Towards a Security Policy for Ubiquitous Healthcare Systems (Position Paper)

Joonwoong Kim, Alastair R. Beresford and Frank Stajano

University of Cambridge Computer Laboratory
15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom
{joonwoong.kim,alastair.beresford,frank.stajano}@cl.cam.ac.uk

Abstract. U-Healthcare promises increases in efficiency, accuracy and availability of medical treatment; however it also introduces the potential for serious abuses including major privacy violations, staff discrimination and even life-threatening attacks.
In this position paper we highlight some potential threats and open the discussion about the security requirements of this new scenario. We take a few initial steps towards a U-Healthcare security policy and propose a system architecture designed to help enforce the policy's goals.

1 Introduction

Granny Alice is so pleased with the special "U-Health Shirt" she received this weekend from her son Bob: it monitors her vital signs and sends them wirelessly to a medical centre. Thanks to this ongoing monitoring, she will be able to continue to live in her own flat instead of having to move into one of those horrible, crowded nursing homes. She feels safe, independent and empowered.

On Monday morning, on his way to his office, Bob checks his schedule on his PDA and is pleased but surprised to see that all his meetings have been cancelled. In the office, he finds Carol sitting at his desk: something feels wrong. "Didn't you get any messages?", she asks with a hint of embarrassment. He checks email and discovers that he has been transferred to the post room and that Carol, not him, is now leading the department. There is also another message, from healthcare services, booking him in for a detailed medical check-up: his body sensors show high levels of nitric oxide, suggesting the possibility of cancer.

Bob feels dizzy. He knows that, even if the check-up later reveals no cancer, he has now lost his prestigious position in the company. There are too many candidates for that post: his own turn may only come back in several years. While he drives back home, he nervously removes all of his body sensors, only regretting he can't easily get rid of the implanted ones. Soon his mobile phone rings: a voice mail tells him he should replace and reconnect his sensors or he will lose his insurance discount, and that he should contact customer service if this is a sensor failure.

Despite the obvious exaggerations, with the rapid evolution of sensor technologies most of the pieces of the above scenario are rather close to feasibility. Healthcare projects using body sensors as remote monitoring devices are already under way[1]. Body-sensor-based 24-7 monitoring will enable remote diagnosis without the patient having to visit a hospital, thus providing cheaper healthcare services. At the same time, more detailed body sensor data, combined with data from infrastructure sensors, will provide a "life log" or "activity diary" of the patient. If such sensitive personal data is shared among interested parties—employers, insurance companies, drug companies and the government, to name a few—the possibility of abuse is great.

Most readers (except perhaps those who were recently downsized) will point out that real companies can't afford to be as ruthless and nosey as Bob's if they wish to retain talent; but Bob might have given his explicit consent to his doctor about sharing his body sensor data with his employer, to prove to the employer that it was safe to promote him to division head because he was a healthy employee who would be able to work hard and handle severe stress levels. And he could have also granted access to his health insurance company for a discount. Or his employer might have done it for him, assuming Bob lives in a country where the employer customarily pays for the employee's health insurance.

The fact that the inappropriate disclosure of private medical information can harm the patient has been clear since at least the 4th Century BC, as the Hippocratic Oath indicates. But what exactly are the security requirements in this age of increasing computerization? Ten years ago, Anderson's BMA Security Policy [1] described the protection goals of clinical information systems: its motivation was that storing patient medical records on a nationwide distributed computer system endangered the principle of patient consent and increased the possibility of data aggregation. U-Healthcare, in turn, brings new threats and vulnerabilities, as illustrated by Bob's story above, which are not all adequately covered by the BMA policy.

The first contribution of this position paper is to point out such new threats and to open up the debate about security for U-Healthcare. Secondly, in order to clarify the protection goals, we propose and discuss some possible principles for a U-Healthcare security policy. Thirdly, we suggest a system architecture consistent with the proposed policy.

1.1 Terminology

Electronic Healthcare Systems, or *(Electronic) Clinical Information Systems* are the existing healthcare information systems that use networked computing systems for recording and accessing medical records. *Ubiquitous Healthcare Systems*, instead, adopt ubiquitous computing as an enabling technology, with sen-

[1] See for example the Codeblue paper [14] and the web sites of the PIPS, MyHeart and Proactive Health projects (`http://www.pips.eu.org/`, `http://www.hitech-projects.com/euprojects/myheart/` and `http://www.intel.com/research/prohealth/` respectively.

sors monitoring the patient continuously, and include Wellness Systems, Disease Care Systems, and Independent Living Systems [11].

We prefer to say *Patient* rather than *User* because a *Clinician*, too, is a user of the U-Healthcare system. There are several *Healthcare (Service) Providers*, including but not limited to *Clinicians* and GPs: for all of them we may also use the term *Caregiver*. The more general terms are preferable if we consider that the clinics can be replaced by other healthcare services such as gyms and healthcare web services.

As for sensor devices, there are body sensors and infrastructure sensors. Examples of the latter include scales and sensors of ambient temperature, light or movement. These sensor devices transfer data to base stations such as PDAs, Smartphones and PCs. The union of these sensors and base stations forms a *Personal Healthcare System* which is used and controlled by an individual patient, or by some trustee on behalf of the patient. The sensor data is then transferred to a *Clinical System* for further analysis, if needed.

Lastly, for economy of expression, we will use the same gender convention as the BMA Policy [1]: the clinician is female, and the patient male.

2 Threats and Vulnerabilities

Currently, most patient medical records are accessed through a standard desktop workstation which requires the caregiver to be in a particular place at a specific time. Therefore the environment and architecture of the hospital or surgery provides some additional social control to prevent unauthorised access to medical data. The use of PDAs and laptops in ubiquitous healthcare to access patient records on the move, or from a remote location, is likely to improve the timeliness of patient care, but may represent a greater temptation for an underpaid caregiver who is offered a bribe by a pharmaceutical company or a private investigator.

Many monitoring systems in hospitals today use physical access control to provide privacy. For example, a heart rate monitor is typically situated beside the patient, and the device only provides data to a caregiver who is standing near the machine. In addition, data might only be available in real-time—any historical data is lost unless a caregiver explicitly records it separately. Ubiquitous healthcare extends the computerisation of medical records to the domain of monitoring and diagnosis—monitored data will be recorded and the historical record used in subsequent analysis.

In a ubiquitous healthcare system, remote access to patient data by a caregiver may become normal. Because sensors will be cheap and portable, a personal healthcare system is likely to be used to record sensitive medical data continuously during everyday life, not just whilst the patient is in a hospital. This record of data will be of great interest to third parties, such as insurers, medical researchers and employers and therefore, without adequate control, the ability to data mine this resource becomes compelling. The recorded data is also likely to contain many personal facts (such as dates, times and durations of the patient's

sexual intercourses) which, whilst irrelevant to any specific medical diagnosis, are hard to remove from the dataset without reducing the quality of the sensor data itself.

It is also likely that a caregiver, or even a computer program, will be able to remotely administer drugs through a body area network. This scenario requires integrity of sensor data and rules used to decide when to administer drugs.

In current out-patient practice, a caregiver will typically engage in a short consultation with the patient and ask a series of questions about his health. The questions must necessarily be on a level that the patient understands. In this scenario the patient is able to query the relevance of any question with the caregiver, ask what the consequences of failing to answering the question might be and, if the patient feels it is necessary, provide a false answer. In contrast, ubiquitous monitoring of physiological signs will generate a large amount of data which the patient cannot interpret without help: the dataset will be too large for manual analysis, and it is likely to require a good deal of technical skill to understand.

The BMA Policy [1] was concerned that the aggregation of many patient records may lead to abuse. In ubiquitous monitoring, a combination of sensor readings of a *single* patient may also be problematic. For example, the symptoms of depression may be inferred from changes in body weight and sleeping patterns, even if this conclusion was not the original intention of monitoring. Disclosure of such medical histories might be unwelcome or used against the patient. For example, the medical history of a US politician who had suffered from depression was disclosed just before an election [18].

A patient may also configure a body sensor network to record data for other purposes. For example, Bell's MyLifeBits project [7] records a wide variety of audio, visual and location data which can be used to aid memory recall of specific events. The patient will probably not want to give unconditional access to these data, yet a caregiver may be able to give a better diagnosis with access to some information contained within the dataset.

The additional problems presented by a ubiquitous healthcare system can be summarised into four broad areas:

Ubiquitous access: easy remote access to data amplifies the vulnerability of medical records to unauthorised access;

Ubiquitous monitoring: monitoring and diagnosis will be computerised and sensors will travel with the patient wherever he goes, potentially providing the caregivers with the ability to record, search and archive sensor data remotely;

Ubiquitous care: patients will receive tele-prescription and tele-infusion of drugs and receive professional advice remotely;

Ubiquitous sensor data: The collection and recording of medical sensor data will be useful to researchers but may contain many personal facts.

For the extensive security analysis, we need to consider confidentiality, integrity and availability. In a sense, confidentiality is more related with privacy,

and the latter two with safety and dependability. However, in the remainder of this paper we focus on addressing the privacy issues of ubiquitous monitoring as a starting point. Because we believe this will become the most prominent part of ubiquitous healthcare systems in the near future, and is something which is missing from the existing discussions on security in healthcare systems, such as those found in the BMA Policy [1]. For the availability and integrity of healthcare systems including the ubiquitous care part might be remained for the future work.

3 Towards a security policy

3.1 Monitoring

Traditionally, health status was measured either directly by the caregiver or the patient; more recently such measurement may have received some form of technological assistance. Such collected data is usually analysed in real-time and is summarised and discussed before being recorded. In contrast, when a ubiquitous monitoring environment is used, computing devices may create a permanent record in much greater detail. To protect the privacy of the patient we propose:

> **Principle of self care**: Data collected in a ubiquitous monitoring environment must be processed and stored on a personal healthcare system under the sole control of the patient. No sensor data shall leave the personal healthcare system without the patient's consent.

This principle reflects our current notion of healthcare: a patient will contact a caregiver only after a medical problem is discovered and caregivers only receive medical facts from the patient or perform an examination with the informed consent of the patient.

In some cases we will want the ubiquitous monitoring environment to analyse, report and automatically execute actions based on the sensor data. For example, a diabetes patient may use a body sensor network to keep him informed of his current glucose level and perhaps even automatically trigger the delivery of insulin. In this case, the principle of self care means that glucose level readings and insulin delivery must operate within the personal healthcare system and run independently of all clinical systems under the control of the caregiver. It is worth noting that this type of autonomous operation may be sensible from a safety and reliability perspective too.

3.2 Consultation

There will be times when a patient will seek the advice of a caregiver. This might happen at pre-defined intervals, whenever the personal healthcare system reports a potentially life-threatening reading, or during an emergency. In these cases, the patient (or, in the case of the young or seriously ill, their next-of-kin)

will require some help interpreting the data recorded by the personal healthcare system. Since the patient cannot know what facts the sensor data contains, he cannot give his *informed* consent to the release of all sensor data directly into his medical record. Therefore, to protect the privacy of the patient we propose:

> **Principle of non-disclosure**: The patient may transfer sensor data from his personal healthcare system into a temporary repository which is also accessible by a caregiver. Only data useful in assessing the state of health of the patient is transferred. By default, data may not be transferred out of the repository, which shall exist for a limited time.

In practice it is impossible to delete all traces of the analysis since the caregiver and patient may mentally recall some of the information. Nevertheless, this principle means that at the end of any consultation between a patient and the caregiver, there should be no electronic record of either the raw sensor data or any derived data.

Some forms of analysis may require several caregivers to collaborate and this might make it difficult to arrange for all the specialists and the patient to meet at once. In this case, the principle of non-disclosure means that as the data is analysed, the patient is kept informed of what data is collected from his personal healthcare system. It is important to limit both the amount of time data can be kept, and the number of caregivers who may access the repository. If this is not the case, the lifetime of the repository may last as long as the lifetime of the patient, and it becomes a medical record in all but name. The length of time data can be held in a repository will depend on the medical condition under analysis; for complex situations this is something which needs to be reviewed by caregiver and patient at regular intervals.

3.3 Permanent records

The principle of non-disclosure means that, whilst caregivers can analyse data from a personal healthcare system, they cannot maintain a summary of the results of the analysis. Such a record may be needed to provide a prescription, charge a fee or provide continuity of care. We believe it is important that the patient controls and understands the meaning of any data which is written to a permanent medical record as the result of analysis in a temporary repository.

> **Principle of limitation and necessity**: Any results from the analysis of sensor data stored in a temporary repository may only be transferred into the patient's permanent medical record if the patient's *informed* consent for the transfer is obtained and the long-term storage of such data is required to protect the patient's future well-being.

Or in other words: record the outcome of the analysis (if it is relevant and useful) rather than the raw sensor data itself. In some sense this principle is nothing new: caregivers have previously summarised information written into a medical record rather than transcribing the entire conversation. The aim of this

principle is to prevent the raw sensor data from being written into a permanent record; this is important since raw data may contain lots of hidden personal facts the user did not consent to releasing, but may be obtained later by data mining.

In many cases, data may be summarised on the personal healthcare system itself. For example, a diabetic may provide the caregiver with a summary of the highs and lows of their glucose level. In other cases, the caregiver will need to see the raw data: an electrocardiogram (ECG) trace provides much more information than simply the heart rate—the data requires an expert to interpret it.

4 Architecture

In the last section we derived a security policy which provides the patient with a method to control access to any sensor data recorded by a personal healthcare system. We believe, from a computer science perspective at least, that it is practical to build a system which conforms to this security policy. In order to support a temporary repository, we envisage a software mediator which logically sits between a personal healthcare system used by the patient and any clinical system used by the caregiver. The concept of an intermediate component exists already in many diverse research areas of computer science, and includes proxies, agents, guardians, Trusted Computing Bases, etc.

The mediator (Figure 1) should provide an interactive environment in which a patient and a caregiver can explore the data recorded by the body sensor network, extract the relevant medical facts from the collected data and, with the patient's informed consent, append those facts to the medical record. In order to meet the criteria set in the security policy, it is important that the patient be in control and be able to limit: (1) the raw sensor data sent to the mediator and (2) the derived facts transferred from the mediator to the medical record. Obviously it is paramount that all data stored by the mediator be deleted at the end of any period of consultation.

5 Related Work

The BMA Security Policy [1] was developed by Anderson for the British Medical Association to protects patient records in clinical informations systems. It is based on nine principles, including access control, consents, audit, information flow and data aggregation. A few updates [2,3,4] also appeared.

In the 1990s, threats to privacy in Electronic Patient Record (EPR) were widely recognized in the U.S. As a result, a few reports [8,13,17] about security in EPR were released. Besides these works, most security research in healthcare systems [16,19] have been based on variants of the Role Based Access Control (RBAC) model. Gostin [9] discussed healthcare information from an ethical perspective, while Health Privacy Project [15] provided a small collection of privacy-breaching incidents in U.S. medical systems.

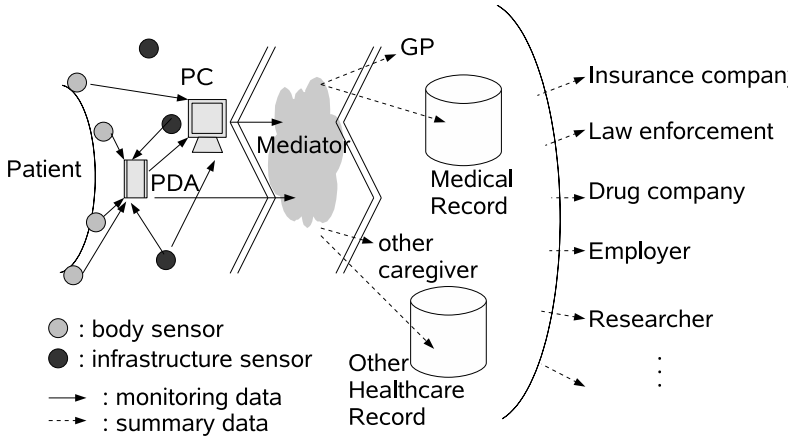

Fig. 1. Ubiquitous Healthcare Systems Architecture Option

Many researchers have worked on privacy in ubiquitous computing environments, including at least [12,10]. Langheinrich [12] proposed the infospace concept as the trust boundary, and the privacy tag. Jiang [10] discusses how the user is notified about data collection by sensors and how a policy can be negotiated. However these two privacy frameworks refer to a general ubiquitous computing or context-aware computing context and are not directly applicable to healthcare information systems. More relevantly, Bohn [6] analysed the dependability issues in U-Healthcare, and Beckwith [5] discussed the perception of privacy based on the case study of a sensor-rich, eldercare facility.

To the extent of our knowledge there has not yet been any proposal of a formal security policy to regulate ubiquitous healthcare systems, along the lines of what the cited BMA policy did for in clinical information systems. Hence our work.

6 Conclusion

U-Healthcare introduces great convenience, but at the same time equally great risk. The shift to 24/7 patient monitoring via body sensors is not just an incremental improvement over the existing practice: it is a qualitative step change. So is the shift to remotely-activated drug dispensers implanted in the patient's body. The main message of this paper is that such major paradigm shifts demand a rethinking of the security and privacy aspects: solutions that were appropriate for yesterday's situation are insufficient for tomorrow's. We pointed out some of the new threats.

We believe it is still too early to propose a complete technology solution: what is most needed at this stage is instead an informed debate. We wish to engage all parties, including clinicians and patients, and understand what is acceptable

and desirable before the coming generation of U-Healthcare systems is deployed. This is why we presented our principles in natural language rather than using equations or formal security terminology. There will certainly be tension between security and usability, between patient privacy and clinician convenience, and we don't presume to have got the balance exactly right at the first attempt; we solicit opinions and corrections, particularly from practicing clinicians, but we all need to understand the issues at stake.

In this context, a security policy is first of all an instrument of communication. By writing down, at least as a working draft, the protection goals of future U-Healthcare systems, we allow the community of stakeholders to think, agree, disagree and debate. We hope that the outcome of this process will be a strong specification upon which to build U-Healthcare systems that, like Isaac Asimov's brilliantly imagined robots, can never be misused to cause harm to their patients.

References

1. Ross Anderson. *Security in Clinical Information Systems*. BMA Report. British Medical Association, Jan 1996. ISBN 0727910485. `http://www.cl.cam.ac.uk/~rja14/Papers/policy11.pdf`.
2. Ross Anderson. "A security policy model for clinical information systems". In "IEEE Symposium on Security and Privacy", 1996. `http://www.cl.cam.ac.uk/~rja14/Papers/oakpolicy.pdf`.
3. Ross Anderson. "An Update on the BMA Security Policy". In "Cambridge workshop on Personal Information — Security, Engineering and Ethics", 1996. `http://www.cl.cam.ac.uk/~rja14/Papers/bmaupdate.pdf`.
4. Ross Anderson. "Healthcare Protection Profile — Comments", 1998. `http://www.cl.cam.ac.uk/~rja14/Papers/healthpp.pdf`.
5. Richard Beckwith. "Designing for Ubiquity: The Perception of Privacy". *IEEE Pervasive Computing*, **2**(2):40–46, 2003.
6. Jürgen Bohn, Felix Gärtner and Harald Vogt. "Dependability Issues of Pervasive Computing in a Healthcare Environment". In "Security in Pervasive Computing 2003", vol. 2802 of *Lecture Notes in Computer Science*. 2004. `http://www.vs.inf.ethz.ch/res/papers/bohn_pervasivehospital_spc_2003_final.pdf`.
7. Steven Cherry. "Total Recall". *IEEE Spectrum*, **42**(11), Nov 2005. `http://www.spectrum.ieee.org/nov05/2153`.
8. Paul D. Clayton (ed.). *For the Record: Protecting Electronic Health Information*. National Academy Press, 1997.
9. Lawrence Gostin. "Health Care Information and the Protection of Personal Privacy: Ethical and Legal Considerations". *Annals of Internal Medicine*, **127**(5), Oct 1997. `http://www.annals.org/cgi/content/full/127/5_Part_2/683`.
10. Xiaodong Jiang and James A. Landay. "Modeling privacy control in context-aware systems". *IEEE Pervasive Computing*, **1**(3), 2002. `http://guir.cs.berkeley.edu/projects/ubicomp-privacy/pubs/infospace.pdf`.
11. Ilkka Korhonen, Juhan Pärkkä and Mark Van Gils. "Health Monitoring in the Home of the Future". *IEEE Engineering in Medicine and Biology Magazine*, **22**(3):66–73, May 2003.

12. Marc Langheinrich. "Privacy by Design — Principles of Privacy-Aware Ubiquitous Systems". In "Ubicomp 2001", 2001. http://www.vs.inf.ethz.ch/publ/papers/privacy-principles.pdf.
13. William W. Lowrance. *Privacy and health research a report to the U.S. Secretary of Health and Human Services.* U.S. Department of Health and Human Services, 1997.
14. David Malan, Thaddeus Fulford-Jones and Matt Welsh. "CodeBlue: An Ad Hoc Sensor Network Infrastructure for Emergency Medical Care". In "International Workshop on Wearable and Implantable Body Sensor Networks", Apr 2004. http://www.eecs.harvard.edu/~mdw/papers/codeblue-bsn04.pdf.
15. Health Privacy Project. "Meidcal Privacy Stories", Nov 2003. http://www.patientprivacyrights.org/site/PageServer?pagename=True_Stories#True_Stories.
16. Jason Reid, Ian Cheong, Matthew Henricksen and Jason Smith. "A Novel Use of RBAC to Protect Privacy in Distributed Health Care Information Systems". In "Eighth Australasian Conference on Information Security and Privacy (ACISP 2003)", 2003.
17. Thomas C. Rindfleisch. "Privacy, information technology, and health care". *Communications of the ACM*, **40**(8), Aug 1997.
18. A. Rubin. "Records No Longer for Doctors' Eye Only". *Los Angeles Times*, 1 Sep 1998.
19. Longhua Zhang, Gail-Joon Ahn and Bei-Tseng Chu. "A role-based delegation framework for healthcare information systems". In "The Seventh ACM Symposium on Access Control Models and Technologies (SACMAT'02)", 2002.