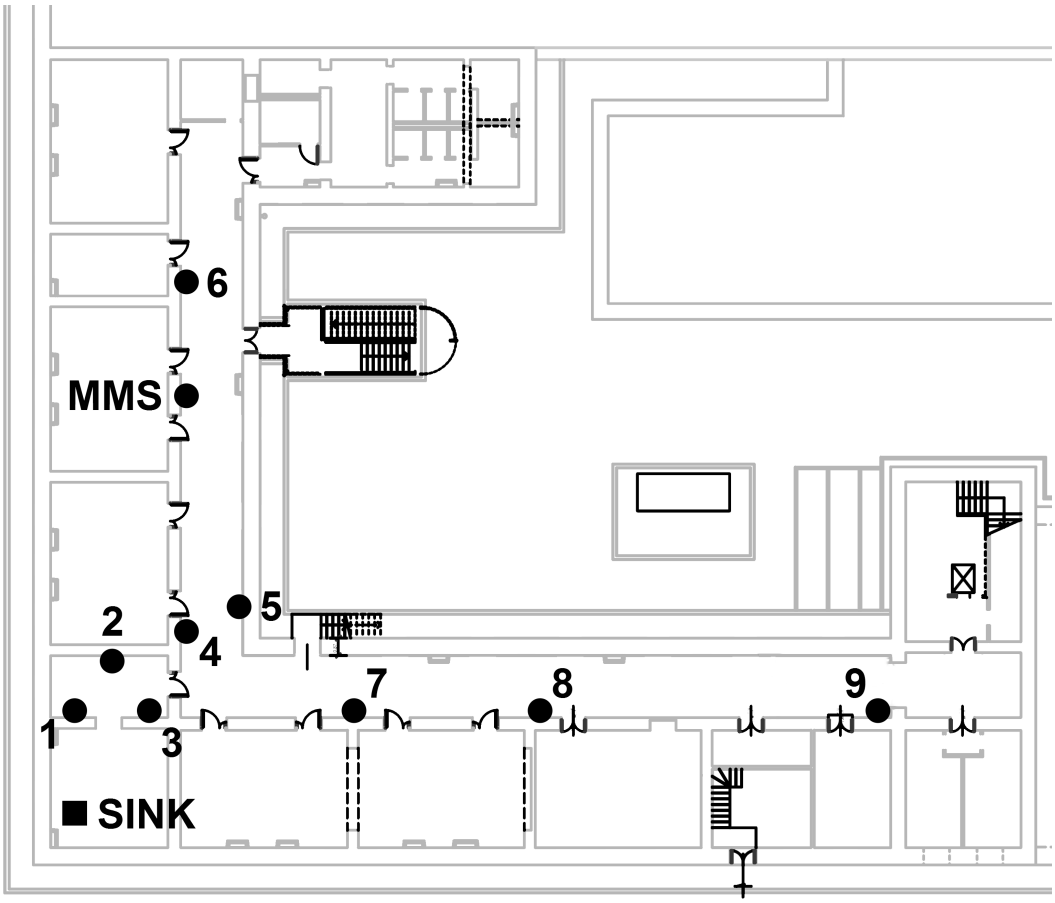

Figure 3.7. Testbed map.

accordingly. In this section, we evaluate the overhead of the new MMS design in terms of power consumption when compared to the original design.

Wake-up, *polling* and *latency* are the three main factors that cause the consumption overhead. We consider a configuration made of a master module and 4 extension modules performing a local measurement cycle as the one described in section **??** to measure the contribution of each of those factors to the overall power consumption.

Wake-up

When the master module starts an I^2C communication, it writes on the bus the address of the module it wants to communicate with. This activity wakes up all the extension modules that start an address matching procedure. Thanks to the new architecture of the M0+ MCU of the extension modules, the address matching procedure is carried out only employing the I^2C peripheral integrated in the chip, without the need to awake the MCU (this mode is called *sleepwalking* [**?**]). This feature wakes up the MCU only upon an address match event, reducing the overall power consumption of the address matching procedure. At the standard I^2C 100kbit/s datarate, the address transmission requires 90μs. Assuming 3.0V and 50μA current consumption of the sleepwalking peripheral, the energy required by an extension module to perform an address matching procedure on the low-power shared bus is 14nJ. Considering that for each request there are 3 extension modules that do not match the address, the overall energy wasted for each request is 42nJ.

Polling

The master module polls each extension module by looking at the data flag area which notifies whenever a new data is available. On the low-power shared bus this operation is performed as in figure **??**: the master first send a write request followed by the register address of the data flag area and successively a read request. The extension module replies with the content of the data flag area and, if no data is available, the communication ends. The whole procedure, which is 8 bytes long, is performed in 840μs on a standard 100kbit/s datarate I^2C channel. The measured current consumption of the MMS during a polling procedure is 4mA that, assuming 3.0V supply, brings to an energy of 10.1μJ per extension module. In the configuration under test, there are 3 polled extension modules over 4 that do not provide new samples, hence, the wasted energy is 30.3μJ .

Latency

We observed that the new design of the low-power shared bus requires, on average, 6 additional control bits each 10 data bits transmitted with the old implementation. This constraint increases latency in communications and, in turn, power consumption. The reason why more control bits are needed relies on the more complex I^2C bus management compared to the SPI protocol used by the old design. Assuming a local data measurement procedure as the one described in section **??**, the protocol transmits 6 additional control bytes. Considering 4mA current consumption of I^2C communication, 3.0V supply and a standard 100 kbit/s I^2C datarate, the additional energy for each extension module measurement introduced by latency is 5.8μJ.

Note that we are not considering the latency introduced by the different supported datarates of each protocol since it is limited by noise immunity requirements.

Based on these measurements, the overall energy overhead of the low power shared bus in a local measurement procedure is 36.14μJ which represents a consumption overhead of 9% when compared to the overall energy consumption measured in a single extension module configuration of table **??**. Recalling that the sampling activity accounts for less than 1% of the overall energy consumption of the MMS, we can conclude that the energy overhead of the new MMS design is totally negligible.

Figure 3.8. Polling I^2C communication detail.

3.7 Low power monitoring device for IoHT applications

Starting from the technology shown in the previous sections, we decided to design and implement a low power board for IoHT applications. We know from Section **??** that ideally the personal health information should be encrypted from the source of data (thus, from the lowest level devices) and we also know that these devices must be as low power as possible. Since we wanted as much control as possible on both hardware and software, we decided to realize by ourselves the device shown in Figure **??**, called *"RealTMS"*.

The RealTMS measures about 5cm by 7cm and it is powered by two alkaline batteries type AA that guarantee long lasting operation. The heart of the device is an ARM Cortex-M4 ore-based microcontroller with 32 bits architecture clocked at 80MHz. It also integrates a DSP (Digital Signal Processing unit), 128kB of RAM and 1MB of flash memory and some advanced low power capabilities (like low power timer and low power UART interface - Universal Asynchronous Receiver-Transmitter). It is produced by ST Microelectronics (the exact model is STM32L476RG). We also added on-board sensors for: temperature, relative humidity, 3-axial linear acceleration, 3-axial angular velocity and 3-axial magnetic field sensor. These sensors are

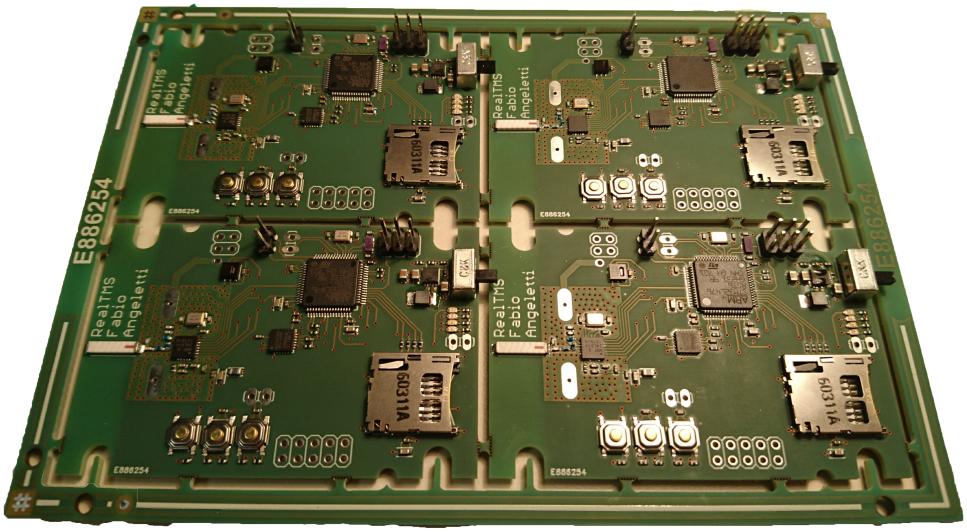

Figure 3.9. RealTMS

MEMS (Micro Electro-Mechanical Systems). The temperature and relative humidity sensor is produced by Sensirion (model SHT25) and has excellent performances: about two tenth of Celsius degree accuracy and 1.8% accuray in relative humidity measurement. It comes fully calibrated from factory. The 3-axial sensor for lineal accelerations, angular velocities and magnetic fields is the MPU9250 (produced by InvenSense) that integrates a complete IMU (Inertial Measurement Unit) within a single chip. The radio interface is provided by the NRF24L01+ chip produced by Nordic Semiconductor and is one of the cheapest radio chip available. It operates in the ISM (Industrial, Scientific and Medical) radio band at a frequency of 2.4GHz. Additionally, a slot for microSD card is provided in the need to store a very large amount of data to further processing.

The MCU (microcontroller) selected is supported by many operative systems for embedded systems, such as FreeRTOS, RiotOS, Mbed OS and more. In our application we used RiotOS firstly, then switched to FreeRTOS. The sensors on-board and the processing power of the MCU allow for complex activity recognition with the integrated DSP, such as walking, falling and many more. The temperature and relative humidity sensor permits a precise monitoring of the surrondings and produces high quality data that is possible to exploit for further processing.

Chapter 4

Conclusions

In an increasingly data-driven world, where information sharing, machine learning, and social networking leads the way, the IoT will be a key technology for person-centric mobile eHealth. This work looks into the usage of IoT technologies as an integral part of clinical trials so that data residing in an IoT world can enable drug developers to get better insights and streamline the overall clinical trial processes. The digital clinical trial assumed moves beyond existing solutions that follow a simplistic approach by limiting to the connection of patients to trials that they might be interested in. The study carefully examines the privacy issues that arise throughout all the phases of a *Digital Clinical Trial* and examines the pertinent security requirements.

The dominant approach for the delivery of IoT ecosystems is providing the majority of the services within cloud-based infrastructures. Consider for example the well established fitness applications[1,2] that use a smart band or smart watches to monitor users' physiological data (e.g., heart rate). It is evident that cloud-based IoT platforms simplify the interconnection of smart devices, the collection of data generated to the cloud, and the central processing of the information utilizing other cloud-based services. However, transmitting confidential data to cloud-based services and storing them over third-party infrastructure poses significant risks for the privacy of users. In principle, the only two entities that should handle the data of a clinical trial are the patient, the provider of personal data, and the *clinical research institute*, the consumer of the data provided by the community of patients participating in the clinical trial. In this dissertation, several approaches have been presented to deliver systems that enable the secure exchange of data between the two parties (i.e., patients and *clinical research institutes*) even when other third parties are involved.

This work aims to protect confidentiality by controlling disclosure of private data. However, this protection may turn out to be insufficient in the light of predictive techniques such as data mining: not only should data be protected but the predictions mined out of these data, i.e., data patterns should be filtered too to forbid their mischievous use as the possession of data patterns and a few medical data about someone may allow to predict part of their medical future and treat

[1] `https://www.fitbit.com/`
[2] `https://www.strava.com/`

them accordingly [**?**].

We applied a random-projection (RP) approach to privacy-preserving data mining of medical data, demonstrating the usefulness of RP in increasing privacy of personal health data handling. The projected data are useful for machine learning algorithms (for example, in clustering) while allows the sharing of information between parties without revealing the patients' clear data. In this particular application, this is of notable importance since allows entities involved in different health branches to cooperate effectively without sharing clear data. Second, we investigated to what extent an attacker can discover additional information starting from leaked data. As long as the projected space is smaller than the original space, and as long as the amount of data leaked is small, than the proposed approach is robust and mantains very good performance in both accuracy and privacy. We also analyzed the ratio behind and the performances (in terms of accuracy) of the RP applied on sensible healthcare data. The results show that the use of RP offers great enhancements in privacy protection. This was a first step into developing a full-fledged platform that allows the effective share of medical data. A proof-of-concept implementation was provided to evaluate the proposed technical solutions and their performance within a complex distributed system. The results indicated that by suitably combining blockchain and IoT technologies it is possible to design a digital clinical trial.

In the Chapter **??**, we presented the initial realization and the further developments of the Modular Monitoring System, a novel low-power wireless modular architecture designed for SHM applications. Our platform, the MMS, takes advantage of a low-power shared bus connecting slotted extension modules that interact with a master in a master/slave communication abstraction. The extension modules, which can be combined as needed, allows the MMS to face the continuously evolving needs of most SHM scenarios. Thanks to its peculiar characteristics, the MMS overcomes commercial state-of-the-art WSN solutions for SHM, like the NI WSN, which do not offer enough flexibility to fulfill requirements of many application contexts in a both cost-effective and efficient way. The newest implementation overcomes the size and costs issues of the first prototype presented in [**?**] offering an improved architecture based on DIN standard modular housing. This solution also offers an increased flexibility by supporting interchangeable power sources. We validated the effectiveness of the system low-power design, performing energy measurements during data acquisition from actual transducers. In addition, we tested the MMS within a real WSN deployment, to demonstrate the compliance of the system in such applications.

In the last Section **??** of the previous chapter we introduced the RealTMS, the embedded system that we refined in order to validate the architecture presented in Chapter **??**. It proved to be capable in handling low power task and signing packet involved in a digital clinical trials scenarios. However the lack of specific sensor for biometrical data (such as heart rate, blood pressure and other) and the size strongly limited its possible applications.

As future work, we suggest to incorporate techniques such as data randomisation in combination with data mining methods, thus limiting the accuracy of the results provided. This action will build stronger basis in enabling access to the private data without violating the privacy of individuals, while allowing to exploit its full potential. Moreover, we suggest to run an audit to check privacy protection against real third parties. Our device, the RealTMS, can greatly benefit from further miniaturization effort, in order to become a seemless wearable. It is strongly suggested to add more sensors on-board, such as heart rate, pulse oximetry and other. This will allow to retrieve better insights about the health of the patients.