

Anonymous E-Prescriptions

Giuseppe Ateniese and Breno de Medeiros
Department of Computer Science
The Johns Hopkins University
Baltimore, MD 21218
{ateniese, breno}@cs.jhu.edu

ABSTRACT

This paper studies issues related to privacy protection of medical data, arguing that the topic is suitable for applied cryptographic research.

We present the problem of medicine prescription privacy and describe a practical system that employs standard cryptographic techniques to achieve several improvements over current practices. We also introduce a very simple tool: Online group signatures which can be built via simple primitives implemented in commonly employed cryptographic libraries.

Keywords: Medical Information Privacy, Public-key Cryptography, Privacy-preserving Cryptographic Techniques.

1. INTRODUCTION

The field of health care has several characteristics that make it worthy of individual consideration from the point of view of privacy. It is a vital societal goal to sustain health care systems that are widely available, affordable, and as inclusive of services as possible. In order to achieve these often conflicting goals, a heterogeneous set of institutions, each pursuing different objectives, must somehow collaborate.

The experience of providing or receiving medical care is a personal and private one. From early times the medical profession has recognized the consequential ethical implications. Quoted from the Hippocratic oath: *“Whatsoever I shall see or hear in the course of my dealings with men, if it be what should not be published abroad, I will never divulge, holding such things to be holy secrets.”*

The main goal of this paper is to study the problem of maintaining the confidentiality of medicine prescriptions. In the following section, to illustrate important and yet not obvious aspects of health privacy, we discuss the general medical record privacy problem, where we include references to mass media articles and other non-technical sources. Next, we identify the relevant entities involved with issuing and filling a prescription and describe the minimal interactions

between them. Finally, we propose a practical system, using only off-the-shelf cryptographic tools, that would better protect patients’ confidentiality in current prescription practices. The proposed system would provide further benefits to patients, by making it harder for insurance companies to engage in discriminatory practices, without hindering their legitimate activities.

1.1 Electronic Medical Records

In the past, a patient’s medical records were kept in paper files at the patient’s family doctor practice. However, with technological evolution, medical practice has now changed to an extent that patients receive care from medical teams. Medical records became the means to support collaboration and information exchange between a patient’s primary provider, and secondary providers, the latter including hospital surgical teams, referral specialists, laboratory personnel, and nursing staff, to name a few ([30]). As a result, once unified medical records were ported into electronic format as a collection of records: Images captured in radiological and other imaging laboratories, encoded diagnostics from a specialist, hospital admittance, discharge and treatment records (ADTs), and an array of other sources of data. These records are created and stored in separate databases, often belonging to distinct institutions ([30]). Upon request of (authorized) personnel this data is gathered, often in real time, from a medical intranet or over the Internet, and combined by the user interface into an organized entity. When addressing security and privacy concerns, it is important to keep in mind that the electronic data about a patient is scattered in different repository systems. For an example of one such system, see [34].

The prevalence of electronically stored information has had tremendous impact in reducing the costs of research based on the review of clinical data. The positive side of this is potential for faster assessment of new treatments, and greater dissemination of outcome effectiveness of health care practices. On the other hand, without adequate measures to protect the privacy of patients, there will be (and there has been) potential for real harm to patients. For instance, some patients are likely to hide health information from their doctors if they do not trust that it will be kept confidential. This is particularly acute when they fear their employer will become aware of their past medical history ([22, 4, 11]).

The current prevalence of electronic medical records (EMRs) has also facilitated a revolution in the financing of health care. Insurers and providers of managed health care services now engage in extensive and detailed research on the cost-effectiveness of specific medical treatments and practices.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WPES’02, November 21, 2002, Washington, DC, USA.
Copyright 2002 ACM 1-58113-633-1/02/0011 ...\$5.00.

The cost of such research was often prohibitive before electronic patient records were available and computer networks allowed the broad dissemination of such data ([36, 37]). In the United States, this research has fueled the success of Health Management Organizations (HMOs) over traditional indemnity systems. The HMOs competitive advantage resides in their active management of the benefits offered to keep costs low. From the privacy viewpoint, it would seem desirable that this research be conducted with anonymized EMRs. The simplest approach, namely stripping EMRs of all identifiers before making them available to researchers, is not always the best solution; there are instances in which patients benefit from being traceable by researchers, such as in the assessment of treatment safety ([40]).

2. MOTIVATION

Assuring security and privacy of personal information in health care systems has acquired greater urgency in the United States since the Congress approval of the Health Information Portability and Accountability Act (HIPAA) in 1996, which established the legal basis for patients' rights to privacy at a federal level. Compliance with this law has required considerable investment by health care organizations. Yet, it is widely predicted that many such organizations will not reach full compliance soon, and will eventually be vulnerable to lawsuits. In this respect, much has been written on how to write organization-specific privacy policies that meet the standards set by the HIPAA rules, and how to certify information systems as compliant with the formulated privacy policies. Such preoccupations, critically important as they are, should in our view be complemented by research and development of technological solutions to specific problems arising in the management of confidential health care information. There are in our view, two primary advantages to such an approach: cost-effectiveness and flexibility.

The HIPAA rules were designed as a *minimum* set of requirements. Many American states, not to mention other countries, have more stringent rules regarding the privacy of patients. Adopting solutions that too closely follow the letter of the HIPAA rules may result in systems that are tied to the legal requirements of particular jurisdictions. Beyond specific requirements, HIPAA introduces a novelty in American federal law by establishing the patient as the owner of her medical records – a principle also recognized in other countries. As a result, judges may interpret these laws as endowing courts with authority to redress personal harm traceable to a breach of confidentiality, even if the accused institution complied with all existing rules ([3]).

Apart from being more general, a technology-driven approach can make privacy breaches impossible rather than merely improper. Moreover, the application of rules can be automated, with potential cost-savings. One of the most comprehensive technical studies in the subject, prepared by the Committee on maintaining privacy and security in health care applications, of the National Research Council, has also advocated research into technical solutions to privacy issues ([41]).

2.1 Identity and Prescriptions

Much of the need to access medical data is dictated by health practice workflow. This includes availability of records to personnel in a patient's regularly visited clinic, but also to staff in medical emergency facilities, to technicians in labo-

ratories, to pharmacists, etc. Moreover, each of these facilities must communicate with paying organizations for billing purposes. In the case of medical prescriptions, more interactions are necessary. Clearly, the pharmacist handling the prescription collects information. Today this information explicitly includes the patient's and doctor's identity, and this is done not only for purposes of billing. Laws and regulations concerning distribution of prescribed medicine give law enforcement agencies access to this data, and require pharmacists to maintain records readily available for inspection by authorized personnel ([38]). If medicine prescriptions were substituted by truly anonymous tokens, there is good reason to believe that abuse would be widespread: Already today there is a thriving black market on prescription medicines, some of which have been obtained through unregulated web sites or mail order suppliers ([24]). Another legitimate reason pharmacists have to ask for patients' identities is to prevent harmful drug interactions. Several American states require pharmacists to keep record systems that can be used to identify and prevent known drug interactions ([38]).

Apart from the reasons above, pharmacists have potent financial incentives to keep records about patients' drug purchasing and doctors' prescribing patterns. This data is often collected in large databases, owned by business enterprises known as Pharmacy Benefits Management Systems (PBMs). According to a study conducted by the Health Care Financing Administration, an agency of the United States government, over 99% of all prescription claims are processed by PBM claims systems ([42]). Health insurance plans offering prescription benefits have agreements with PBMs, which purchase drugs in large volumes at discount prices, passing some of the cost savings on to the health insurers. PBMs enter in agreements with pharmacies to obtain information from prescriptions, such as doctor specific prescription patterns, patients' medical and prescription histories, drug interactions, and outcome effectiveness data. This information is later resold to the HMOs, where they help cost-containment research and prescription guideline enforcement. PBMs also sell this information to drug manufacturers, which use them as resource for drug development efforts, but also for marketing research and even direct marketing purposes ([50]).

2.2 Expectations of Confidentiality

Most people assume that pharmacists are bound by the same ethical rules of confidentiality as doctors and must obey laws that protect the privacy of prescriptions. In reality, there is no federal law in the United States specifically to protect the privacy of records kept by pharmacies, nor is there a legal presumption of equality with other medical records. Individual states have adopted regulations asserting the confidentiality of records kept by pharmacists. We give as an instance the State of Ohio ([49]). The differences in standards between states is large, with some states having no confidentiality provisions whatsoever. As a result, breaches of medical privacy are commonplace. In a 1997 survey of Fortune 500 companies, it was found that 35% of surveyed companies used medical records in employment related decisions, such as promotions or hiring. This practice was disclosed to employees in 90% of the cases ([26]).

These preoccupations are hardly hypothetical, as there are many reported cases of damaging use of medical infor-

mation. A Maryland banker which was part of a state medical board was able to procure state-wide records of cancer patients and called in their loans ([25]). In case that made headlines in the United States, employees of a railway company were requested to provide blood samples, and without their knowledge or approval, their blood was tested for a gene associated with a predisposition for carpal tunnel syndrome ([51]). The results were given to the employer. In yet another report, a Philadelphia man taking HIV suppressant medication paid his prescription with the health benefits insurance provided by his employer and found later that the information percolated to his company's human resources department, which contacted his supervisor ([21]).

The Privacy Rule of The Health Information Portability and Accountability Act (HIPAA) establishes that all medical records (independently of how the data was gathered) are property of the patient. Each access to a patient's medical information must be explicitly authorized by the patient, though this authorization can be assumed in cases where there is expectation of legitimacy and requesting explicit authorization would be excessively burdensome. The law has enough caveats for such legitimate uses that it allows current practices in health care to continue. For instance, the aforementioned brokering and trade of prescription data will probably continue with few if any imposed modifications. On purely technical grounds, one of the main changes imposed by the law is the requirement of auditability of record access. Each access to a patient's health records (whether explicitly authorized or not) must be recorded and added to the patient's records. The patient has the right to request disclosure of his/her records at any time (including the access logs) and to request that corrections be made in case erroneous information is present in the records. For reference, see the Office for Civil Rights website ([47]) and their various guidelines and introductions to the law ([48]).

The impact of the law is likely to go beyond its letter. As mentioned, the law has affirmed the patient ownership of medical data. One may expect that courts of law will interpret these rights as placing constraints in what constitutes legitimate data management and exchange. For instance, while HIPAA does not mention prescription records per se, a New York judge has found that patients have reasonable expectations of confidentiality of the records kept by pharmacies ([3]).

Doctors also have a privacy stake in prescriptions. In fact, doctors have worried that aggregate information about all of a doctor's yearly prescription data are routinely compiled by HMOs. This data is then compared with prescription patterns of other doctors and used in cost-control research. Some HMOs have proceeded to write guidelines on the frequency of prescribing certain drugs, and in a at least one instance promoted compliance with those guidelines by giving cash bonuses to doctors ([2]). Other doctors worry that they may be dropped from the plans if they do not follow the guidelines. Aggressive direct marketing strategies by drug manufacturers pull doctors in the opposite direction, offering perks to doctors that prescribe large amounts of certain medicines.

3. PREVIOUS WORK

As we shall discuss later in more detail, privacy of medicine prescriptions may involve protecting the identities of both doctor (medicine prescriber) and patient ([2]). From the pa-

tient's viewpoint, the confidentiality of a medicine prescription is similar to the confidentiality of other arbitrary items in his medical record. We start by reviewing work that has addressed the security and confidentiality of general medical records.

Important contributions are to be found in the work of Ross J. Anderson, including a description of a formal security model for clinical information systems ([6].) In this paper, the author proposes that such systems should implement multilateral forms of access control and information flow control, including support of access control lists, lattice-type hierarchical access structures, append-only write access supporting audit trails, and strong authentication. The suggested policy model has elements in common with both Bell-LaPadula and Clark-Wilson security models, specially with the latter. Other important works by the same author includes review of proposals by the British Government for the integration of clinical data systems ([5, 7].)

Again, we refer the reader to the comprehensive report elaborated by the National Research Council to provide guidance to the formulation of the rules and standards called for by HIPAA: [41].

A different aspect of the privacy of prescriptions is that involving the identity of the prescriber. The article ([28]), by Matyáš Jr., Václav, explored how to protect doctors' identities while at the same time allowing statistical data to be available for the purposes of research and analysis. This work describes ways in which data could be aggregated for each doctor and then released in a summarized fashion for each drug prescribed. The identity of the doctors could be released or kept hidden, according to the doctor's preference.

We should finally mention existing electronic prescription systems which, while not providing real privacy, demonstrate the viability of such implementations. Doctors are being urged to adopt these systems, as they cut administrative costs and greatly decrease the chances for medical errors, thus reducing potential liability and improving patient care ([33].) In a typical example, doctors connect to a clearinghouse to verify that a particular medicine is covered by the patient's benefits plan and to check for possible negative interactions. The prescription is filled electronically and forwarded to the PBM associated with the patient's plan. The patient later goes to a pharmacy affiliated with the PBM, where the prescription can be retrieved. We shall use a similar transaction sequence in the architecture of our system.

4. THE DRUG PRESCRIPTION PROBLEM

Within the field of medical privacy, there are good reasons to focus on drug prescriptions. Medicine prescriptions are clearly defined both as entities - a token signed by a doctor - and as a process: Doctor issues prescription, prescription is taken to the pharmacy where it serves as evidence of authorized use of medication, medicine is delivered to patient and a payer (patient or insurer) is billed for the purchase. It is an easy problem to communicate but which, within itself, demonstrates most of the relevant issues in the general field of medical records privacy.

As for relevance, the privacy of medicine prescriptions is quite important within the overall context of medical confidentiality. We offer two reasons as to why this importance: On one hand, the knowledge of a prescription is often quite revealing of (at least some aspects of) the medical history

of the patient. Secondly, medicine prescription benefits are the fastest rising budget item in insurance plans that cover medicine prescriptions ([27]).

The procedure of issuing a prescription may entail:

- Adding entries to a patient’s medical records.
- Performing queries in expert systems for possible drug interactions or medical conditions which, if present, may counter-indicate the use of the medication. These expert systems may belong to the clinic where the prescription was issued, but also to the pharmacy filling the prescription or to the insurer.
- Collecting legal evidences that prove authorized use of the medicine, which must be kept by pharmacies for the purposes of compliance with laws and regulations.
- Issuing of claim forms for billing purposes.

In specific situations further events may use this information. For instance, the prescription, in combination with other parts of a patient’s medical records, may be disclosed at future points in time, together with the patient’s identity, for legal investigative purposes (such as in the context of a malpractice lawsuit) or to comply with other legal requirements, such as the right of review enjoyed by patients, or for other purposes of law enforcement.

4.1 Toward a Prescription Model: Entities and their Relationships

When a prescription is issued, a few parties are necessarily involved: A patient, a doctor, a pharmacist to fill the prescription, and an insurance agency to pay the claim (if applicable). In practice, several other entities have a role to play. The first task to model the medicine prescription problem is to identify all these participants and their responsibilities.

- **Patient** – The patient is the entity to whom the prescription is issued. One central issue of the prescription privacy problem is to protect the confidential identity of the patient. The security requirements of a prescription system may include that patients provide proof-of-enrollment in insurance plans, proof of consent to the release of medical information, or proof of consent to a course of treatment.
- **Doctor** – The entity issuing the prescription. A possible goal of the prescription privacy problem is to protect the identity of the doctor. The doctor must be able to provide a proof-of-capacity, which authorizes the doctor to issue the prescription.
- **Pharmacist** – The entity filling the prescription. In automated prescription systems the pharmacist is a server within a Prescription Benefit Management system. The pharmacist must collect enough evidence to successfully process prescription claims and satisfy requirements of law enforcement agencies. Pharmacists may also engage in activities related to fraud prevention, such as statistical analysis and profiling.
- **Insurer** – The insurer provides health benefits for enrolled patients. Insurers also engage in medical profiling of patients and doctors to estimate risks and set

premiums, and for investigative purposes. One of the goals of a prescription privacy system is to remove or reduce opportunities of use of medical information for discriminatory purposes.

- **Privacy officer** – Privacy officers are responsible for establishing compliance with all laws concerning the privacy of patients and/or doctors and with regulations on disclosure of medical information. The privacy officer of a organization keeps a database which translates pseudonyms into patient names.
- **Enforcement agent** – An enforcement agent works for a government institution, such as the FDA or FBI, and has legal authority to oversee the prescription of medicines and/or certain controlled substances used for medical treatment. Prescriptions should be linkable per-patient and per-doctor basis from the enforcement agent’s point-of-view.
- **Judge** – A judge is any entity which has the legal authority to revoke a patient’s or doctor’s confidentiality in a transaction.
- **Certification Boards and Certification Authorities** – These are entities such as medical boards that grant doctors powers to issue prescriptions, and agents that issue digital certificates stating these capacities, possibly under pseudonyms. Certification Boards may also act as Certification Authorities in a prescription system.

5. SYSTEM DESIGN

We wish to define clearly the scope of our design. For instance, current practice determines that once a patient is admitted to a hospital, the hospital creates for him an admission record linked to some sort of label or Id. The processes that involve the patient’s records within the hospital, including the handling of medicine prescription records associated to this patient *are out of the scope of the system described in this paper*. We are concerned with the prescription as a token that is transferred between two business domains: From the health provider to a pharmacist. The entry and exit points where information is transferred between systems are known to be prime sources of information flows compromising privacy.

In order to reduce this information flow between domains, we must have in mind that individuals are identified within systems primarily for two distinct purposes. The first is access control, and its goal is to ensure that individuals can only access those system objects to which they have been granted privileges. The second is to manage the use of resources. These two aspects of authentication have been discussed at length in [1].

It is often the case that access privileges are defined in terms of roles. Thus authentication of a user’s role can often substitute for authentication of a user’s identity. On the other hand, resource usage management can be done under pseudonyms, as the real identities of each user are of no consequence – the system only needs to know that the user in question has used less than his/her allotted resources.

In the context of this paper, personal identification is necessary to establish a patient’s enrollment status with the insurer. Financial concerns, such as whether payments are

up-to-date, must be linked to the patient’s real identity, and are useful for billing purposes.

Several other concerns depend on a patient’s history of use of his account. For instance, available lifetime or annual benefit maximum, available per-disease or per-accident benefit maximum, year-to-date out-of-pocket maximum, remaining annual deductible. These items are relevant to claim managers and to provide a basis for payment of benefits. The real identity of the patient is not truly relevant for this type of accounting, and it can as well be made under a pseudonym.

Furthermore, review of claims and determination of eligibility for particular benefits should be made under the guidance of policies that apply on an equal basis to all the members of a plan, independent of knowledge of a patient’s real identity.

Health insurance claim managers control the use of resources through a pre-approval process. Currently, a patient wishing to receive a health care benefit submits a pre-approval request. The insurer reviews this request, determining both the patient’s account status and the patient’s eligibility. The first task assumes knowledge of the patient’s identity, while the second involves only verification against general plan policies and guidelines. Clearly, guidelines and policies may make references to the presence or absence of medical conditions. In the interest of fairness, all such references should be explicit; as such, they should be suitable to automatic verification from the knowledge of a few items of a patient’s medical history. A vague or overly specific policy, or a policy that requires close scrutiny of an individual’s lifetime medical history is conducive to discriminatory and arbitrary enforcement.

We now justify our assumptions about the security and confidentiality requirements of our system. Some of these assumptions are conventional and commonly used. Some others are results of our observations and have been abstracted from loosely specified privacy goals that were called for by medical organizations and consumer protection groups. Yet some others are forced upon us by legal requirements which we feel are unlikely to be changed as they have roots in concerns about public safety or potential for economic fraud. We expect that not all will agree with our choices, and hope that our work will encourage further discussions and developments in this subject.

The goal of protecting patient confidentiality must be balanced against potential for fraud in a truly anonymous system. Thus the privacy of the patient must be **revocable** under provisions of the law. For the same legal reasons, it is desirable that patient participation in transactions should result in **non-repudiable** evidence of patient engagement. More restrictively, transactions by the same patient should be **linkable** by the pharmacist/PBM. Otherwise current fraud-prevention investigative practices using statistical treatment would be rendered useless. One could argue that linkable anonymity is no anonymity at all. Our counter-argument is that patients should have the right to request a change in pseudonyms if they have reasons to believe their privacy is under risk of being compromised, and privacy officers may place reasonable restrictions on how often such pseudonym changes may take place. Another solution would be issuing different pseudonyms every time enrollment is renewed. That would limit histories to shorter periods of time, reducing risks of privacy breach while still allowing investigative profiling to take place.

Confidentiality of the doctor’s identity must be similarly revocable, and her participation non-repudiable. However in our view there is no good reason for doctors’ transactions to be linkable by the pharmacist/PBM. Instead fraud on the part of doctors could be investigated by the insurer. In other words, doctors’ transactions could (and probably should) be **unlinkable** for the pharmacist, but linkable from the perspective of the insurer.

Finally we address a point which is not one of privacy, but which protects patients against arbitrary enforcement of policies by a discriminatory insurer. Today pre-approval is used by insurers to enforce compliance with guidelines and to prevent excessive litigation. We argue that such practices, often perceived as arcane and an inconvenience, can bring benefits to both patients and insurers when automated within a secure electronic prescription system. Before issuing a prescription the doctor would contact a clearing house and verify that the drug is indeed in the list (formulary) of the patient’s plan. She would also have the ability to check for harmful drug interactions with other drugs taken currently by the patient. Only when satisfied that the drug could be safely prescribed and reimbursable, would the doctor issue the prescription, which would be directly deposited in the PBM system affiliated with the patient’s plan. The insurer would be under obligation to honor the claim later; refusing to do so would in most circumstances be used by the patient as proof of insurer’s arbitrary or dishonest behavior. Exceptional situations where the insurer would be allowed to refuse payment include suspicion of fraud – to be settled later through legal proceedings – or when the patient has reached the benefit limit, for instance.

PBMs would not be allowed to dynamically substitute medicines once the prescription has been approved by the clearinghouse. Such substitutions happen today, even when the particular medicine prescribed by the doctor is covered in the patient’s plan. For instance, PBMs sometimes stop supplying drugs – with the exception of painkillers and “comfort drugs” – to terminal patients admitted in hospices, using the rationale that the dying patients does not benefit from the expensive drugs. This medicine withdrawal can undermine the quality of life of these patients ([44]).

We are now ready to describe the confidential prescription system in terms of its sub-protocols. We first mention one assumption about the deployment model. We replace the traditional insurance identification cards patients carry today by smart-cards. These smart-cards need only to store information and to issue (not verify) standard digital signatures.

5.1 Patient enrollment and enrollment verification

The patient P enrolls in a health insurance plan by contacting an agent of the insurer. This agent processes identifiable information, such as name and address of the enrollee/patient, and may perform other tasks, such as verification of employment or payment. Once an account for the patient is created, account data is forwarded to the insurer’s privacy officer, denoted here by PO_1 . The PO_1 generates a random pseudonym $Label_P$ and enters it together with the account data in her personal, secure database. PO_1 also initializes a smart-card storing $Label_P$, a descriptor of the patient’s policy, a signed expiration date and some key material. The PO_1 mails this smart-card directly to the pa-

tient/enrollee's address.

Later the patient can prove his enrollment status to his doctor by presenting the smart-card. The doctor can read the expiration date in the smart-card, signed with the insurer's public key.

Since the smart-card alone is sufficient to determine eligibility for benefits, patients must be made aware of the need to communicate the loss or theft of smart-cards. On the other hand, it provides flexibility in some instances, such as when carried by the **legal guardians** of incapacitated patients instead of by the patients themselves.

5.2 Doctor certification

Today several medical professional organizations certify doctors with power to prescribe medicines. These institutions could maintain Certification Authorities to issue doctors digital certificates, describing doctors' various privileges such as issuing prescriptions. Alternatively, doctors could obtain digital signature certificates by joining health care provider organizations. In this second case, the certificates should identify which organization originally certified the doctor as capable of issuing prescriptions, with the provider organization basically serving the function of a public notary in issuing the digital certificate.

Doctors could join one or more health care provider organizations. For instance, a doctor might join several different HMOs, PPOs, or become a member of a hospital or of a doctor organization with pre-negotiated contracts for reimbursement of benefits from insurance providers. Of course, doctors could negotiate terms of membership or reimbursement directly with a particular insurance plan. After a doctor had entered one such agreement, she would be contacted by the privacy officer (PO_2) of her provider organization. This PO_2 would issue the doctor a pseudonym, and one-time use certificates or a group signature certificate. Using these certificates doctors could issue anonymous, unlinkable signatures.

Remark: For the purposes of this paper, a doctor is any entity with legitimate capacity for issuing prescriptions. If a paramedic or nurse is certified or licensed to write certain types of prescriptions, she could also become a participant in such a system – *though her type of digital certificate could be different from that of a full doctor, to indicate limitations in her capacity to prescribe medicines.*

5.3 Pre-approval

In order to issue the digital prescription, the doctor connects to a digital clearinghouse to obtain pre-approval. She locates a form to fill for the prescription. This form should mention any policies that will be enforced by the patient's benefit plan. For instance, a prescription for an antibiotic could require a diagnosis for a condition within specified categories of diseases. The doctor fills all the required fields, including any diagnosis fields. The form does not require identification of the patient, only of the type of plan he is enrolled in. The form would also display possible harmful drug interactions and require the doctor to vouch that the patient is not at risk for such interactions. This form is used to generate a prescription token, certified by the clearinghouse. This certification acts as evidence that the patient can use for reimbursement, a guarantee that the PBM will not later deny the claim or substitute the medicine prescribed.

The doctor then concatenates the prescription token with

the patient's pseudonym and other transaction identifiers, plus a **conditionally linking token** and issues her unlinkable signature. The complete prescription is forwarded to the PBM affiliated with the patient's plan.

5.4 Filling

The patient gives his smart-card to the pharmacist. The patient's pseudonym is read, together with a description of his benefits plan. The corresponding PBM is contacted and the appropriate prescriptions retrieved. The prescriptions are submitted, one at a time, to the patient's smart-card, which signs them for the pharmacist. The pharmacist forwards the digital tokens to the PBM, collects an electronic payment, charges any remaining co-pays to the patient and delivers the medicines.

A possible modification to this scheme would be to require that all current prescriptions be also stored at the patient's smart-card. While not necessary for our purposes it could be helpful in a medical emergency or in the situation of a communication breakdown.

Remark: In case a patient is in continuous care, such as when admitted to a hospital or long-term care facility, this process is different. In that case prescription issuing and filling is done "in-house" within the network or intranet of the medical facility, and it is not convenient to assume the need for the patient's smart-card to authenticate individual transactions. Still at the time of closing the bill, the patient or his guardian must review the hospital charges and sign claim forms; at that point, prescriptions are handled separately and are submitted to PBMs for claim processing, on an itemized basis. So the *model* of the process remains unchanged for the purposes of this design.

5.5 Anonymous Profiling

The PBM would profile patients and health care providers profile doctors for investigative purposes and to prevent fraud. They would engage in such actions on their own initiative or under the request of an officer for due process of law enforcement. The patient's records, while anonymous, are collected together under the same pseudonym by the PBM. Correspondingly, transactions by the same doctor can be linked together by the health provider organization to which the doctor belongs via the conditionally linking token, see section 6. Law enforcement officers would have ready access to these anonymized profiles, and in the cases prescribed by law, could request the privacy officer to have pseudonyms revoked.

5.6 Pseudonym Revocation

Only privacy officers would be able to open pseudonyms, and would do so according to pre-established privacy policies and guidelines disclosed to all participants, doctors and patients. In practice, suspicious results collected through anonymous profiling would be a typical starting point for proceedings leading to pseudonym revocation. It is to be hoped that lawmakers, industry leaders, consumer defense groups and judges would take initiatives to establish parameters under which such revocations would be considered either legitimate or abusive.

6. THE PROTOCOL

From the previous system description, it is clear that we need a cryptographic tool that would allow doctors to gen-

erate signatures which are anonymous and unlinkable but, at the same time, it should be possible to revoke anonymity in case of dispute. A cryptographic primitive that provides exactly what is needed is the group signature. In contrast to ordinary signatures, group signatures provide anonymity to the signer, i.e., a verifier can only tell that a member of some group signed. However, in exceptional cases, any group signature can be “opened” by a designated *group manager* to reveal unambiguously the identity of the signature’s originator. At the same time, no one, including the group manager, can mis-attribute a valid group signature. We denote a group signature on a message m by $GS(m)$, the rest of the notation used throughout the paper is shown in Table 1.

A doctor Dr must become member of a health care provider organization by joining the group of doctors and receiving the appropriate group public key and a private group certificate that can be used to prove actual membership. The patient, P , can enroll in a given health plan by interacting with the insurer’s privacy officer, PO . To accomplish this task, P generates a public key and prove to PO knowledge of the corresponding secret key by signing a random message and a time-stamp.

Patient Enrollment:

$$P \longrightarrow PO : PK_J, S_J(\text{Enroll_Req}, \text{Nonce}_0);$$

$$P \longleftarrow PO : C_J = S_I(PH_0, J, PK_J, E(P));$$

The protocol header PH_0 is defined in Table 1.

To correctly prove the enrollment in a approved health plan, the patient P releases to the doctor the signature $S_J = S_J(\text{Nonce}_1)$, computed under the public key PK_J , along with the certificate C_J , where J is the pseudonym generated for P . Notice that, the privacy officer will make sure that a single patient is not provided with multiple pseudonyms, i.e., pseudonym uniqueness is preserved and guaranteed by sharing data, or by employing the mechanisms presented in [32].

The time-stamp in Nonce_1 guarantees that the proof of enrollment cannot be reused by the doctor. The value $E(P)$ is optional and represents an encryption of the identity of the patient with an emergency key. Such an encryption may turn out to be useful when medicines are recalled and patients have to be contacted urgently while PO ’s database is not functioning properly or has been compromised.

The pre-approval phase requires the doctor to acquire a signature from the insurer stating that a prescription Rx is approved under a particular plan; this approval request may include information about a patient’s type of plan, some diagnostics, the prescription requested, and other information as required. This protects the patient against later disputes with the PBM or the insurer. To guarantee privacy, the insurer’s signature should contain only a general statement so to exclude any information that could link later the approval notification with a patient’s pseudonym.

Pre-approval:

$$P \longrightarrow Dr : S_J = S_J(\text{Nonce}_1), C_J;$$

$$Dr \longrightarrow I : GS_{Dr}(\text{Approval_Req});$$

$$Dr \longleftarrow I : A = S_I(PH_1, Rx);$$

$$Dr \longrightarrow Ph : D = GS_{Dr}(PH_2, A, S_J, C_J, PE_I(LP));$$

In the table above, P refers to the patient who owns public key PK_J . We use the symbol Rx for the prescription token, as this is an internationally accepted convention for “prescription.”

In section 5, we mentioned a privacy officer within the doctor’s provider organization. This privacy officer acts as the group manager in the group signature scheme. The doctor identifies himself as a member of the group of doctors by computing a group signature over a standard request message. The signature A can actually be pre-computed and could be stored in a database. The practice we recommend is to store all the approval signatures in a database managed by the health care provider. This would relieve the doctor from computing the group signature on the approval request message since he would only have to access a local database.

The doctor produces a group signature over the approval of the insurer and the patient’s proof of enrollment. The value LP is explained in the next section and allows the insurer to link different group signatures from the same doctor in order to enable anonymous profiling. For this reason, it is encrypted under the insurer’s public key to ensure that no one but the insurer could link group signatures. The signature D is sent to the PBM system in order to make the prescription available to any pharmacists. Alternatively, D could be stored in the patient’s smart-card. The doctor also keeps a local copy of D , preventing the PBM from maliciously or erroneously deleting it.

Filling:

$$P \longrightarrow Ph : S = S_J(Ph, \text{Nonce}_2);$$

The pharmacist (PBM) may decide whether to forward the prescription to the insurer or close the transaction offline without the intervention of the insurer.

The privacy officer may notify the patient about the transaction by mailing to P ’s address an appropriate notification form.

7. PRIVRX

In this section we describe the technical details about the cryptographic infrastructure. All systems described herein can be implemented efficiently using widely available cryptographic libraries.

The general framework we designed is called PrivRX. The system comes in two versions, one applicable when the group of doctors is a static group, the other for large and highly dynamic groups.

7.1 PrivRX v1: Static Groups and the LP value

The case of static groups of doctors can be managed via traditional group signature schemes. Group signatures are a relatively recent cryptographic concept introduced by Chaum and van Heyst [17] in 1991. Several schemes have been proposed, most notably in [18, 15, 14, 8]. Our system is based on the group signature in [8] where the group public key and the signatures are both of constant size.

The group public key is of the form $\mathcal{Y} = (n, a, a_0, y, g, h)$, where $n = pq$ with $p = 2p' + 1$ and $q = 2q' + 1$. The values a, a_0, g, h are randomly chosen in the set of quadratic residues $QR(n)$.

$PH_i, i = 0, 1, 2$	Protocol header that should minimally contain inception and expiration date, insurance and health plan identifiers, and a transaction ID
PE_X	semantically secure encryption under X 's public key
$Nonce_i$	$r t$, where r is a random number and t a time-stamp
$S_X(m)$	digital signature on m produced by X
$GS_X(m)$	group signature on m produced by X
$E_X(m)$	symmetric-key encryption of m that can be opened by X
J_X	pseudonym generated on behalf of X
P	Patient
PO	Privacy Officer
Dr	Doctor
I	Insurer
Ph	Pharmacist/PBM system

Table 1: Notation

The group manager GM , which acts as privacy officer in our protocol, sends to each participant/doctor Dr_i a membership certificate $[A_i, e_i]$, where $A_i = (a^{x_i} a_0)^{1/e_i} \bmod n$. In order to anonymously sign, the participant/doctor has to prove to possess a member certificate without revealing it.

In particular, Dr_i computes:

$$T_1 = A_i y^w, T_2 = g^w, T_3 = g^{e_i} h^w, SK.$$

The value SK represents a signature of knowledge of (1) a value x_i such that $a^{x_i} a_0$ is the value that is ElGamal-encrypted in (T_1, T_2) under GM 's public key y and of (2) an e_i -th root of that encrypted value, where e_i is the first part of the representation of T_3 w.r.t. g and h and that e_i lies in a certain interval. (See [8] for details.)

The doctor must allow the insurer to link several group signatures for anonymous profiling. Building a group signature scheme which allows signatures to be linked under some conditions is quite straightforward when using the scheme in [8]: The insurer I can generate a fixed base f that must be a quadratic residue modulo n . Then, the doctor computes the value LP in our protocol so that it contains f^{e_i} and a proof that e_i in T_3 is the same e_i in f^{e_i} . Such a proof of equality of discrete logarithms can be efficiently performed as shown in [16]. The value LP is encrypted so that only the insurer is able to use it to link group signatures. The value f^{e_i} can be seen as a pseudonym of the doctor and the proof of equality of discrete logarithms shows that the doctor who owns the pseudonym is also the same who signed the prescription.

Remark: Anonymous profiling comes with a price. As mentioned earlier, a signature of a particular doctor on a prescription can be opened by the privacy officer of the doctor's health care provider in case of dispute. For instance, law enforcement agents may accuse the doctor of misconduct and use the signature as proof in a court of law. However, revealing both the prescription and the identity of the doctor publicly may allow the insurance company to associate the identity of the doctor to all the prescriptions he signed even if they are legitimate and thus not under investigation.

Observe that this is an intrinsic problem of any system allowing anonymous profiling, thus it is technically impossible to solve. However, we may observe that the details of the opening procedure could be treated as confidential and never revealed to the insurance company.

Alternatively, one can mitigate the problem by dividing the time in periods and let a trusted third party generate a different base f , as defined above, for each time period. In this way, only the privacy of the prescriptions signed in the same time period of the ones that are being investigated will be compromised.

7.2 PrivRX v2: Dynamic Groups

Group signatures work smoothly as long as the group of doctors is stable. In case of dynamic groups, revocation of group certificates becomes a real bottleneck and it is in general a costly process. The problem of revocation in group signatures was initially pointed out in [10]. What makes this process difficult is the requirement that even if a member is revoked at a certain time, all his previously computed signatures must still remain anonymous and unlinkable. Several solutions have recently been proposed [12, 31, 9, 13] but none of them is effective in a highly dynamic environment.

In a dynamic group, doctors may join or leave the group frequently, forcing the group manager to continuously modify the membership list and recalculate group signature parameters. Any solution to the revocation problem has to minimally require updating the group public key when members are revoked. Since this may happen very frequently in dynamic groups, any existing or future solution would not be acceptable for our system. Moreover, notice that the group manager would still be involved each time an user is revoked which may suggest that an on-line approach is preferable.

In Appendix A we introduce on-line group signature schemes on which the version 2 of our system PrivRX is based. On-line group signatures are much simpler than their off-line counterparts however they are useful only when efficient and instant revocation of group members is needed. They may be applied to other more general contests and might be of independent interest.

PrivRX v2 is very practical since the overall cost of computing and verifying group signatures is constant, i.e., it does not depend on the number of users, either legitimate or revoked. It is also possible to employ standard and commonly employed cryptographic technology which makes the implementation process easier.

Each doctor has to have an on-line connection with the privacy officer/group manager who is the one that is actually computing the signature on the prescription on the doctor's behalf. In particular, by using direct on-line signatures (see

Appendix A), the following step in the pre-approval phase:

$$Dr \longrightarrow I : GS_{Dr}(\text{Approval_Req})$$

becomes ¹:

$$Dr \longrightarrow PO_2 : S = S_{Dr}(e = E_K(M)), Dr, PI = PE_I(M, K),$$

$$PO_2 \longrightarrow I : PI, S_{PO}(E(j), E_j(Dr, S), e),$$

where $M = \text{Approval_Req}$.

The actual signature on the prescription that the doctor sends to the pharmacist must also contain the value LP encrypted under the insurer's public key which allows the insurer to link group signatures from the same doctor for anonymous profiling. When using off-line group signature schemes, we needed a pseudonym and a proof of the equality of discrete logarithms to make sure that the doctor who signed the prescription was the same who owns the pseudonym. With on-line group signatures, we just need the privacy officer to assign to each doctor an unique label, $Label_{Dr}$, which does not contain any information about real identities. Therefore, the following step in the pre-approval phase:

$$Dr \longrightarrow Ph : GS_{Dr}(PH_2, A, S_J, C_J, PE_I(LP)),$$

becomes:

$$Dr \longrightarrow PO_2 : S = S_{Dr}(e = E_K(M), Dr, PP = PE_{Ph}(M, K)),$$

$$PO_2 \longrightarrow Ph : PP, S_{PO}(E(j), E_j(Dr, S), e, PE_I(Label_{Dr})),$$

where $M = PH_2, A, S_J, C_J$. We are assuming that $PE_I(\cdot)$ is a semantically secure encryption algorithm so that two encryptions of the same label cannot be linked.

8. IMPLEMENTATION ASPECTS

Recent developments in smart-card technology have made their use possible in many more contexts than the traditional strong authentication/ dedicated payment systems. For instance, the standardization of smart-card communication/powering protocols (via standard ISO/IEC-7816 [45]) have allowed smart-card systems to support a large number of different programming platforms, including multi-application support in a single smart-card. Similarly, it has permitted the development of multi-platform card readers.

The availability of flexible, expandable smart-card infrastructures has implications for interoperability: Today smart-card based systems can be developed with multiple-vendor support, and when deployed widely the per-transaction costs can be minimized levels similar to, or lower than, those incurred by credit card payment transaction.

In the model deployment scenario, patients would carry credit-card sized smart-cards containing the secret signing key and certificate in a tamper-proof environment. The protocol requires that the smart-cards be able to both verify and issue signatures on messages of pre-specified formats.

¹We are using a simplified notation, see Appendix A for details.

In addition, system designers might want to dedicate storage space in the smart-card for storage of the user's current prescriptions. This would add reliability to current systems, because the information would be available in case of emergencies and in other situations wherein prescription information might not have reached the point-of-care.

Another feature of the smart-card would be to support multiple pins, in case the patient or courts have conferred power of attorney for health care related decisions to third parties, such as a close family member.

A systemic change of the way prescriptions are handled today can only take place if a gradual migration is possible, with new systems accommodating older formats. Fortunately, the applicable communication standards for prescription claim processing – HL7 ([43]), NCPCD ([46]) and X12 ([35]) – have been designed for extensibility, and could be made to accommodate the changes suggested here. For instance, HL7 defines a data record for patient identifier, which comprises of both a patient identification string and an identification provider field - thus permitting a system-specific identification provider and a customized type of identification scheme. There are also provisions for specifying custom encryption schemes for identifying information. These already built-in capabilities of the standard could be used to support a pseudonym-based identification scheme. Similar changes could be adopted to substitute provider identification with the use of a group signature and a group identifier. The lack of a universal scheme for patient and provider identifiers has created a situation in which existing standards have capabilities for extensibility in the way identifiers are handled, facilitating the implementation of privacy-protecting techniques.

From a legal standpoint, the recognition of digital signatures imply that all workflows described in the system could take place in accordance with laws and regulations. Special support might have to be designed to facilitate oversight by local law enforcement agents that cannot count with the sophisticated technological backing now available to pharmacy and health care networks.

9. CONCLUSION

Electronic prescription systems are still at an early stage of adoption, mainly because of its perceived high costs. However, several independent studies of such systems have found benefits from fewer prescription errors and patient deaths. In the Harvard Hospital, adoption of an electronic prescription system reduced the rate of prescription errors from 140 per 100 patient days to just 25. Overall, prescription mistakes are believed to account for about 20% of all patient deaths attributable to medical errors, which are estimated at 44,000 to 98,000 a year. These numbers have lead to the creation of advocacy groups promoting the wide adoption of such systems ([39]). In this paper we have presented a practical system intended to provide privacy services to doctors and patients using electronic prescription systems; the system also provides stronger guarantees to all parties than it is achievable when paper prescriptions are used. The system proposed can be easily integrated into existing electronic prescription systems and takes in account the multilateral security needs of the health care field.

We hope our work will serve as a catalyst of further discussions. There may well be scope for refinement in our security assumptions and in the system design we propose.

There probably exist other services and business processes within the health care field that can be similarly studied. The authors plan to continue investigations in this direction.

10. REFERENCES

- [1] Martin Abadi. Two facets of authentication. In *Proceedings of the IEEE Symposium on Security and Privacy*. Oakland, 1998.
- [2] T. Albert. Doctors ask AMA to assure some privacy for their prescription pads. In http://www.ama-assn.org/sci-pubs/amnews/pick_00/prl11225.htm, American Medical News. 2000.
- [3] T. Albert. Records privacy extended to pharmacies. In http://www.ama-assn.org/sci-pubs/amnews/pick_01/prsb0402.htm, American Medical News. 2001.
- [4] A. Allen. Exposed. In <http://www.washingtonpost.com/wp-srv/national/longterm/exposed/exposed1.htm>, The Washington Post.
- [5] R. J. Anderson. NHS wide networking and patient confidentiality. In <http://www.cl.cam.ac.uk/ftp/users/rja14/bmj.ps.Z>, British Medical Journal. 1995.
- [6] R. J. Anderson. A security policy model for clinical information systems. In *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, 1996.
- [7] R. J. Anderson. Problems with the NHS cryptography strategy. In <http://www.cl.cam.ac.uk/users/rja14/zergo/zergo.html>. 1997.
- [8] G. Ateniese, J. Camenisch, M. Joye, and G. Tsudik. A practical and provably secure coalition-resistant group signature scheme. In *Advances in Cryptology — CRYPTO 2000*, volume 1880 of *LNCS*, pages 255–270. Springer-Verlag, 2000.
- [9] G. Ateniese, D. Song, and G. Tsudik. A Quasi-Efficient Revocation of Group Signatures. In *Financial Cryptography (FC '02)*, 2002.
- [10] G. Ateniese and G. Tsudik. Some open issues and new directions in group signatures. In *Financial Cryptography (FC '99)*, volume 1648 of *LNCS*, pages 196–211. Springer-Verlag, 1999.
- [11] A. Brandt. Privacy watch: A handful of insurance and medical industry companies want to sell your personal data, but you can stop them. In <http://www.pcworld.com/features/article/0,aid,43762,00.asp>, PCWorld.com. 2001.
- [12] E. Bresson and J. Stern. Efficient Revocation in Group Signatures, In *Proceedings of Public Key Cryptography (PKC'2001)*, Springer-Verlag, 2001.
- [13] Jan Camenisch and Anna Lysyanskaya. Efficient Revocation of Anonymous Group Membership. In *Cryptology ePrint Archive*, Report 2001/113.
- [14] J. Camenisch and M. Michels. A group signature scheme with improved efficiency. In *Advances in Cryptology — ASIACRYPT '98*, volume 1514 of *LNCS*, pages 160–174. Springer-Verlag, 1998.
- [15] J. Camenisch and M. Stadler. Efficient group signature schemes for large groups. In *Advances in Cryptology — CRYPTO '97*, volume 1296 of *LNCS*, pages 410–424. Springer-Verlag, 1997.
- [16] D. Chaum and T. Pedersen. Wallet databases with observers. In *Advances in Cryptology — Crypto '92*, pages 89–105, 1992.
- [17] D. Chaum and E. van Heyst. Group signatures. In *Advances in Cryptology — EUROCRYPT '91*, volume 547 of *LNCS*, pages 257–265. Springer-Verlag, 1991.
- [18] L. Chen and T. P. Pedersen. New group signature schemes. In *Advances in Cryptology — EUROCRYPT '94*, volume 950 of *LNCS*, pages 171–181, 1995.
- [19] S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. In *SIAM Journal on Computing*, vol. 17, number 2, 1998.
- [20] S. Halevi and S. Micali. Practical and provably secure commitment schemes from collision-free hashing. In *Advances in Cryptology — CRYPTO '96*, volume 1109 of *LNCS*, 1996.
- [21] C. Jabs. The myth of privacy: Technology is putting your medical history on public view-and you in jeopardy. *FamilyPC*, 2001.
- [22] N. Keene, W. Hobbie, and K. Ruccione. Childhood cancer survivors. In <http://www.patientcenters.com/survivors/news/jobs.html>, OncoNurse.com.
- [23] Hugo Krawczyk. The order of encryption and authentication for protecting communications (Or: how secure is SSL?). In *Advances in Cryptology — CRYPTO '01*, Springer-Verlag, 2001.
- [24] J. Ledbetter. Is buying drugs on the web too easy? In <http://www.cnn.com/TECH/computing/9906/29/drugs.idg/index.html>, CNN.com.
- [25] S. Lehrman. Keeping your genes private. *GeneLetter*.
- [26] D. F. Linowes and R. C. Spencer. How employers handle employees' personal information. In <http://www.kentlaw.edu/ilw/erepj/v1n1/lino-main.htm>, 1997.
- [27] S. Martin. Insurers project big increases in prescription drug spending. In <http://my.webmd.com/content/article/3645.103>, WebMD.com.
- [28] V. Matyáš Jr. Protecting doctor's identity in drug prescription analysis. *Health Informatics Journal*, 4.4, 1998.
- [29] T. P. Pedersen. Non-interactive and information theoretic secure verifiable secret sharing. In *Advances in Cryptology — CRYPTO '91*, volume 576 of *LNCS*, 1991.
- [30] E. H. Shortliffe and G. O. Barnett. Medical data: Their acquisition, storage and use. In E. H. Shortliffe et al., editors, *Medical Informatics*, chapter 2. Springer, 2nd edition, 2000.
- [31] D. Song. Practical forward-secure group signature schemes. In *Proceedings of 2001 ACM Symposium on Computer and Communication Security*. November 2001.

- [32] S. Stubblebine and P. Syverson. Authentic Attributes with Fine-Grained Anonymity Protection. Financial Cryptography 2000, LNCS Series, Springer-Verlag, 2000.
- [33] D. R. Weinstein and H. J. Worman. Electronic prescription services offer potential savings and improvements in efficiencies. In <http://www.temple.edu/gisection/aganew19.html>.
- [34] A⁴ Health Systems. <http://www.a4healthsystems.com/>.
- [35] Accredited Standards Committee (ASC) X12, charter of the American National Standards Institute (ANSI). <http://www.x12.org>
- [36] Agency for Health Care Research and Quality. Managed care plans develop research capabilities and partnerships. In <http://www.ahcpr.gov/research/may98/ra2.htm#head9>.
- [37] Agency for Health Care Research and Quality. Managed care research has come of age. In <http://www.ahcpr.gov/research/jul99/799ra11.htm#head3>.
- [38] California Board of Pharmacy. California pharmacy law. In http://www.pharmacy.ca.gov/lawbook_table_of_contents.htm.
- [39] Editorial. Computerized prescription systems cut error deaths—expensively. In http://cancer.med.upenn.edu/cancer_news/reuters/2001/apr/20010412publ004.html, Oncolink Cancer News. 2001
- [40] Food and Drugs Administration. Medwatch: The FDA safety information and adverse event reporting program. In <http://www.fda.gov/medwatch/>.
- [41] For the record: Protecting Electronic Health Information. Computer Science and Telecommunications Board, National Research Council 264 pages. Washington, DC: National Academy Press 1997.
- [42] Health Care Financing Administration. Study of pharmaceutical benefit management. In <http://www.hcfa.gov/research/pharmbm.pdf>, 2001.
- [43] Health Level Seven. <http://www.hl7.org>
- [44] Hospice Patients Alliance. When getting prescriptions filled, beware of medication substitutes! In <http://www.hospicepatients.org/substituteRx.html>.
- [45] International Organization for Standardization. <http://www.iso.ch>
- [46] National Council for Prescription Drug Programs Standard Claims Billing. http://www.hipaonet.com/hisb_ncdpd.htm
- [47] Office for Civil Rights. National standards to protect the privacy of personal health information. In <http://www.hhs.gov/ocr/hipaa/>. 2001.
- [48] Office for Civil Rights. Standards for privacy of individually identifiable health information. In <http://www.hhs.gov/ocr/hipaa/finalmaster.html>. 2001.
- [49] Ohio State Board of Pharmacy. Confidentiality of patient records. In <http://www.state.oh.us/pharmacy/rules/4729-05-29.html>. 1999.
- [50] RX News. Prescription privacy. In http://www.findarticles.com/cf_dls/m0815/1999_April/54425466/p1/article.jhtml?term=. Gale Group, 1999.
- [51] The U.S. Equal Employment Opportunity Commission. EEOC settles ADA suit against BNSF for genetic bias. In <http://www.eeoc.gov/press/4-18-01.html>, 2001.

APPENDIX

A. ON-LINE GROUP SIGNATURES

A *group-signature scheme* is a digital signature scheme comprised of the following five procedures:

SETUP: An algorithm that outputs the initial group public key and all system parameters.

JOIN: A protocol between the group manager and a user that results in the user becoming a new group member.

SIGN: An algorithm that on input a message m , and group member's secret key, outputs a group signature on m .

VERIFY: An algorithm for establishing the validity of an alleged group signature on a message with respect to a group public key.

OPEN: An algorithm that, given a message, a valid group signature on it, a group public key and a group manager's secret key, extracts a proof of the identity of the signer.

A secure group signature scheme must satisfy the following properties:

Correctness: Signatures produced by a group member using **SIGN** must be accepted by **VERIFY**.

Unforgeability: Only group members are able to sign messages on behalf of the group.

Anonymity: Given a valid signature, identifying the actual signer is unfeasible for everyone but the group manager.

Unlinkability: Deciding whether two different valid signatures were computed by the same group member is unfeasible.

Exculpability: Neither a group member nor the group manager can sign on behalf of other group members.

Traceability and Coalition-resistance: The group manager is always able to open a valid signature and identify the actual signer. In particular, a colluding subset of group members cannot generate valid signatures that the group manager cannot link to one of the colluding group members.

We assume that communication channels are private and authenticated. We also assume that each user has a universal public key which is authenticated and publicly available. Let M_i denote a user and $PK(M_i)$ his universal public key. The signature on a arbitrary message m under the public key $PK(M_i)$ is denoted by $Sign_{M_i}(m)$. Such a signature can be verified by anyone and in particular by a judge. We always assume cleartext signatures: If a party receives a message of the form $Sign(x)$ then that party can read x without the help of further cryptographic keys. (Except perhaps for publicly available information.) Any digital signature algorithm may be transformed into a cleartext signature algorithm by simply prepending the cleartext message to the signature.

During the **SETUP** phase, the group manager (GM) selects a digital signature algorithm and generates the corresponding public and secret keys. The public key is authenticated by a certification authority and published so to represent the public key of the entire group managed by GM . To summarize:

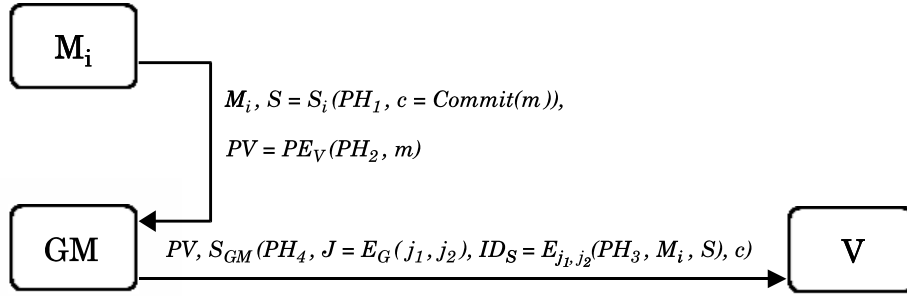


Figure 1: Direct On-line Group Signature

SETUP:

1. *GM* selects a digital signature algorithm and generates the corresponding public key and secret key. Moreover, *GM* selects a standard and widely accepted symmetric encryption algorithm, $E(\cdot)$ (for instance, Triple DES, Blowfish, or AES are all good candidates).
2. The secret key is kept private whereas the public key is authenticated and published, along with a description of $E(\cdot)$, as the group public key;
3. Only *GM* will be able to sign messages under the group public key. A signature on the message m will be then denoted by $S_{GM}(m)$.

Once the group public key has been published, users can start joining the group by interacting with *GM* during the JOIN phase. Initially, each user selects some standard digital signature algorithm and generates the appropriate parameters for it.

JOIN:

1. M_i selects a digital signature algorithm and generates the corresponding public key, PK_i . A signature on a message m , under PK_i , is denoted by $S_i(m)$;
2. M_i computes the membership certificate, $MC_i = \text{Sign}_{M_i}(PH, PK_i)$, and sends it to *GM*;
3. *GM* verifies the signature MC_i and stores it in a private database. Finally, *GM* sends a signed receipt to M_i .

The value PH is a protocol header that contains relevant information about the scheme. In particular, it should minimally contain the identity information of M_i and *GM*, a group identifier, a time-stamp and a *nonce* (a random number), a transaction identifier, and other pertinent information such as the digital signature algorithm selected (to be used with PK_i) and a sentence stating the intention of M_i to join the group according to the terms agreed upon with *GM*.

In case of dispute, the identity of the original signer must be revealed which requires the group manager to collect enough evidences to incontestably prove that a specific member indeed signed a certain message. The process of collecting evidences can be performed in a distributed way by storing into the final signature (and not locally) relevant

information about the signer. Another important aspect is confidentiality: The group manager should not know the content of the message signed by group members. The group manager, for example, may decide not to forward the final signature because of the content of the message received. To avoid this clearly undesirable case, the group member encrypts the message m and *instructs* the verifier on how to open such an encryption to retrieve the original message. The group member uses a *commitment scheme* to fix the values of the message to be signed. A very efficient (little more than a secure hash computation) commitment scheme is described in [20], with other frequently cited schemes being [29, 19].

We have designed two on-line group signature protocols. The first is called *direct* on-line group signature and assumes that the group member has direct access to the group manager. The second one, called *indirect* on-line group signature, assumes that the group member cannot communicate with the group manager but can, however, deliver messages to the verifier.

We will make use of the following notation:

$E_k(\cdot)$	symmetric encryption in CBC mode
$\text{MAC}(\ell, \cdot)$	cryptographic message digest with key ℓ
$E_{k, \ell}(m)$	$E_k(\text{MAC}(\ell, m), m)$
PH_i	protocol header
$PE_U(\cdot)$	encryption under U 's public key
G	master secret key owned by <i>GM</i>
$S_i(\cdot)$	signature under the key PK_i
$S_{GM}(\cdot)$	signature under the group public key

The SIGN protocol for direct group signatures is shown in Figure 1 and is run as follows:

SIGN (direct):

1. M_i computes the commitment $c = \text{Commit}(m)$ to the message; signs PH_1 and the commitment under the public key PK_i . The resulting signature, S , is sent to *GM* along with an encryption PV of the message m , to be forwarded to the verifier V .
2. *GM* checks whether S is a signature on c . If that is the case then *GM* signs the commitment c and includes inside the signature information about M_i and the original signature S . This information is authenticated and encrypted under freshly selected random keys j_2




Figure 2: Indirect On-line Group Signature

and j_1 . Such keys are also encrypted with a master key, G , known only to GM . Finally, GM 's signature and PV are sent to V .

The **SIGN** protocol for indirect on-line group signatures is shown in Figure 2 and is run as follows:

SIGN (*indirect*):

1. M_i computes a commitment c to the message; signs PH_1 and the commitment c under the public key PK_i . The resulting signature, S , is encrypted under GM 's public key and the result, PG , is sent to V along with the message m ; ²
2. V forwards PG to GM ;
3. GM checks whether PG contains a signature from M_i on c . If that is the case then GM signs the commitment c and includes inside the signature information (encrypted) about M_i and the original signature S . Finally, GM 's signature is sent to V .

The protocol header PH_1 should minimally contain relevant and publicly available information about the protocol such as the identity of M_i , GM , and the group. A nonce, time-stamp, and transaction identifier are also included as well as a sentence stating that the signature S should be interpreted in a special way as a signature on a commitment in accordance with the publicly available specification of the group signature protocol. Similar information is included in PH_4 so that the signature $S_{GM}(\cdot)$, computed by GM , will be interpreted correctly as a signature on a commitment generated by a group member.

We assume that $E_k(\cdot)$ is a symmetric encryption algorithm computed in CBC mode and secure against a chosen-plaintext attack, and that $MAC(\cdot)$ is a message authentication code built from hash functions, such as HMAC-SHA1 or HMAC-MD5, which are believed to be resistant against a chosen message attack. Under this assumptions $E_{k,\ell}(\cdot)$ itself can be thought of as providing a private and authenticated channel. Finally, we assume that the public-key encryption algorithm is resistant against the adaptive chosen-ciphertext attack. One practical way to achieve this is to compute $PE_U(m)$ in the following way: Generate two fresh random keys k_1, k_2 and encrypt them under any standard public-key

²We are assuming private and authenticated channels, so nobody can intercept and modify the message m being sent.

encryption (such as RSA or ElGamal) and finally compute $E_{k_1, k_2}(m)$.

Once the verifier receives the signature $S_{GM}(PH_4, J, ID_S, c)$ from GM , he can release:

$$S_{GM}(PH_4, J, ID_S, c), m^3.$$

To verify that the group signature on the message m is valid, it is enough to check first that $S_{GM}(PH_4, J, ID_S, c)$ is a valid signature on c , then verify that c is a commitment to m under the specified commitment protocol.

Given a valid signature $S_{GM}(PH_4, J, ID_S, c), m$, to open it and reveal the identity M_i , the group manager would run the following protocol:

OPEN:

1. GM decrypts $J = E_G(j_1, j_2)$ and reveals j_1, j_2 , and the identity M_i along with M_i 's member certificate ($MC_i = Sign_{M_i}(PH, PK_i)$).
2. To check whether GM has opened the signature correctly, it is sufficient to:
 - decrypt ID_S and verify the message authentication code by using j_1 and j_2 , respectively;
 - verify whether S is a signature from M_i on c .

In this way, GM can incontestably prove that the signer of the message was M_i . Only M_i could have produced the signature under PK_i . If GM cannot open a signature or refuses to do so, he will be held responsible for the content of the message.

³We assume that the necessary information to verify the commitment is also released along with m .