

Privacy-Preserving Telemonitoring for eHealth

Mohamed Layouni¹, Kristof Verslype², Mehmet Tahir Sandikkaya³,
Bart De Decker², and Hans Vangheluwe¹

¹ School of Computer Science, McGill University, Montreal, Canada

² Department of Computer Science, K.U. Leuven, Leuven, Belgium

³ Katholieke Hogeschool Sint-Lieven, Gent, Belgium

Abstract. Advances in communication technology have opened a myriad of new possibilities for the remote delivery of healthcare. This new form of service delivery, not only contributes to the democratization of healthcare, by reaching far-away populations, but also makes it possible for elderly and chronically-ill patients to have their health monitored while in the comfort of their homes. Despite all of these advantages, however, patients are still resisting the idea of medical telemonitoring. One of the main obstacles facing the adoption of medical telemonitoring, is the concern among patients that their privacy may not be properly protected. We address this concern, and propose a privacy-preserving telemonitoring protocol for healthcare. Our protocol allows patients to selectively disclose their identity information, and guarantees that no health data is sent to the monitoring centre without the patients' prior approval. The approval process can be automated, and requires only an initial configuration by the patient.

1 Introduction

The phenomenal medical advances achieved in recent years, as well as the remarkable improvements in overall quality of life, have led to a significant increase in lifespan. This increased longevity is directly reflected in the worldwide emergence of a larger elderly and chronically ill population—a population in need of a special, sometimes round-the-clock, geriatric care. The growth of this type of patient population represents a whole new challenge to existing healthcare infrastructures worldwide. In order to deal with this challenge, countries around the world have experimented with a variety of approaches [Can08, US08, Nag06, AMD08]. One of the most promising among these approaches, is the adoption of telemonitoring. Telemonitoring is the medical practice of remotely monitoring the health of patients in the comfort of their homes, or more generally, outside traditional healthcare environments (e.g., hospitals, nursing homes, doctor's office). The idea of medical telemonitoring itself is not new [Nag06]. In 1906, Willem Einthoven, father of electrocardiography, sought to transmit electrodiagrams over telephone lines [Ein08]. In the 1920s, ship radios were used to link doctors with sailors to assist during medical emergencies at sea. In the 1970s, paramedics in remote Canadian villages were able to perform

life-saving interventions while linked with hospitals in distant towns via satellite. Today, telemedicine is beginning to mature considerably with new advances in communication technology. Although telemedicine is not suitable for all patients in all cases, it is still considered highly advantageous in many situations. For instance, medical telemonitoring has been successful in helping elderly people avoid nursing homes, maintain a more dignified social life, remain productive, stay home longer, and thus incur less healthcare costs. Medical telemonitoring has also helped decrease the burden on the country's healthcare infrastructure and economy as a whole. Besides the above aspects, medical studies have shown that telemonitoring makes patients more health-conscious as a result of being personally engaged in the health monitoring process. The study in [TSI⁺08] for example, found that patients became significantly more proactive and careful about their health, after being required to take their own blood pressure at regular time intervals.

Despite all the above advantages, patients are still showing a certain reluctance to accept the idea of medical telemonitoring. This lack of acceptance is generally attributed to two main concerns:

1. **Dependability.** How efficient is telemonitoring compared to an in-person visit to the hospital? Will someone be there to help in case of emergency?
2. **Privacy.** Is private data properly protected when sent over the wires to the hospital?

In this paper, we assume that sufficient resources are available to make the system dependable, and focus our attention on solving the privacy aspect of the problem. In particular, we present a medical telemonitoring protocol that preserves the patients' privacy.

Organization. We start in Section 2 by describing our system setting and model. In Section 3, we present the list of security and privacy requirements that our protocol achieves. We then give a high-level overview of our solution, along with a description of the building blocks, in Sections. 4 and 5. The details of the proposed protocol are highlighted in Section 6. In Section 7, we discuss the security and privacy features of the proposed protocol. In Section 8, we highlight related work before concluding in Section 9.

2 System Model and Settings

We assume a setting where two main parties are involved: a patient at home, and a health monitoring centre (*HMC*) located at the hospital. We assume that the patient has two types of devices:

- Measurement devices : they are used to measure the patient's vital signs such as his heart rate. Measurement devices can be portable wearable devices that patients carry on them all the time (e.g., a wearable heart-rate monitor in the form of a bracelet or a belt [Zep08]), or stationary static devices that the patient can use at discrete points in time (e.g., a conventional tensiometer)

- A master storage and communication device \mathcal{M} : this device, usually located at the patient's home, collects information from the measurement devices¹.

The collected information is stored and analyzed by the master device \mathcal{M} . Based on this analysis, \mathcal{M} sends a summary report about the patient's condition to the health monitoring centre periodically. However, in case the collected measurements indicate a sudden health deterioration, or a condition that requires immediate attention, the master device automatically triggers an emergency signal and sends an immediate notification to the health monitoring centre. In response to this notification, and following the specific medical practices in place, the health monitoring centre may send an ambulance to the patient's home, or notify the patient's neighbours and family members, etc. Figure 1 gives a summary of the overall setting.

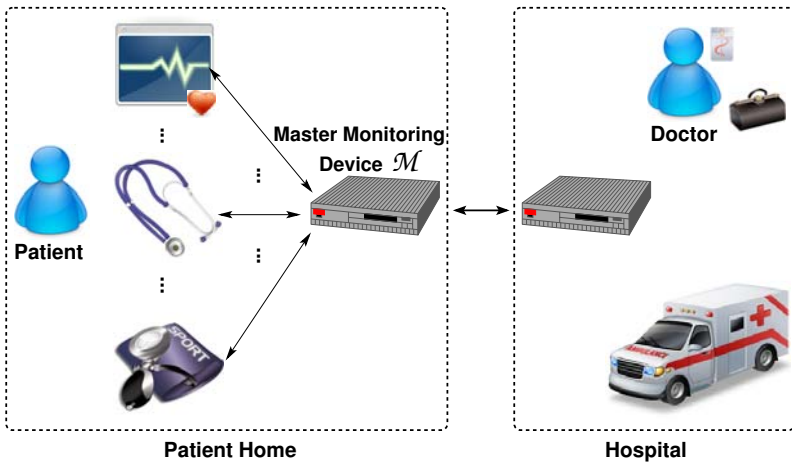


Fig. 1. Overview of the Health Telemonitoring System

From a design point of view, several issues need to be addressed in such a system. For example, the system has to be easy to use; the communications between the measurement devices and the master device on one hand, and between the master device and the health monitoring centre on the other hand, have to be seamless to the patient. The system has to be user-friendly; for example if the patient does not feel well he should be able to *easily* reach out for an emergency button to call for help. More importantly, the system has to be reliable, and the exchange of information should not compromise the patient's privacy.

In this paper, we focus on the security and privacy aspects of the system, and propose a construction that ensures a set of properties that we describe in the following section.

¹ For portable devices this transfer of information takes place once the patient is back at home.

3 Security and Privacy Requirements

- **Selective disclosure.** This property captures the ability of the patient to wilfully disclose *fine-grained* information about his identity attributes, and hide the rest. Selective disclosure implies another useful property called *minimal disclosure*, which describes a user’s ability to disclose the minimum information necessary for a transaction to take place. For example, in order to receive the services of a local hospital, the minimum information required could be a proof that the patient is a resident of a certain postal code area, without the need to reveal his exact address.
- **Patient-centricity.** A system is said to be *patient-centric* if and only if it guarantees that any data disclosed to the monitoring centre, must have received the prior approval of the patient.
- **Pseudonimity.** A system is said to preserve pseudonimity if data records sent from the patient’s home to the monitoring centre are linkable to each other (*e.g.*, via a patient’s pseudonym) but not to the patient’s real identity.
- **Conditional deanonymization.** In cases of emergency, it should be possible to recover the real identity behind the patient’s pseudonym, so that urgent help can be provided.
- **Integrity.** It should not be possible to alter health information on the way between the measurement sensors and the master device, or between the master device and the monitoring centre, without being detected.
- **Confidentiality.** The content of the data sent by the master device should be readable only by the monitoring centre, as intended.

4 Solution Outline

First, we assume that the patient has one or more measurement devices and one master monitoring device \mathcal{M} . The measurement devices communicate only with the master device \mathcal{M} , which in turn communicates with the health monitoring centre. The measurement and master monitoring devices are issued by the health monitoring centre to the patient. Both types of devices are assumed to be tamper-resistant. Each device is assigned some key material before being handed over to the patient. The key assignment process includes the embedding of some key material into a protected memory location on the device.

Without loss of generality and to keep the presentation simple, we assume in the remainder, that the patient has a single measurement device m , in addition to the master monitoring device \mathcal{M} . Let k_m be a symmetric encryption key assigned to the measurement device m , and shared with the master device \mathcal{M} . Let $(PK_{\mathcal{M}}, SK_{\mathcal{M}})$ be the public/private-key pair assigned to \mathcal{M} .

The measurement device m collects health readings from the patient and sends them encrypted under k_m to the master device \mathcal{M} . To ensure the integrity of messages between m and \mathcal{M} , the encrypted information can be signed using a message authentication code (MAC). Alternatively, conventional public key signatures (*e.g.*, RSA) can be used. However, the former option is the more efficient and preferred one.

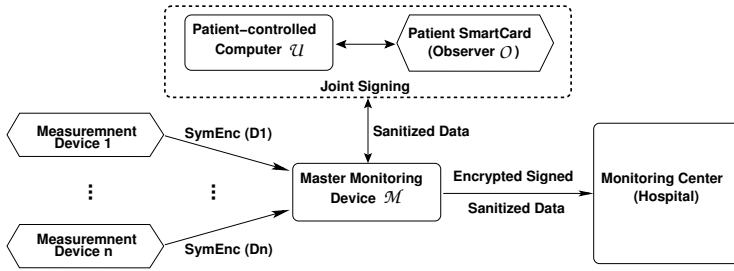


Fig. 2. Interaction between the Master Monitoring Device and the Hospital

The master device first analyzes the collected measurements, and removes any identifying information from them using data sanitization techniques (see Sec. 5.2.) The sanitized data is then signed by a computer under the patient’s control, before being sent encrypted to the hospital. The use of a patient-controlled computer, that is separate from the master device \mathcal{M} , to sign health data, is intended to strengthen *user-centricity*. This design choice also helps guarantee that data releases have been approved by the patient. User-centricity can be further enhanced by using a two-factor message authentication mechanism, where computing signatures requires both the knowledge of a secret key (stored on the patient’s computer) and the possession of a valid smartcard.

The signing procedure does not require the direct intervention of the patient, and can be configured to work automatically.

Privacy-preserving two-factor message authentication. To ensure patient-centricity as well as selective disclosure capabilities, we use the wallet-based version of Brands credential system (WBr) [Bra00, Chap 6]. The WBr system provides a way to authenticate communications between the master device and the health monitoring centre, while allowing the patient to selectively disclose his identity information [Bra00, Sec 6.3]. The WBr is a two-factor authentication system, where performing signatures on data, requires the patient to be in possession of a valid smartcard and to know the secret attributes underlying his credential. This feature makes attacks by impersonation against the patient harder, even when the patient’s computer is compromised.

In our context, the master device sanitizes the patient data, and then sends it to the patient-controlled computer for signing. The latter signs the data using the patient’s anonymous credential. The signed data is then sent over to the hospital, encrypted under the monitoring centre’s public key. Owing to the “selective disclosure” capabilities of anonymous credentials, the patient is able to wilfully disclose, via the computed signature, any self-approved information or property about his identity, while keeping everything else private.² The WBr system has a built-in mechanism to block any possible covert channels between the patient’s smartcard and the outside world (e.g., monitoring centre.) This

² *Selective disclosure* cannot be achieved using conventional X.509 certificates.

mechanism represents an additional guarantee that no data about the patient's identity will be sent without his approval. Figure 2 shows a high-level overview of the interaction between the master device and health monitoring centre.

5 Building Blocks

5.1 Brands Wallet-Based Anonymous Credentials

In [Bra00, Chap 6], Brands presents a credential system (WBr) suitable for the *wallet-with-observer* setting, where the user holds a *wallet* composed of a self-controlled computer denoted \mathcal{U} , and a tamper-proof device called *Observer* and denoted \mathcal{O} . The observer is supplied by a recognized certification authority CA. The WBr system is such that credentials are issued to the pair $(\mathcal{U}, \mathcal{O})$, and can only be shown if the two entities approve the showing, and perform it jointly. The WBr system provides a number of privacy preserving features including anonymity and selective disclosure.

Settings and Assumptions. The WBr system operates in the Discrete Logarithm setting, and its security is based on the DL assumption [Bra00, Chap 6]. Let p and q be two large primes such that $2q|(p-1)$. Let G_q be the unique subgroup of \mathbb{Z}_p^* of order q , and let g_0 be one of its generators. In the setup phase, the CA randomly chooses $y_1, \dots, y_\ell \in_R \mathbb{Z}_q$ and $x_0 \in_R \mathbb{Z}_q^*$, and computes $(g_1, \dots, g_\ell, h_0) := (g_0^{y_1}, \dots, g_0^{y_\ell}, g_0^{x_0}) \bmod p$. Next, the CA chooses a collision-resistant hash function $H : \{0, 1\}^* \rightarrow \mathbb{Z}_q^*$. Finally, the parameters $G_q, H, (g_0, g_1, \dots, g_\ell, h_0), (g_1^{x_0}, \dots, g_\ell^{x_0}, h_0^{x_0})$ are all made public.

Issuing Credentials to a Wallet-with-Observer. A credential consists of a public key h and a signature on it, $\sigma_{CA}(h)$, issued by a certification authority CA. Let $(h, \sigma_{CA}(h))$ denote a credential issued by the CA to a wallet $(\mathcal{U}, \mathcal{O})$. Let $x_1, x_2, x_3 \dots, x_{\ell'}, x_{\ell'+1}, \dots, x_\ell$ denote the attributes embedded in h . The attributes are such that x_1, x_2 are known only to the wallet observer \mathcal{O} , $x_3 \dots, x_{\ell'}$ and x_ℓ are known only to the user-controlled computer \mathcal{U} , while $x_{\ell'+1}, \dots, x_{\ell-1}$ are known both to the CA and \mathcal{U} . The fact that attributes x_1, x_2 are known only to the wallet observer \mathcal{O} , and that $x_3 \dots, x_{\ell'}$ and x_ℓ are known only to the user-controlled computer \mathcal{U} , means that the issued credential can be shown to a verifier only if \mathcal{O} and \mathcal{U} both approve the showing and participate in it. Figure 3 depicts the issuing process.

At the start of the issuing protocol, observer \mathcal{O} computes a message M_1 containing a number of commitments and a signed proof of knowledge. M_1 is sent to \mathcal{U} and then forwarded to the CA.

To provide guarantees that M_1 is originating from a legitimate observer, we assume that all legitimate observers have CA-supplied certificates embedded in them. These certificates are independent from the wallet holder's identity, and their embedding takes place before the wallets are attributed to a particular user. The certificates are of the form $(e := g_{\mathcal{O}}^x, cert_{CA}(e))$, where $g_{\mathcal{O}}$ is a public generator of G_q , and x a secret randomly chosen by the observer in question.

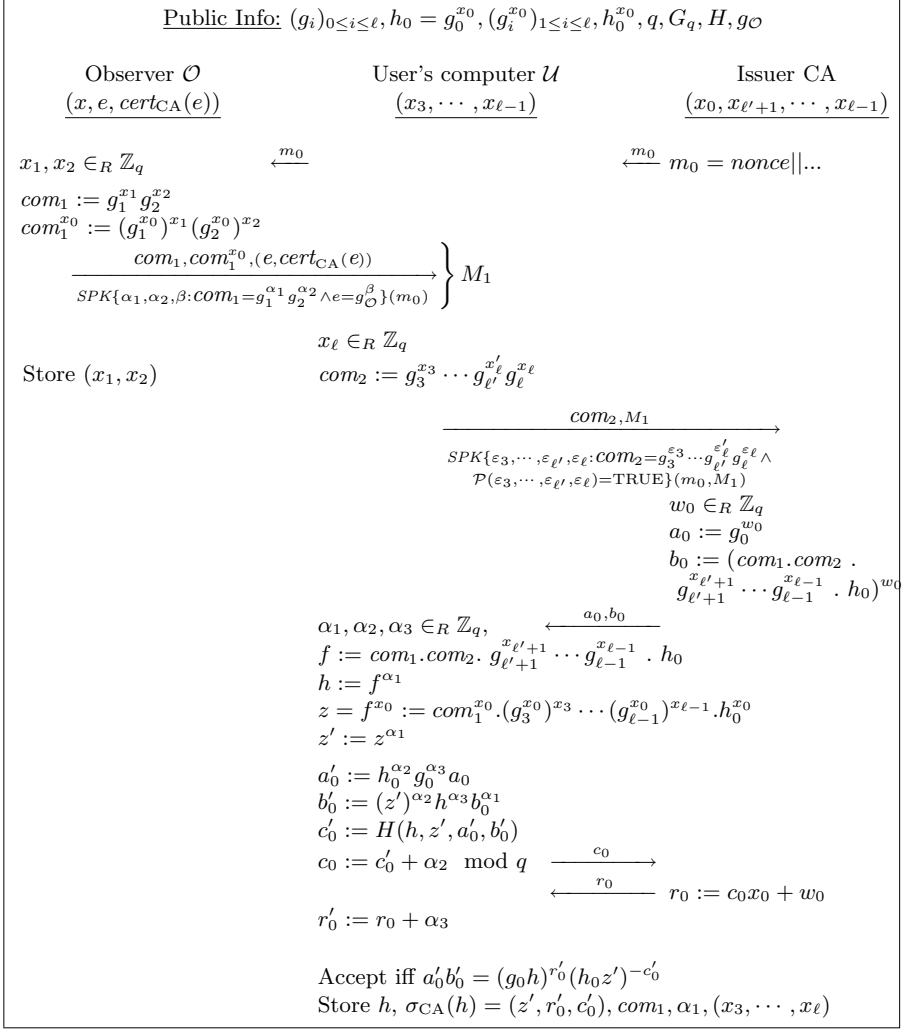


Fig. 3. Protocol for Issuing a Credential to a Wallet-with-Observer (NB. Attributes $x_{\ell'+1}, \dots, x_{\ell-1}$ are known to both \mathcal{U} and the CA. The secret key x , underlying the Observer's certificate $(e, cert_{\mathcal{CA}}(e))$, is such that $e := g_{\mathcal{O}}^x$).

A legitimate observer authenticates its messages, by including its certificate in M_1 , and proving knowledge of the underlying secret x .

It is worth noting that the issuing protocol shown in Figure 3 is a blind signature protocol, since the CA does not learn any information about the credential $(h, \sigma_{\mathcal{CA}}(h))$ obtained by the pair $(\mathcal{U}, \mathcal{O})$. More details on the *blinding* features of the protocol can be found in [Bra00, Chap 6].

It is also worth mentioning that in the basic WBr system, the observer holds only *one* secret x_1 , as opposed to two secrets (x_1, x_2) as in the protocol of Figure 3. We made the latter choice to protect the observer's secrets (x_1, x_2) *unconditionally*, even from the user-controlled computer \mathcal{U} . This is achieved because given the commitment $com_1 := g_1^{x_1} g_2^{x_2}$, there are q solutions $(X, Y) \in (\mathbb{Z}_q^*)^2$ to the equation $com_1 = g_1^X g_2^Y$. Because these solutions are all equiprobable, even an all-powerful \mathcal{U} cannot guess which one is encoded in com_1 , with a probability of success better than $1/q$. Therefore the observer's secrets are *unconditionally* protected. This technique is similar to those used by Okamoto in [Oka92].

Issuing protocol correctness (sketch). By examining the protocol, we can derive the following equalities:

$$\begin{aligned} - a'_0 &= g_0^{(\alpha_2 x_0 + \alpha_3 + w_0)}, b'_0 = (h_0 \prod_{i=1}^{\ell} g_i^{x_i})^{(\alpha_1 \alpha_2 x_0 + \alpha_1 \alpha_3 + \alpha_1 w_0)} \\ - r'_0 &= c'_0 x_0 + \alpha_2 x_0 + w_0 + \alpha_3 \\ - h &= (h_0 \prod_{i=1}^{\ell} g_i^{x_i})^{\alpha_1}, z' = (h_0 \prod_{i=1}^{\ell} g_i^{x_i})^{(x_0 \alpha_1)} \end{aligned}$$

It is then straightforward to check that the following two equalities hold.

$$a'_0 = g_0^{r'_0} h_0^{-c'_0}, \text{ and } b'_0 = h^{r'_0} (z')^{-c'_0}$$

As specified in Figure 3, combining the above equalities leads the user \mathcal{U} to accept the credential as valid.

Wallet-assisted Credential Showing with Inflow and Outflow Prevention. The user holds a credential, with two of its attributes known only to the wallet observer \mathcal{O} , and others known only to the user-controlled computer \mathcal{U} . As a result, the credential showing requires a special protocol where both \mathcal{U} and \mathcal{O} need to cooperate. This protocol is illustrated in Figure 4. To prevent the establishment of a covert channel between the wallet observer and the verifier, the protocol in Figure 4 has built-in mechanisms for inflow and outflow prevention.

The protocol in Figure 4 highlights a simple setting, where users are only required to prove knowledge of the attributes underlying their credentials. There could be scenarios however, where users are required to prove more elaborate predicates about their attributes. It is still possible to handle those scenarios, since the WBr system [Bra00, Sec. 3.6] allows for proving a wide class of predicates on the attributes. Because of space limitations, we do not discuss general predicate proofs³ here, and leave them to the full version of the paper.

Correctness of the showing protocol (sketch). To establish the correctness of the showing protocol, we prove that the following equality holds:

³ For a small example, let us consider a credential h with attributes x_i , $1 \leq i \leq \ell$, and the predicate $\mathcal{P}: 3x_1 - x_2 = 0$. As shown in Fig. 4, h can be written as $h = h_0^{s_0} \prod_{i=1}^{\ell} g_i^{s_i}$, for $s_i = \alpha_1 x_i$, $1 \leq i \leq \ell$, and $s_0 = \alpha_1$. If we take into account predicate \mathcal{P} , then we get $h = h_0^{s_0} (g_1 g_2^3)^{s_1} \prod_{i=3}^{\ell} g_i^{s_i}$. To prove that the x_i 's satisfy \mathcal{P} , we just need to prove *knowledge* of a discrete log representation of h wrt. basis $(h_0, g_1 g_2^3, g_3, \dots, g_{\ell})$. This can be done by using the same technique of Fig. 4.

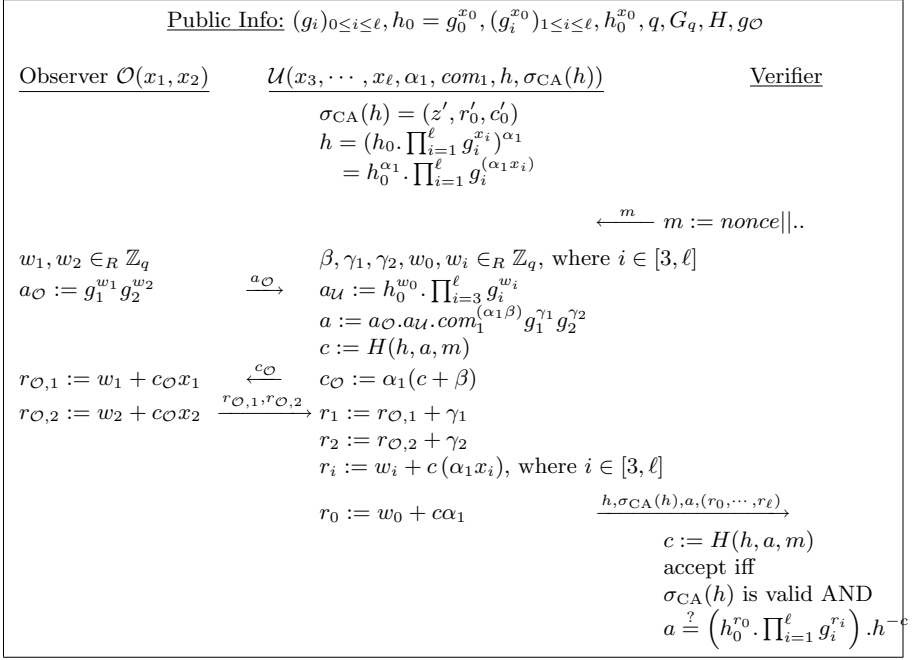


Fig. 4. Wallet-assisted Credential Showing with Inflow and Outflow Prevention (The output tuple (a, r_0, \dots, r_ℓ) represents a signature on the Verifier's message m , using the pair $(\mathcal{O}, \mathcal{U})$'s credential $(h, \sigma_{CA}(h))$. The protocol can also be seen as a signed proof of knowledge of a discrete log representation of h with respect to basis $(h_0, g_1, \dots, g_\ell)$).

$$a = \left(h_0^{r_0} \cdot \prod_{i=1}^{\ell} g_i^{r_i} \right) \cdot h^{-c}$$

Starting with the right hand side of the equation we obtain:

$$\begin{aligned}
 \left(\prod_{i=1}^{\ell} g_i^{r_i} \right) h_0^{r_0} h^{-c} &= \left(\prod_{i=3}^{\ell} g_i^{w_i + c \alpha_1 x_i} \right) \cdot g_1^{w_1 + \gamma_1 + (c + \beta) \alpha_1 x_1} \cdot g_2^{w_2 + \gamma_2 + (c + \beta) \alpha_1 x_2} \\
 &\quad \cdot h_0^{w_0 + c \alpha_1} \cdot \left(h_0 \prod_{i=1}^{\ell} g_i^{x_i} \right)^{-\alpha_1 c} \\
 &= \left(\prod_{i=3}^{\ell} g_i^{w_i} \right) \cdot g_1^{w_1 + \gamma_1 + \beta \alpha_1 x_1} \cdot g_2^{w_2 + \gamma_2 + \beta \alpha_1 x_2} \cdot h_0^{w_0} \\
 &= a_{\mathcal{O}} \cdot a_{\mathcal{U}} \cdot (g_1^{\gamma_1} g_2^{\gamma_2}) \cdot (g_1^{x_1} g_2^{x_2})^{\alpha_1 \beta} \\
 &= a_{\mathcal{O}} \cdot a_{\mathcal{U}} \cdot (g_1^{\gamma_1} g_2^{\gamma_2}) \cdot com_1^{\alpha_1 \beta} \\
 &= a
 \end{aligned}$$

Arguments on Inflow and Outflow Prevention capabilities. Briefly stated, the random number β chosen by the user-controlled computer \mathcal{U} , is used to mask the challenge $c = H(h, a, m)$, thereby preventing the verifier from covertly communicating with the observer \mathcal{O} through m . Similarly, random numbers γ_1 and γ_2 prevent the observer \mathcal{O} from sending out covert information to the verifier.

5.2 Data Sanitization

In addition to controlling the disclosure of identity information (which we handle by using anonymous credentials), the patient needs a *data sanitization* procedure to anonymize the health records he sends to the monitoring center. As described in Sec. 2, measurements about the patient’s health are collected and aggregated in the form of health records. The patient then releases a portion of these records in *sanitized* form to the telemonitoring service. Data sanitization (e.g., [Swe02, CKV⁺03]) which aims at anonymizing the records, consists of either removing fields that contain identifying information, or modifying the values of those fields through *generalization* or the *addition of some random noise*.⁴ The goal of the sanitization procedure, is to blur the direct link between a given record and its owner. Instead of being mapped back to a single individual, a *sanitized record* can generally be associated with a set of possible owners. The latter is called the *anonymity set*, and the actual owner of the sanitized record is said to be *anonymous within this set*. Data sanitization aims at making the anonymity set as large as possible, while keeping the sanitized data useful for subsequent (e.g., statistical) analysis.

Note. It is worth mentioning here, that for the data sanitization techniques above to work, the patient’s master device \mathcal{M} should have an *a priori* knowledge of the overall distribution of health measurements among the population, as well as an approximation of the size of the total dataset. Such information need to be updated regularly, and can be made available by the health authorities managing the hospital, for example.

Note that the data exchange in our setting is pseudonymous, since the public part of the patient’s credential’s ($h, \sigma_{CA}(h)$) is revealed for every data submission. The latter serves as a pseudonym. The pseudonymity of data is important here; it makes it possible for the monitoring centre to aggregate health records into medical histories associated with each patient’s pseudonym. Because of space limitation, we do not discuss the details of data sanitization any further, and leave this topic to the full version of the paper.

⁴ The practice of adding noise from a known probability distribution, is one way to anonymize records while keeping the data useful from a statistical point of view. Anonymization through *generalization* consists of replacing the value of a certain field by a more general representation. For example, a specific value in the “age field” of a record can be replaced by an age interval.

6 Proposed Protocol

In the following, we describe the steps taken by each party involved in the telemonitoring process. Figure 5 depicts the whole process. Let \mathcal{U} and \mathcal{O} denote the patient’s computer and smartcard respectively.

1. Initially, at the setup phase, the hospital (or monitoring centre) authority issues a credential to the pair $(\mathcal{U}, \mathcal{O})$. This is done using the issuing protocol described in Figure 3. At the end of this step, the pair $(\mathcal{U}, \mathcal{O})$ obtain a credential $(h, \sigma_{\text{CA}}(h))$ containing a number of attributes x_1, \dots, x_ℓ , such that x_1, x_2 are only known to the smartcard \mathcal{O} , while x_3, \dots, x_ℓ are known only to the patient—represented by \mathcal{U} —and not to \mathcal{O} .
2. After gathering health readings from the measurement devices, the master monitoring device (also located at the patient’s home) sanitizes the data by removing identifying information.⁵ The master device then sends the sanitized data to the patient’s computer to be approved and signed.
3. The patient’s computer \mathcal{U} checks the received data to make sure it has been properly anonymized. If this is the case, \mathcal{U} initiates the signature process with the patient’s smartcard \mathcal{O} . The signature process is in fact a signed proof of knowledge, of the attributes underlying credential $(h, \sigma_{\text{CA}}(h))$ initially obtained in step 1. The signed proof is performed on the sanitized data. The computation of the signed proof, requires the collaboration of both the smartcard (which knows x_1, x_2) and the patient’s computer (which knows x_3, \dots, x_ℓ). The necessity of this cooperation makes the signing a two-factor⁶ authentication process. The signing is done using the protocol described in Figure 4, where m is chosen to be a concatenation of (1) a *nonce* chosen by \mathcal{M} and (2) the sanitized data to be signed (sEHR). That is $m := \text{nonce}||\text{sEHR}$.
4. Once it receives the signed proof $SPK_{\text{patient}}(\text{sEHR})$, the master device \mathcal{M} checks its correctness, and signs on top of it. \mathcal{M} then sends $\sigma_{\mathcal{M}}(SPK)$ and sEHR to the monitoring centre at the hospital, all encrypted under the latter’s public key. In order to prevent the monitoring centre from identifying the patient via the signature $\sigma_{\mathcal{M}}(SPK)$, we can set device \mathcal{M} to use a group signature scheme (e.g., [CvH91, ACJT00]) to compute $\sigma_{\mathcal{M}}(SPK)$, instead of a conventional public key signature.

In the setting above, the patient has a single credential $(h, \sigma_{\text{CA}}(h))$ which he uses to sign all the data disclosed to the hospital. This credential can be viewed as a pseudonym of the patient, and used to link all of his health records. This linkability is essential to the process of building medical dossiers. In cases where linkability (between records of the same patient) is not needed, the patient may use a multi-show credential such as those in [CL02, CL04]. Such a credential can be shown an indefinite number of times without the showings being linkable to each other or to the identity of the credential holder.

⁵ The data can be sanitized according to rules chosen by the patient and his doctor.

⁶ In order to compute a signature with respect to credential $(h, \sigma_{\text{CA}}(h))$, one needs to know the secrets x_3, \dots, x_ℓ and to hold a smartcard containing x_1, x_2 .

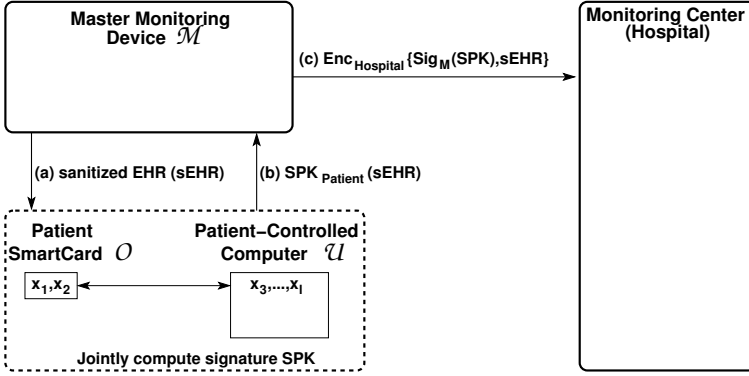


Fig. 5. High-level Protocol Architecture (with two-factor message authentication)

7 Security and Privacy Analysis

In the following, we briefly discuss the security and privacy of our protocol. Our analysis assumes that all the underlying building blocks are secure. A more complete analysis will be given in the full version of the paper.

- **Selective disclosure.** This follows from the *selective disclosure* capabilities of the WBr credential scheme, and the assumption that the tamper-proof device \mathcal{M} will behave according to the protocol specifications.
- **Patient-centricity.** This is achieved due to the fact that it is infeasible to compute a signed proof of knowledge, with respect to a credential, on behalf of its owner. In particular, the presence of a valid signed proof of knowledge (*SPK*) in the data received by the monitoring centre, indicates that the patient has approved the data release. The data could not have been sent without the patient’s approval since (1) no party, except the patient, is able to compute *SPK*, and (2) \mathcal{M} is trusted to follow the protocol.
- **Pseudonymity and Conditional deanonymization.** This is achieved by the combination of three mechanisms: (1) limiting the disclosure of health data to the sanitized form only, (2) the use, by the patient, of *anonymous credentials* to sign the sanitized data, and (3) the use of a *group signature* scheme by \mathcal{M} . The group signature computed by \mathcal{M} convinces the monitoring centre at the hospital, that the signature is generated by a valid master device, without revealing which one. This prevents the monitoring centre from identifying the patient through \mathcal{M} . Note however that using the deanonymization mechanism of the group signature scheme, it is possible in case of emergency, to recover the identity of \mathcal{M} , and consequently that of the patient.
- **Defense against covert channels.** In the proposed protocol, we use the smartcard-based version of Brands credentials (Fig 4) which has mechanisms to prevent covert inflow and outflow of information between the patient’s

- smartcard and the monitoring centre at the hospital. This prevents the patient's smartcard from sending out information not approved by the patient.
- **Integrity.** This breaks down into two sub-properties: (1) integrity wrt. alterations caused by the patient, and (2) integrity wrt. alterations occurring during the wire transmission. The first sub-property is satisfied following the soundness⁷ of the signed proofs of knowledge, and the assumption that \mathcal{M} will follow the protocol specifications. That is, the fact that \mathcal{M} added its signature on the patient's *SPK*, implies that \mathcal{M} accepted the patient's signature as valid, and found no alteration to the data. The second sub-property is ensured by the unforgeability and soundness of both signature schemes: the group signature and the credential-based signature.
 - **Confidentiality.** Assuming that \mathcal{M} is tamper-proof and that it follows the protocol in section 6, the confidentiality property is insured by the fact that all data sent to the hospital is encrypted using a secure encryption scheme.

8 Related Work

The topic of telemedicine has received a lot attention from researchers both in the industry and academia [DMDB07, DDB08, SK08, RLNB03, INT08, AMD08]. Most of the previous works in this area however seem to concentrate mainly on the usability, interoperability, and communication efficiency aspects. The question of protecting the patients' privacy has also been addressed (e.g., [DMDB07, SK08]), but to a lesser extent. In the following we highlight some works from the literature that are most relevant to this paper.

In [DMDB07, DDB08], Durresi et al. propose a ubiquitous health monitoring system which uses existing wireless telecommunication networks to carry data. Owing to the pervasiveness of wireless networks, the system put forward by Durresi et al. makes it possible to send patient data continuously and in real time, to a central medical database. The protocols in [DMDB07, DDB08] address a number of security-related problems such as authentication, confidentiality, and non-repudiation, by taking advantage of the authentication, public-key encryption and signing mechanisms, which come by default with the communication infrastructure (e.g., GSM network.) With respect to privacy, the protocols in [DMDB07, DDB08] achieve only pseudonymity (towards a network observer.) This is done by requiring the patient's cellphone to assume a temporary identity assigned to it by the monitoring centre.

As widely recognized in the literature (e.g., [Bra00, CL02]), pseudonymity *alone* is not sufficient to protect the privacy of patients. In fact, patients' privacy can be guaranteed only if the patient is able to decide (1) what information about his identity is disclosed and to whom, and (2) which portions of his health data can be sent to the monitoring centre. The protocol we propose in this paper satisfies these requirements through the use of *data sanitization* and privacy-preserving credentials, which allow patients to selectively disclose their identity information.

⁷ By soundness, we mean that a signed proof that is incorrect, will be detected with overwhelming probability.

In [SK08] Sufi and Khalil propose an efficient data sanitization method for electrocardiograms (ECG). Their method, which includes encoding and compression algorithms, is designed to conceal cardiovascular details, as well as features in ECG data that could identify an individual. The method in [SK08] achieves compression ratios neighbouring 20:1, which makes it suitable for real-time telemonitoring. Our system can use their method as a building block.

Other systems such as [RLNB03, AMD08, INT08, Can08] have dedicated most of their focus to important aspects such as deployability and interoperability, without giving much attention to the privacy issues surrounding telemonitoring systems. This paper addresses medical telemonitoring in a global way, and proposes concrete solutions to protect the privacy of patients.

9 Conclusion

We have presented a protocol for monitoring patients in the comfort of their homes. Our proposed protocol protects patients' privacy at two levels: (1) the *identity information level*: patients are able to selectively reveal information about their identity and to hide the rest, and (2) the *medical data level*: health measurements collected from the patient are sanitized according to patient-approved privacy policies before being sent to the health monitoring centre (*HMC*). The data is sanitized in a way that keeps it useful from a medical perspective, while preventing it from being directly linkable to the patient's identity.

The protocol we propose, provides security against impersonation attacks, even when the patient's computer is compromised. This is achieved thanks to a smartcard-based two-factor authentication mechanism. The same authentication mechanism allows the monitoring center to recognize and accept data records that have been approved for release by the patient, and decline others. Furthermore, and in line with our stated concerns for usability, the procedure we propose for handling disclosure approvals is automated; patients are only required to specify their disclosure policies in an initial configuration phase.

In addition to our ongoing prototype implementation, this work can be extended in a number of ways. For example, in cases where privacy requirements are less stringent, one could use simpler two-factor authentication methods, in particular those based on one-time passwords, generated by a tamper-proof device. We can also improve the communication efficiency of our protocol by using sanitization and compression techniques such as those in [SK08].

The issue of *liability* also deserves further investigation. For example, if a patient dies, the monitoring center should be able to prove that everything that could be done to save the patient has been done. A simple way to achieve liability would be to require the *HMC* to send an acknowledgement token back to the patient's master device, every time it receives data from the patient. These tokens can be used later to prove that the *HMC* was aware of the patient's condition. In addition, the *HMC* should keep a record of all the efforts it made to help the patients (e.g., ambulance calls etc.) All of these records, as well as the

acknowledgement tokens, can constitute the basis of any service audits that may follow. More details on the question of liability will be presented in future work.

Acknowledgements. This research is partially funded by the Interuniversity Attraction Poles Programme Belgian State, Belgian Science Policy and the Research Fund K.U.Leuven and by the IWT-SBO project (ADAPID) “Advanced Applications for Electronic Identity Cards in Flanders”.

References

- [ACJT00] Ateniese, G., Camenisch, J., Joye, M., Tsudik, G.: A practical and provably secure coalition-resistant group signature scheme. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 255–270. Springer, Heidelberg (2000)
- [AMD08] AMD Telemedicine Inc. Telemedicine becoming more relevant to childrens health within school systems (December 2008)
- [Bra00] Brands, S.: Rethinking Public Key Infrastructures and Digital Certificates: Building in Privacy. MIT Press, Cambridge (2000)
- [Can08] Canada Health Infoway. Patients manage health at home with telehealth (November 2008), <http://www.infoway-inforoute.ca>
- [CKV⁺03] Clifton, C., Kantarcioglu, M., Vaidya, J., Lin, X., Zhu, M.: Tools for privacy preserving distributed data mining. ACM SIGKDD Explorations 4(2) (2003)
- [CL02] Camenisch, J., Lysyanskaya, A.: An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, pp. 93–118. Springer, Heidelberg (2001)
- [CL04] Camenisch, J., Lysyanskaya, A.: Signature schemes and anonymous credentials from bilinear maps. In: Franklin, M. (ed.) CRYPTO 2004. LNCS, vol. 3152, pp. 56–72. Springer, Heidelberg (2004)
- [CvH91] Chaum, D., van Heyst, E.: Group signatures. In: Davies, D.W. (ed.) EUROCRYPT 1991. LNCS, vol. 547, pp. 257–265. Springer, Heidelberg (1991)
- [DDB08] Durresi, A., Durresi, M., Barolli, L.: Secure ubiquitous health monitoring system. In: Takizawa, M., Barolli, L., Enokido, T. (eds.) NBiS 2008. LNCS, vol. 5186, pp. 273–282. Springer, Heidelberg (2008)
- [DMDB07] Durresi, A., Merkoci, A., Durresi, M., Barolli, L.: Integrated biomedical system for ubiquitous health monitoring. In: Enokido, T., Barolli, L., Takizawa, M. (eds.) NBiS 2007. LNCS, vol. 4658, pp. 397–405. Springer, Heidelberg (2007)
- [Ein08] Einthoven, W.: (2008), <http://nobelprize.org>
- [INT08] Intel Digital Health Group Inc. Addressing the challenges of chronic illness with personal health system technology (2008)
- [Nag06] Nagy, B.: Telemedicine’s depth now going beyond rural areas. Managed Healthcare Executive (2006), <http://www.managedhealthcareexecutive.com>
- [Oka92] Okamoto, T.: Provably secure and practical identification schemes and corresponding signature schemes. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 31–53. Springer, Heidelberg (1993)

- [RLNB03] Rialle, V., Lamy, J.-B., Noury, N., Bajolle, L.: Telemonitoring of patients at home: a software agent approach. *Computer Methods and Programs in Biomedicine* 72(3), 257–268 (2003)
- [SK08] Sufi, F., Khalil, I.: Enforcing secured ecg transmission for realtime telemonitoring: A joint encoding, compression, encryption mechanism. *Security and Communication Networks* 1(5), 389–405 (2008)
- [Swe02] Sweeney, L.: Achieving k-anonymity privacy protection using generalization and suppression. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.* 10(5), 571–588 (2002)
- [TSI⁺08] Trief, P.M., Sandberg, J., Izquierdo, R., Morin, P.C., Shea, S., Brittain, R., Feldhousen, E.B., Weinstock, R.S.: Diabetes management assisted by telemedicine: Patient perspectives. *Telemedicine and e-Health* 14(7), 647–655 (2008)
- [US08] The assisted-living project (November 2008), <http://lion.cs.uiuc.edu/assistedliving>
- [Zep08] ZephyrTM: Smart Fabric for physiological and bio mechanical monitoring (2008), <http://www.zephyrtech.co.nz/technology>