How to Collect Consent for an Anonymous Medical Database

Emmanuel Benoist and Jan Sliwa
RISIS, Bern University of Applied Sciences - TI, Quellgasse 21, CH-2501 Biel, Switzerland

Keywords: Medical Databases, Anonymous Consent, Privacy, PKI Infrastructure.

Abstract: The goal of some medical databases is not to support the actual treatment of individual patients, but to provide the platform for medical research. Health data collected in such databases have to be anynomized - they should be analyzed only statistically and should not permit to retrieve the patient's identity. Medical data collected for research should be anonymized to protect the patients' privacy. In many countries it is mandatory. In many cases, not only one person treats a patient for a given illness. The documentation of a case requires the collaboration of different physicians that share information. This sharing of information requires the patient to authorise the access to the data stored by one physician by another one. We need therefore to implement a system for collecting the consent of an anonymous person. We present a novel solution to allow the practitioner to collect the consent of the patient in order to access the data recorded for that person. This solution is based on existing infrastructure, such as X509 certificates (present in e-ID or e-Health cards). Patients do not require to acquire any new hardware or to remember any new secret. We produce the fingerprint of the private key of the patient that can be used to re-identify the patient without having to know the identity of the patient (for instance the certificate) or even the patient's public key.

1 INTRODUCTION

The goal of this paper is to present a practical solution for efficient, secure and privacy preserving sharing of anonymous medical information stored in registry used for medical evaluative research. The aspects of constructing and managing registries are presented in (Gliklich and Dreyer, 2010).

Important work is currently being done in the area of the Electronic Health Records, where information is combined from many distributed and autonomous sources. Often such information is heterogeneous, in various formats, including unstructured notes. Many networks assuming mutual trust and permitting the exchange of medical data are currently in use or are being implemented. We can mention here e-toile[1] (Geneva, Switzerland), Clalit Health Services[2] (Israel) or GCS EMOSIST-FC[3] (Franche-Comté, France). Such decentralized systems do not create a new central database, but rather let data be stored locally where they have been produced and provide the methods for remote access. Data are used for treatment only and cannot be used for statistics, since they are mainly heterogeneous (each data provider having its own format).

The setting we examine in this paper is different. We consider the case of a database (registry) used to collect data for medical evaluative research. On the contrary to Electronic Health Records, a medical registry contains a limited set of data, but coherent for all patients, as its main goal is to allow to perform meaningful statistics. Medical data are collected in a centralized database where they can be compared and analyzed. Unlike for the EHR, data in a registry can be anonymous, since they are not used for treatment of the patient but only for statistical purposes. As personal information is necessary to retrieve a patient in order to add supplementary information (e.g. a follow-up record), it is also stored, but separately, so that connecting medical cases with personal data is impossible. We will describe the architecture used to handle both parts of data in a privacy protecting manner.

The physician who treats the patient can access all data records he/she has created. The same patient may be also treated by another physician who has no direct access to the records created by other doctors. The physician may however need this information in order to apply correct treatment or to further document the

[1] www.e-toile-ge.ch/etoile.html
[2] www.clalit-global.co.il/en/
[3] www.ch-dole.fr/contenu.php?idR=1

Benoist E. and Sliwa J..
How to Collect Consent for an Anonymous Medical Database.
DOI: 10.5220/0004902404050412
In *Proceedings of the International Conference on Health Informatics* (HEALTHINF-2014), pages 405-412
ISBN: 978-989-758-010-9

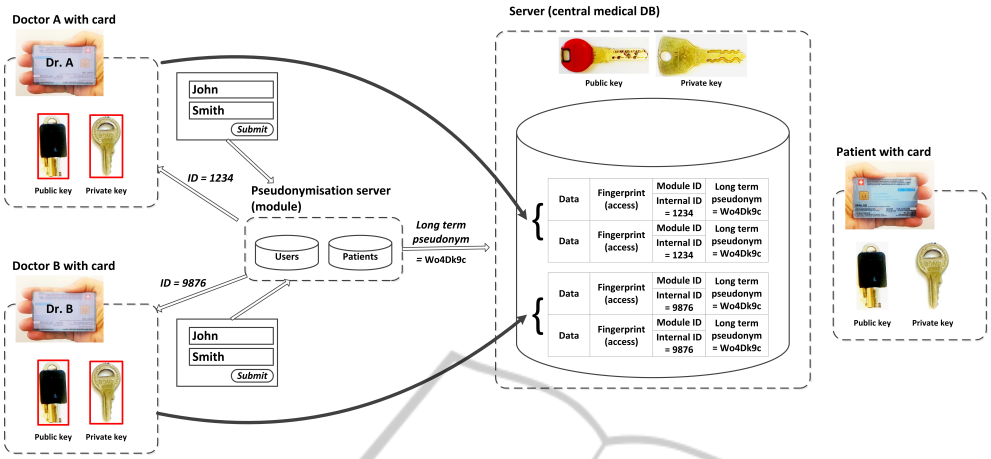

Figure 1: Configuration.

case. As we assume that this happens in the presence of the patient during a consultation, the patient may express the consent to make it available to the doctor. The item used to control the access is the patient's health related smartcard. The data records concerning a given patient are stored in the database along with the *"fingerprint"* generated from the data on his/her smartcard. This fingerprint marks the data ownership and protects them from an unauthorized access. The same patient's card may be used to unblock the access. The patient expresses the consent to use his/her data by allowing the physician to use his/her smartcard. No other items are necessary. This solution is practical and provides simultaneously an adequate security level. This protocol is the main subject of this paper, along with the presentation of the necessary environment and data structures.

As the smartcards play the essential role in our protocol, we assume their widespread use by the patients and by the health professionals. Currently, in many countries such cards are being deployed: European Health Insurance Card[4] with the related project NETC@RDS[5], Carte Vitale[6] in France, Versichertenkarte[7] in Switzerland, and many others.

The advantages of the proposed scheme are:

- the patient data are stored anonymously and the patient nevertheless retains control over it
- the patient needs neither to acquire any new token nor to remember any new secret

We first show the basic scenario that we want to handle. We then present the published work related to the considered case and show what distinguishes our approach from theirs. We discuss the various risks faced by such an application. Then, we will expose the details of the proposed protocol and comment on the advantages and the problems of the scheme. In the following, we discuss the related problems not treated in this paper and finally suggest the directions for the future research.

2 BASIC SCENARIO

The environment we consider in this paper is a medical database (Fig. 1). Its goal is evaluative research, i.e. assessing the efficiency of various therapy methods and devices. The database is centralized and contains homogenous data, because only in this way valid statistical comparisons can be performed. Medical data are anonymous, identity of the patients is irrelevant. As the patient's case consists of a sequence of events scattered over many years, they have to be internally connected in the database. If the patient appears for a subsequent consultation, his/her case has to be retrieved based on the his/her identity. Therefore a pseudonymisation server (later called *module*) is used. On this server the patients as well as the physicians are registered. The same patient may be registered by many physicians, every relation physician-patient receives a distinct ID. This permits to delimit the data directly accessible to a specific physician. On the other hand, the patient him-/herself has a unique long term pseudonym that permits to connect anonymously all events in his/her medical history for statistical analysis. When a physician logged into the system requires to access a patient with use of his/her personal data (name, etc.), the pseudonymi-

[4] ec.europa.eu/social/main.jsp?catId=559&langId=en
[5] www.netcards-project.com/web/frontpage
[6] www.sesam-vitale.fr/index.asp
[7] www.bag.admin.ch/themen/krankenversicherung/07060/

sation server responds with tho numbers. One is the internal ID that denotes the items in the central medical database created by this physician to which he/she has a direct access. The other one is the patient's long term pseudonym that is not seen by the physician but can be later used to access other items created by other physicians. This access depends on the patient's consent, as implemented in the protocol presented in this paper.

In practice, the need to treat the same patient may arise when he/she moves to another city or needs consultations at many physicians: general practitioner (family doctor) and various specialists. The patient may be treated by an oncologist and the surgeon, his/her tissues may be analyzed in a laboratory. At a certain moment, a doctor has to collect all this information in order to assess the case and to define further therapy. The patient, trusting the doctor, can allow him/her to access the relevant data created by all specialists.

A registry (with the corresponding pseudonymisation server - module) is organized around a specific medical problem, like cancer, orthopedical prostheses, or else. The central database can support many distinct modules. We intend not to connect data from different modules. It could be of value for the medical research, but would pose new security and privacy related problems.

Each module uses a specific function to construct the long term pseudonym. The principle is always the same: the module combines a selection of partial identifiers, adds salt (specific for this module) and computes a hash on this information. The selection of the partial identifiers and the way they are combined may vary from module to module. For instance, for one module the social security number birth year and gender can be used. For another module we can use last name at birth, first name at birth, birth date, city and country of birth. Each time a patient visits the same module, the same pseudonym will be computed. In some very rare cases two patients will receive the same pseudonym, but only if they have exactly the same identifiers.

Using (secret) salt prohibits the administrators of the central database to discover the patients' real identity.

Until now, we have discussed the long term pseudonym that is used to connect anonymously the patient's data records. Another item, stored along with the medical data and controlling the access to them is the fingerprint. The fingerprint is created using the data stored on the patient's health smartcard. If the patient presents the same smartcard to the physician, the same fingerprint can be generated what is understood as the consent to access the data given by the patient present in persona. In following sections the handling of the fingerprint will be presented.

3 RELATED WORK

As medical information is more and more stored electronically, studying various aspects of this process has created a vast research area. Our interest is directed towards privacy and security of medical data, especially in the interaction of the systems used for health support (Hospital Information Systems, Electronic Health Records) and the systems used for medical research (clinical trials repositories, medical registries).

We will present here recent literature regarding this subject.

In our article, we do not enter in detail into the way the pseudonym is created to link all the cases of the same patient in the database. This subject has already been covered by other publications. For instance (Elger et al., 2010) present the aspects of the reuse of health data for clinical research, especially the anonymization of data and the construction of pseudonyms. (Wilson, 2005) addresses also the problem of pseudonymization, suggesting the use of the PKI smartcards. We opted for the computation of a hash based on information that remain stable (identity at birth for instance) and a salt specific for each study. Our system is much simpler and does not require any new token or information from the patient.

It has to be stressed that retrieving and connecting data is a different problem from controlling the access to them. Data items have to be marked with the long term pseudonyms if they have to be treated as a set, what is a necessity in the case of medical research. Consent validation is another problem and we use different means (a fingerprint based on the smartcard) to achieve this goal.

(Kwon, 2011) proposes to use X.509 certificates to provide anonymized session identifiers that could be deanonymized under certain circumstances. This is very near to what we aim, but we concentrate here on the consent. The anonymity or the way how the pseudonym is created is out of the scope of our article.

Another related problem is the one stressed by (Camenisch and Lysyanskaya, 2001). They focus on anonymous credentials delivered by a central authority to anonymous users. We try to solve the inverse problem, where an anonymous user (for us a patient) gives a credential to a central system, while remaining anonymous.

4 RISK ASSESSMENT

In this section we will present the different risks for the application and how we mitigate them.

The first risk to consider is the attack from an outsider. The protection is done according to OWASP guide lines and protects mainly against the OWASP Top 10 flaws[8]. The details of the protection are not included in the scope of this article.

An outsider cannot become a legal user of our system, since the module administrators verify the identity of the users registered in their module. Since the modules are in most of the cases operated by medical societies, the user is typically a member of such a society. This fact should also limit the motivation of the users (physicians) for misusing their access rights and disclosing confidential data. The price for misbehaving would be a rejection from the community and an end of the professional career. It does not make a data theft impossible but raises substantially the bar for it. In any case, they cannot browse freely in the database - they have only access to the data they entered themselves or to which they have obtained explicit consent from the patient.

Administrators of a module do just have access to the data of the module (i.e. the identity of the patients). This information is important and can already be stigmatizing, like being registered in a HIV database. However, the module administrators do not have access to the central database, so they do not know the medical details of the case.

Administrators of the central database have only access to anonymized data. They cannot infer the identity of the patients from the hash they have. Even a dictionary attack is not possible, since they do not have access to the salt used in the module for computing the hash.

We propose a system that allows to collect consent of a patient to share his/her medical data between different physicians without disclosing the identity of the patient to the system.

The system must rely on a preexisting public key infrastructure. We can not access the certificates of the patients in this PKI, since the certificates contain the identity of the patients. We will therefore produce a *fingerprint* of this certificate, that does not reveal it. The consent to access is considered as given, if the patient's certificate produces the same fingerprint as was previously stored with the data. In this process, the identity of the patient remains unknown for the server.

In a public key infrastructure, the changing and revocation of keys is always of crucial importance. (Ferguson et al., 2010) is here a good reference and discusses also other practical aspects of key management. Our system does not have the possibility to handle revocation lists or expiration dates. We propose therefore a way to update the fingerprint of the key when the key is changed, e.g. when the card is lost or renewed. In this process, the physician takes the role of the trustee. Even if our system does not trust physicians in general (they should only access to their own data), we will rely on them for renewing the fingerprint of the key of their patients. Since this step is central in our system, we will require the physician to sign any modification in the fingerprint with his/her health professional's card. This will allow the administrators to monitor any misuse of the system and to react accordingly to protect the data they have in custody.

[8]www.owasp.org

5 PROPOSED PROTOCOL

The protocol is separated in two parts. The first part concerns the visit at the physician when the medical data are collected and stored on the server. The second part refers to another visit when the patient allows another health practitioner to access the previously stored data.

As the medical data are stored, a "fingerprint" controlling the access is created and stored along with the data. This fingerprint is a shared secret, based on the pair of the private keys (of the server and of the patient) in the Public Key Infrastructure (PKI) scheme. It does not require the server to know the patient's certificate, and not even to know the patient's public key. It can be later activated by the security keys stored on the patient's smartcard. It is stored on the server and the patient just uses his standard card and needs not to remember or store any new information. The doctor (or another health professional) plays also an important role in the process, for example verifying the identity of a physically present patient. Therefore he/she is included in our scheme, together with his/her Health Professional's Card that can be used to sign and certify his/her actions.

In the remainder of this article, we make no deeper analysis of the communication between the physician and the client on one side and the module on the other side. The creation of the internal ID and the long term pseudonym is also out of the scope of this article.

5.1 Enrollment (Fig. 2)

The patient visits a physician and the physician generates a record that must be inserted in the central

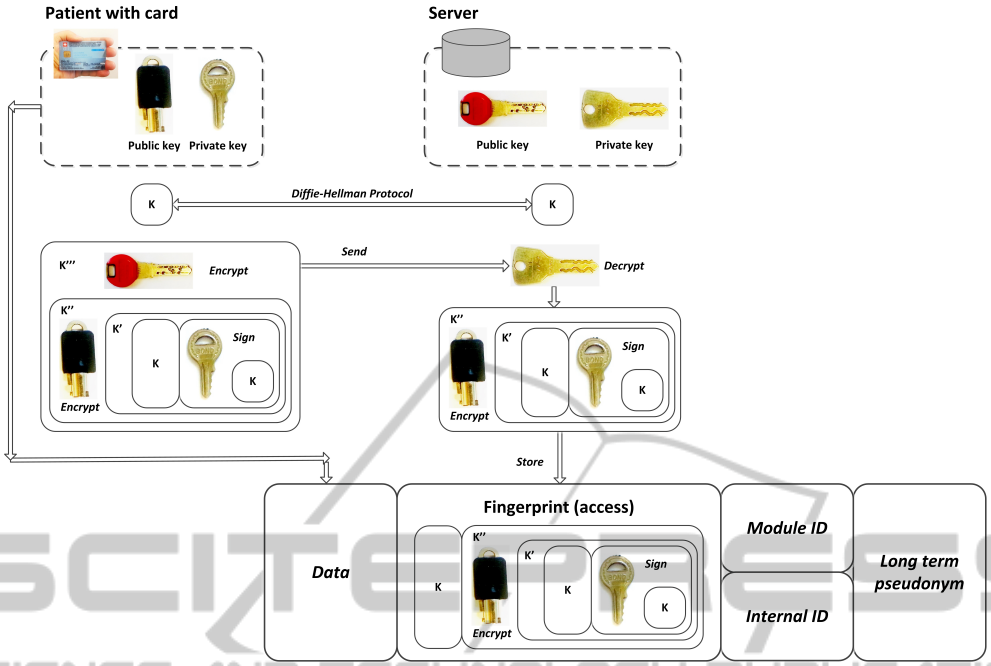

Figure 2: Enrollment.

database. He/she has his/her patient's smartcard and inserts it into a reader (client device). The communication between the physician's computer and the central server uses a secure channel. Since both the physician and the central server know each other, this channel does not require to offer any anonymity. We use a HTTPS (i.e. TLS) channel for securing the communication that can neither been intercepted nor modified by a third party.

In our protocol, we will use S to denote the server containing data and P to denote the patient (patient's smartcard). S has a private-public key pair $(Pub_S, Priv_S)$ and a certificate $Cert_S$ containing Pub_S and signed by a trusted certificate authority. P has also a private-public key pair $(Pub_P, Priv_P)$. The patient should remain anonymous on the server, therefore neither the patient's certificate cannot be known by S nor can $Priv_P$ (the public key of P) be known by the server since it could be used as a unique identifier.

We do not discuss the PKI infrastructure design for this protocol, we assume it simply exists and is adequately deployed.

Protocol:

1. S and P create a shared message K using the Diffie-Hellman protocol, so that S and P both know K
2. P signs the message K and produces K' - the concatenation of K and $sign_{Priv_P}(K)$:

$$K' = K + sign_{Priv_P}(K)$$

3. P first encrypts K' using Pub_P

$$K'' = enc_{Pub_P}(K')$$

4. P then encrypts K'' using the public key of S

$$K''' = enc_{Pub_S}(K'')$$

so that the message is encrypted with both keys

5. P sends K''' to S
6. S decrypts K''' and gets K''

$$K'' = decrypt_{Priv_S}(K''')$$

7. S stores the pair (K, K'') together with the user's data

5.2 Re-Identification (Fig. 3)

The patient visits another physician that participates in this project (i.e. having an access to the server and its research database) for a consultation or a treatment. The health practitioner indicates the need to access to already stored data. The patient accepts the necessity of retrieving the data and gives a consent to do so. The server S is confronted with a patient P' pretending to be P and in order to accept the consent, has to verify his/her rights, without revealing his/her

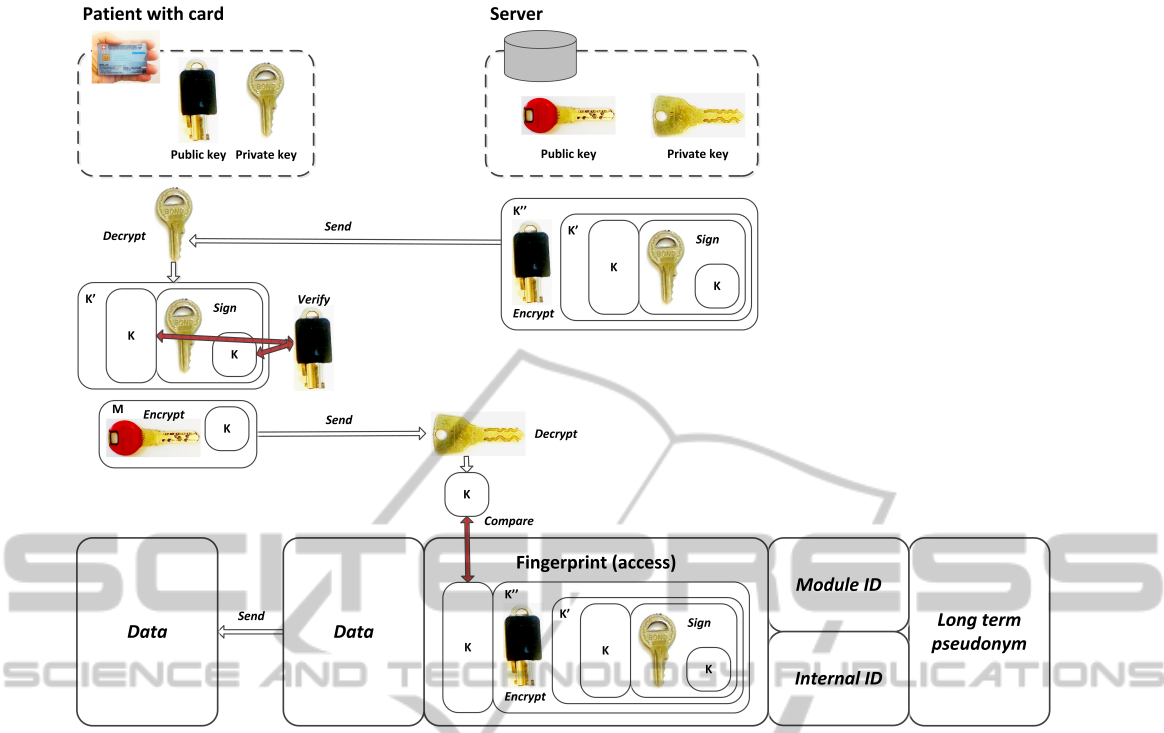

Figure 3: Re-Identification.

real identity. We speak of the re-identification of the anonymous patient.

S starts the second part of the protocol:

1. S sends K'' to P'
2. if P' is not P, the message can not be decrypted. Since P' does not know $Priv_P$. If $P' = P$ then the private key is known, and the message can be decrypted.

$$K' = decrypt_{Priv_P}(K'')$$

3. P separates the message in two parts: the message itself and its signature

$$K + sig = K'$$

4. P verifies the signature of the message to assert that the message has not been modified.
$verif_{Pub_P}(sig, K)$
5. If the signature is valid, then P sends back K to S, encrypted with the public key of S.

$$M = enc_{Pub_S}(K)$$

6. If the value received from P is the same as the value stored for the user, then S accepts the re-identification of P.

$$test\ if\ (K = decrypt_{Priv_S}(M))$$

5.3 Re-deployment of the Security Keys

Server key pair (K, K'') can only be used once, otherwise a replay attack would be easily successful. Therefore at the end of the re-identification process, the two partners will generate a new pair using the same protocol as in the enrollment.

5.4 Changing the Keys (Fig. 4)

In our protocol we have assumed that the patient's smartcard is of critical importance. It is used to create the access keys and to verify them in order to access the data. In real life such dependence is risky as the card can be lost, exchanged or upgraded. In such a case the access to the data would be irreversibly lost. Therefore we propose a backdoor procedure to transgress this limitation. Naturally, it is a trade-off between security and usability.

If the patient for any reason receives a new card, the health institutions cannot expect to be informed about it. A hospital will just be confronted with the situation that the patient does not own anymore the card that has been used to protect the data. The patient's identity should be however reliably verified to a seasonable degree with use of other documents, like a national identity card. In this case the access should

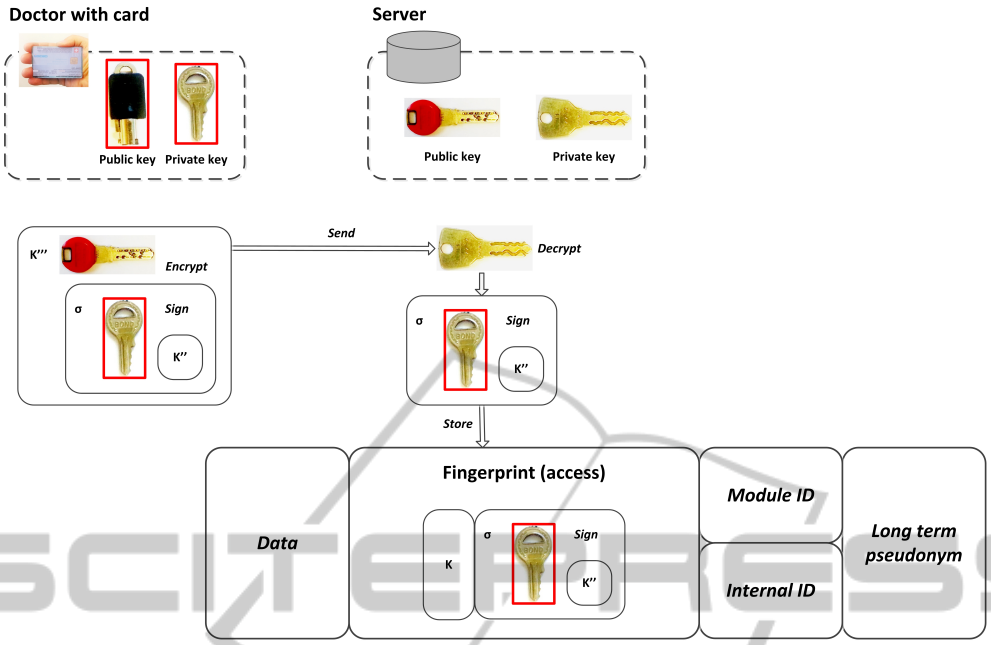

Figure 4: Signing by the doctor.

be refreshed by creating new fingerprint based on the new patient's card.

The doctor is the sole person verifying the identity of the patient. So a doctor could use this feature to gain access to undue records. In order to prevent abuse, the new fingerprint has to be signed by the doctor with his/her Health Professional's Card. This card also contains the private key and a certificate with the public key.

Handling of the case of card exchange - as discussed above - will induce an extension of the enrollment procedure. In addition to K and K", the server will store the identity of the doctor and the signature ($\sigma_{PrivM}(K'')$) created when signing the message K" with his/her private key. As during the re-identification a new access key (key pair (K, K'')) is created, all accesses will be signed by the doctors involved. This means that our database will contain the certificates of all participating doctors.

In this process, the doctor plays a role similar to certification authority in the PKI architecture as the trust in the doctor's integrity asserts the trust in the stored access key (fingerprint).

Since the doctor receives more power, the plausibility of the change of the key has to be verified. For instance, an alarm should be raised if a physician accesses a case without entering new data.

5.5 Discussion

The only requirement is that the patient has a X509 compatible card. Technically, it can be a e-Health card, a national identity card or any other valid card accepted by the physician. This card has to contain a private and a public key and a certificate signed by a Certificate Authority. Currently in many developed countries such card are being deployed. We need no supplementary features or software to be loaded on the card. The entire algorithm is implemented on the database server which is under control of the institution hosting the medical registry.

The proposed protocol has following important security features:

- Impossible for someone to register using someone else's identity (re-identification is not possible, since the wrong signature will not be accepted by P)
- Impossible to send a message to P in order to let P decrypt it, since P only accepts messages that were signed by him- or herself.

The scheme has to ensure a reasonable protection level but also has to be robust in practical situations. We assume that refusing the access to patient's own data is a real threat to the his/her health. Therefore although elevated security standards are important, they should not be fulfilled at any price. We have to consider not only the point of view of a computer security specialist, but also that of a medical practitioner.

6 FUTURE WORK

There is a number of problems related to our case that may be studied more profoundly.

In this paper we propose only a data access protocol - we do not discuss here how data are actually stored. We assume as self-evident that a semantic compatibility of data formats have to be ensured. An important question is if and how they are encrypted. If data were readable (or easily decryptable for some parties), it would be necessary to eliminate identifying information from the data content. This is for example the case in the DICOM headers of medical images, as presented in (Elger et al., 2010).

In the use case described here, the patient gives consent to access his/her data in the presence of the doctor. A set of records is retrieved and displayed, the relevant information is accessed. As long as the entire set concerns a specific disease, we may assume that the doctor can be trusted and can see it all. If the set covers various diseases, it may be useful to divide them in groups, possibly of difference confidentiality level. If the patients is HIV-positive and visits an orthopedist, the patient needs to have the freedom to decide if even the headers of the data records are visible.

The data accesses have to be logged in order to prevent and detect the cases of data theft. Also the failed accesses have to be logged. It is not so much the case of a malicious patient, because it one failed trial can happen and it would be difficult for the patient to try to read many data sets using many forged cards. On the other hand, a malicious doctor can do it quietly, not being disturbed.

7 CONCLUSIONS

In this paper we have considered a realistic case of retrieving valuable medical data stored in an anonymous registry with the consent of the patient concerned. The protocol we propose is a trade-off between security and privacy protection on one hand and usability on the other hand. If we devise a scheme to be massively used by patients, we have to remember that we deal with "common" people, many of them elderly, many of them having little experience in using computers. Therefore we should exclude following from our design:

- carrying special items
- remembering special secrets, like username/password
- upgrading standard items, like loading Java applications on the smartcard, especially by the patients themselves

Our protocol meets these requirements and provides a practicable solution to be used in the scope of the existing infrastructure. We do not expect anything special from the patients except that they have the identity token they normally use.

We have described the use of this protocol in the medical context, but it can be equally applied in other situations. We can think about any collection of anonymously stored data, where a data originator wants to recall records related to him/her. Moreover, he/she could trace its secondary use, if such information were stored in the collection.

In general, such scheme is useful when a large amount of anonymous data is collected for an acceptable goal, and the originator is allowed to retain the relation with his/her data. The main advantage is a reasonable privacy protection (and tracing the actions if the strict rules are loosened) and the simplicity of deployment.

REFERENCES

Camenisch, J. and Lysyanskaya, A. (2001). An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *Advances in CryptologyEUROCRYPT 2001*, pages 93–118. Springer.

Elger, B. S., Iavindrasana, J., Lo Iacono, L., Müller, H., Roduit, N., Summers, P., and Wright, J. (2010). Strategies for health data exchange for secondary, cross-institutional clinical research. *Computer methods and programs in biomedicine*, 99(3):230–251.

Ferguson, N., Schneier, B., and Kohno, T. (2010). *Cryptography Engineering: Design Principles and Practical Applications*. Wiley.

Gliklich, R. E. and Dreyer, N. A., editors (2010). *Registries for Evaluating Patient Outcomes: A User's Guide*. Outcome Sciences, Inc., AHRQ Publication No.10-EHC049.

Kwon, T. (2011). Privacy preservation with x. 509 standard certificates. *Information Sciences*, 181(13):2906–2921.

Wilson, S. (2005). A novel application of pki smartcards to anonymise health identifiers. In *AusCERT Asia Pacific Information Technology Security Conference Refereed R&D Stream*, page 64.