The Use of Communication Networks to Increase Personal Privacy In a Health Insurance Architecture

N. F. Maxemchuk
S. H. Low

AT&T Labs - Research
Murray Hill, New Jersey 07974

1. Introduction

If a national health insurance plan is implemented, the government has an obligation to monitor its use and to discourage any abuses. Individuals, on the other hand, expect a high degree of privacy in their health records. Communication networks make it possible to achieve both monitoring and privacy by separating an individual's identity and health records, and only associating the two when possible abuses are detected.

Recent advances in computing and communications make it reasonable to assemble and process the vast amounts of information that would be needed to administer and monitor a national health insurance network. Using computing and communications technology to operate on medical information has many side advantages, such as making an individual's complete health records available any place, almost instantaneously, and conducting research on correlations between treatments and diseases. However, once an individual's medical information is made available, it may also be misused.

Legislation may be used to eliminate uses of the information which are considered inappropriate. Society can decide that it is not appropriate to publish lists of individuals with communicable diseases, to provide medical histories to potential employers, or to sell information on diseases to medical supply houses for targeted advertizing, as is currently done with credit card purchases. However, legislation cannot prevent illegitimate uses of information. For instance, a system operator or computer snooper may obtain and resell medical information on celebrities or political candidates. In addition, flaws in an implementation may inadvertently release information.

Traditionally, communication networks have been used to bring information together. Recently, it has been shown that the same networks can keep information apart and increase personal privacy[1]. The objective is to construct systems that require sensitive information to perform their functions, but to make the total information unavailable at any single location. In these systems, we trust the various parties in the network not to collude with one another to join the information that we are trying to separate. However, we assume that the trust that we place in the parties is not always warranted. The parties may be compromised and release the information that they are protecting either because of human intervention or implementation error. The number of parties that must collude in order to extract information is a measure of how well the system separates that information. In reference 2 collusion analysis is reduced to determining the shortest paths on a graph.

A useful characteristic of these systems is that they can be designed to bring the information that has been separated together. The parties in the system are forced to collude by presenting them with the digital equivalent of a subpoena. The ability to subpoena information makes it reasonable to keep information private unless a law is violated. In effect, instead of deciding what information an individual has a right to keep private, and passing legislation to enforce that decision, an individual has the right to keep all information private unless the law is used to obtain that information.

The first application of information separation is an anonymous credit card[3]. This credit card separates an individual's identity and purchases. This mechanism can be used in electronic publishing[4] to preserve privacy except when a law is broken.

A concern in electronic publishing is that a document recipient can easily copy and redistribute a document. In reference 5, a technique is demonstrated that makes each copy of an electronic document different, so that a publisher can register the document with the original recipient. If an illegitimate copy is found, the publisher can determine the original owner. The privacy concern is that publishers can compile reading profiles on individuals. When the anonymous credit card is used to pay for the electronic document, the publisher does not know the recipient's identity. However, if an illegitimate copy is found, the publisher can obtain a subpoena to determine the identity of the credit card holder.

The anonymous credit card is a simpler version of the system that is being proposed to preserve privacy in a national health plan. Understanding the anonymous credit card can make the national insurance plan easier to understand. The architecture and functioning of the credit card is described in section 3.

In section 4 a candidate architecture for a national health insurance plan is described. This architecture is used to:

— make insurance payments from an employer, the government or an individual to an insurance carrier,

— make medical historys available, at an individual's discretion,

— provide health care providers with a subset of a person's medical history when that individual is unable to authorize access,

— submit claims for services from an authorized health care providers,

— make payments for services with a combination of insurance coverage and individual funds,

— audit a individual's or a health care provider's complete medical history, and

— identify the individuals or health care providers when grounds for insurance fraud are detected.

With this architecture, an employer or the government does not know the insurance provider that an individual selects, the health care provider does not need to know the identity of an individual being treated, and insurance companies can pay valid claims without knowing the individual or the health care provider.

This health care architecture is only one of a number of architectures that can perform these functions. Different types of information can be hidden and different degrees of protection against collusion can be designed into the system. The system can also be designed with different collusion paths, so that identities are not permanently hidden by corrupt or destroyed agents who do not reply to a subpoena. Implementing this type of system does not prevent a personal doctor-patient relationship. It just leaves it up to the patient to decide if she wants a personal relationship or privacy.

2. Party-to-Party Communications

Communications between two parties is conducted through an intermediary, using a double-locked box protocol. The use of intermediaries to hide information is an extension of a technique described in reference 6, in which an intermediary forwards electronic mail between two communicating parties. In the systems considered here, the intermediary

— hides the identity of the source and destination of a message from one another,

— hides the message destination from a vender or health care provider that owns the entry device being used by a customer or patient, and

— transfers trust between unknown entities,

  — to move funds between two banks that do not know each other's identity, or

  — to assure an insurance company that it is being billed by an unknown, but authorized health care provider, and that the charges are within the appropriate limits for the procedures performed.

A funds transfer between account i in bank 1 to account j in bank 2 is demonstrated in figure 1. The customer deposits a double-locked box in his account i in Bank 1. The first box can only be opened by the federal reserve, and contains the name of Bank 2 and a second locked box. The second box can only be opened by bank 2 and contains the account number j. When bank 1 transfers amount M for the customer, it withdraws the money from account i and sends message A to the federal reserve. The message instructs the federal reserve to deposit amount M in the account in the double-locked box. Bank 1 signs the message with its digital signature, so that the federal reserve can verify that it is from a bank that it trusts. The federal reserve opens the first box and sends a signed message to bank 2 to deposit amount M in the account in the locked box. The second bank verifies that the order was signed by the federal reserve, opens the second box, and deposits amount M in account j. The federal reserve is responsible for settling the accounts between banks.

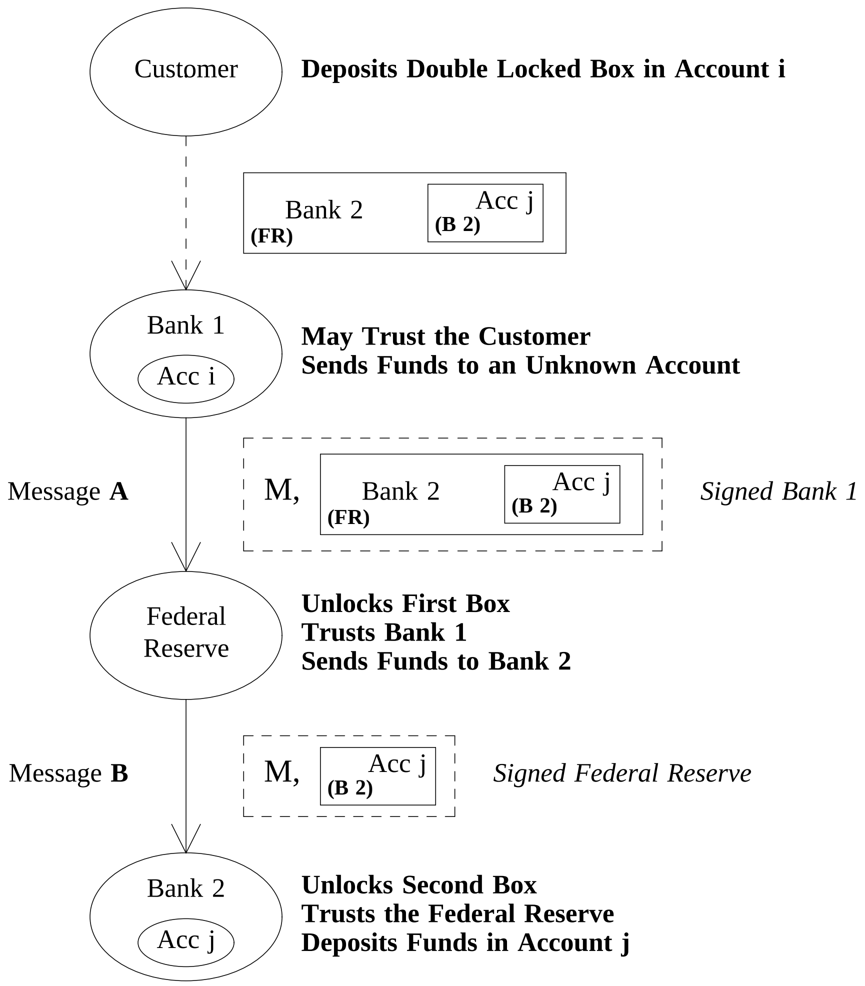

Figure 1. The Double Locked Box Protocol

When a customer must store the double-locked box for an account in several different locations, he places a different random number in each box, in addition to the information. This makes the double locked box that he places in the different locations different. If the double locked boxes were the same, the different locations might realize that they were communicating with the same account, and would be able to join the information that they had about their accounts.

The point-to-point communications between the parties and the communications exchange uses conventional cryptographic protocols. In this work, we do not explain these techniques, but assume that they can guarantee that:

a. an eavesdropper who intercepts a message can not decrypt it to determine its content,

b. the receiver is certain that the message was transmitted by the source, and not someone spoofing the source,

c. the receiver is certain that the message is not a replay of an earlier message,

d. the receiver can identify messages that were modified after the source transmitted them, and

e. the source is certain that the receiver has received the correct message.

How these characteristics are achieved, and how the double-locked box is constructed using public key cryptography is discussed in reference 3. How to perform the same functions with shared secrets is discussed in reference 2.

3. Anonymous Credit Card

The architecture for the anonymous credit card is shown in figure 2. The credit card company trusts the individual to repay his debt. Therefore, it is assumed that the credit card company will know the individual's identity. The store knows the merchandise that is purchased. The objective is to distribute the information so that a number of players must collude in order to associate an individual's purchases and identity.

The credit card company places credits in the individual's anonymous account using a double-locked box that the individual has placed in his credit card account. The bank with the anonymous account does not know or trust the individual. This bank trusts the federal reserve, which trusts the credit card company to pay the bill. At the end of the month, or when credit is exhausted, the bank with the anonymous account presents a bill to the credit card company, using a double-locked box that the individual has deposited in his account. The credit card company presents the individual with a bill, and when it is paid, deposits additional credits in the anonymous account.

An individual makes purchases in two phases, first he convinces the bank with the anonymous account that he is authorized to draw on that account, then he instructs that bank to transfer funds to the store's bank. Trust is transferred from the bank with the anonymous account, to the federal reserve, to the store's bank. The store's bank notifies the store that it as been paid. Since the store does not have to trust the individual to pay a bill, the store does not have to know the individual's identity.

In the anonymous credit card, a user must encrypt information at the store, and must use a double-locked box to communicate with the bank that maintains its anonymous account. It is assumed that a smart card[7], or similar technology will be used. In order to use the card, the individual proves his identity to the card. This can be accomplished by providing a PIN, answering a random set of personal questions, the answers of which have been stored on the card, or providing a biometric identifier[8], such as finger prints, that is stored in the card. Once activated by the authorized individual, the card can prove its identity to the bank. The card shares a secret key with the anonymous account. The card can prove its identity by encrypting a random number from the bank with the shared secret key. Communications between the customer at the store and his bank is through the communication exchange, and uses a double-locked box that can access the anonymous account. By communicating in this way, the customer does not reveal the identity of his bank to the store.

The customer receives a bill from the store, and a double-locked box that can access the store's bank account. He presents the bill to his bank. If there are enough credits left in the individual's account, the bank transfers the credits to the store's account, specified by the double-locked box. When the

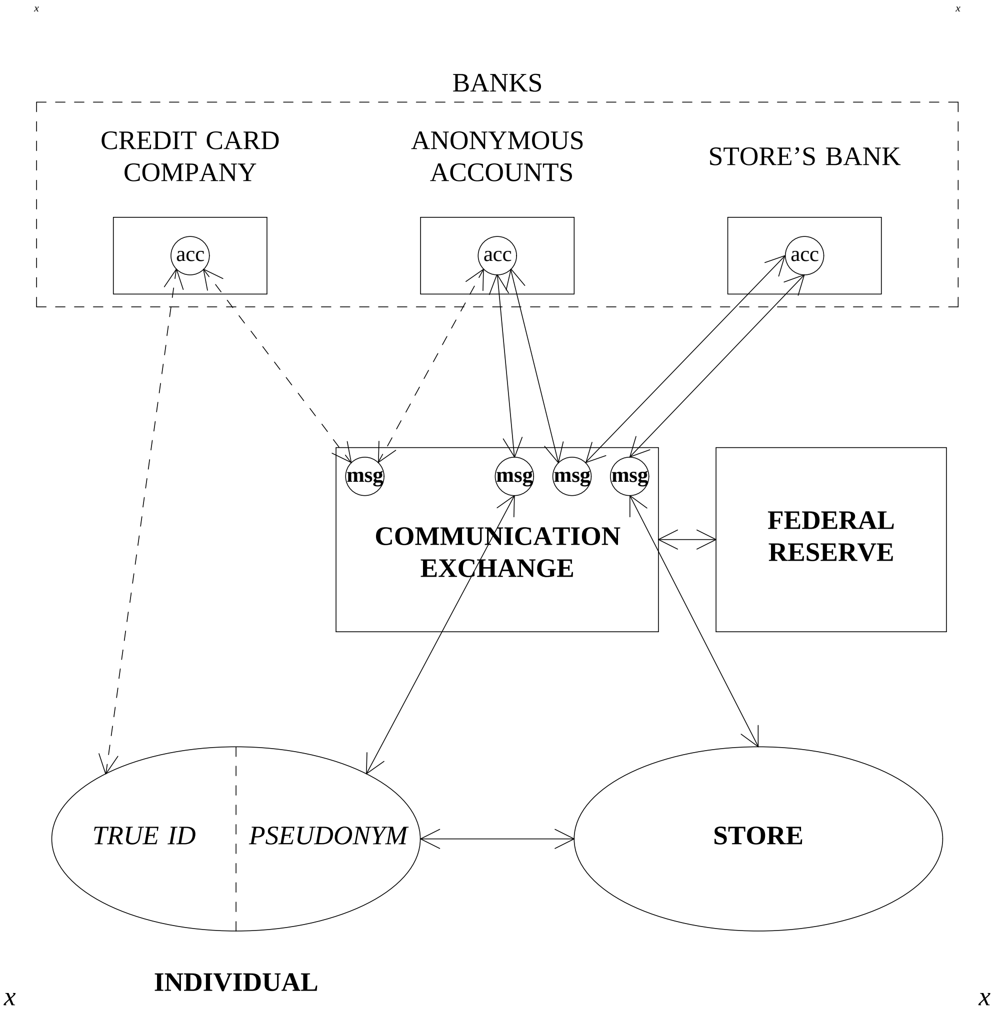

Figure 2. The Participants in an Anonymous Credit Card

store's bank receives the deposit, it notifies the store. When the store receive's confirmation of the deposit, it gives the customer the goods.

In addition to making purchases, with credit cards we expect

— to cancel lost cards,

— to detect unusual spending patterns,

— to receive detailed bills that augment our record keeping,

— to be able to return purchases, and

— to be able to challenge purchases that are charged to us.

These functions can all be implemented within this framework, as described in reference 1.

All of the messages in the credit card system that transfer funds are signed and must be uniquely identifiable. Otherwise, a replay attack that might cause funds to be transferred more than once. In addition, the source and destination of a message know each other's identity. It is assumed that messages that transfer funds are saved for a period of time, as is the current banking practice. If the equivalent of an electronic subpoena is issued against an individual, tracing the billing messages that the credit card company receives can locate the anonymous account. Tracing the charging messages that the anonymous account receives can locate the stores. Similarly, a subpoena against a store can force the store's bank to trace the messages transferring funds into the store's account back to the customer's anonymous accounts and from there back to the individuals. This is a secure tracking method, since the individual cannot erase messages in the trusted banks or the communication exchange.

Two parties that have been compromised can only collude to combine information about an account, if the accounts at both parties have a common, unique piece of information. For instance, if an individual's credit card account and anonymous account both kept his social security number, then the banks could determine that both accounts belong to the same individual. It is important to design the information distribution to eliminate as much common information as possible.

It is assumed that banks have many account, and that there are many purchases or bills of the same amount. Knowing that an individual has an account in a specific bank is not sufficient to identify the account. Similarly, knowing that a purchase was made for a specific amount and a charge for the same amount occurred, at about the same time, is not sufficient to link the two. Under these assumptions, the information in the anonymous credit card is separated so that the only unique information that can link accounts is the messages that are exchanged.

Figure 2 can be considered to be a graph. The nodes of the graph are the accounts in the banks, the store, and message exchanges in the communication exchange. The links of the graph are the lines that indicate that messages are passed. For each message that passes through the communication exchange, the communication exchange can associate the source of the message, the message that arrived from that source, the destination of the message and the message that was sent to that destination. At each bank, all of the messages that arrive at the bank or are transmitted to the bank are associated with a specific account. The bank that extends credit, also knows the individual's identity. The store associates messages with purchases. If a bank colludes with the communication exchange, for each message that was received or transmitted, it can determine the party that the communication exchange forwarded the message to or from, and the specific message received from or transmitted to that party. If the bank next colludes with that other party, it can find any information that the other party associates with that message. From this graph, it can be seen that an individual's identity and purchases can only be associated if all five parties, the credit card company, the communication exchange, the bank that maintains the anonymous account, the store's bank, and the store collude.

If the store eavesdrops on the customer, it can associate the customer's message to its anonymous account with the purchase. This additional piece of information allows the customer's identity and purchases to be linked when four parties, excluding the store's bank, collude.

4. National Health Insurance

In this section we consider one architecture for introducing both privacy and accountability into the national health insurance plan. The primary objective is to make it possible for an individual to control access to his medical history. Auditing operations can be used to identify individuals who are making claims for excessive treatments, health care providers who are claiming excessive services or services that they are not certified to provide, and insurance companies who are making excessive profits. By minor modifications, different types of information can be made more difficult or easier to join. It is also possible to introduce redundancy into the system to make misuse more difficult to conceal. For instance, if there is one player who must cooperate in order to identify an individual, but that player has been corrupted, then it may not be possible to identify the individual. However, if there are two independent ways to learn an individual's identity, then corrupting a single player will not hide the identity.

The parties in the insurance plan are shown in figure 3. An individual has a personal account with an organization that knows his identity, and anonymous accounts with a health insurance company, a database of medical histories, and an organization that can provide access to his medical history during a medical emergency. The individual also has the accounts associated with an anonymous credit card. It is assumed that each individual has one and only one personal account, as with a social security or internal revenue accounts. The individual's credit card company and employer are assumed to know the individual's identity. The individual's employer may have many employees and a health care provider may have many patients. Both place all of the information they have about a particular individual in an account.

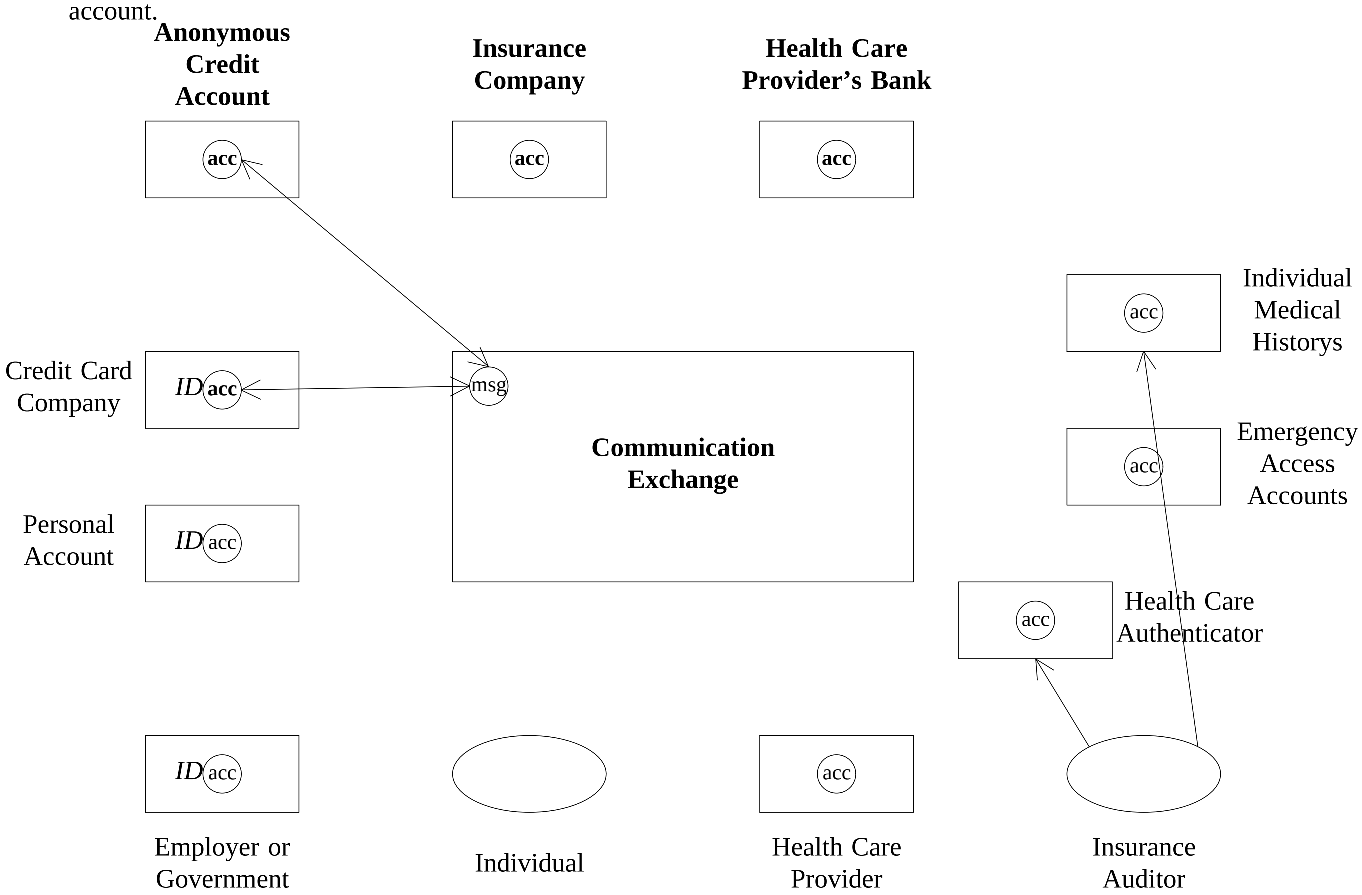

Figure 3. The Participants in the National Health Insurance

Each health care provider has accounts with a bank and with a health care authenticator. The health care authenticator knows the health care provider's identity, and is responsible for keeping track of the procedures he is licensed to perform and the history of the procedures that he has performed. In this architecture, a health care provider does not receive the same degree of privacy as an individual. However, if there is reason to give more privacy to the health care provider, the system can be made more symmetric.

An insurance auditor examines the histories in the individual accounts and the health care providers accounts. It can detect excessive charges, abnormal numbers of treatments by providers or for individuals, individual treatments that are inconsistent with medical histories, or any other indicators in the histories that may indicate that the insurance system is being abused.

In this architecture all of the payments to the insurance company and payments by the insurance company are made through the communication exchange. When funds transfers occur, the communication exchange operates as the federal reserve. The communication exchange has an account for each bank or organization that is authorized to transfer funds through it, and it is responsible for settling the accounts. By auditing an insurance company's account, it should be possible to determine if the company is making excessive profits, without revealing any information on the individuals that the company covers.

4.1 Paying for Insurance Coverage

An individual selects his own insurance company and coverage. The insurance company is given access to the individual's medical history, but does not need to know his identity. When the individual opens an account with an insurance company, he places two double-locked boxes in that account, one that accesses his personal account and another that accesses his medical history account. He also places double-locked boxes in his personal account and medical history account to access his insurance account.

The insurance company uses the link with the medical history database to determine the insurance premiums and the link with the personal account to present a bill to the individual. A filter can be placed in the medical history database to remove any information that policy dictates cannot be used to determine insurance premiums. An employer or the government may deposit credits that can be used to purchase medical insurance in the individual's personal account. Premiums for the health insurance are paid from the individual's personal account to the insurance company using the double-locked box protocol. If an individual selects coverage that exceeds his benefits, he receives a bill from the personal account for the difference. The individual uses his anonymous credit card, in the normal manner, to pay for the additional coverage.

The messages exchanged to pay insurance premiums are shown in figure 4. The messages indicated by dashed lines only exist if the coverage that a person selects exceeds his benefits. In this architecture, the organization that gives an individual medical benefits does not know which insurance company the individual selects. In addition to the functions that the communication exchange performs for funds transfer, it also verifies that the insurance company is a valid insurance provider when it communicates with the personal account or the medical history database, and assures the employer or government that it is placing insurance credits in a valid personal account. In a system in which funds are transferred between trusted parties, rather than being given to an individual, it is straightforward to guarantee that funds intended for health insurance are not used for other purposes.

4.2 Receiving Medical Treatment

A person's complete medical history is stored in an anonymous account in a medical history database. Since the database is connected to the communication network, the medical history is available anywhere, any time.

When a person accesses the medical history database while at a health care provider's location, he uses a smart card and proves his identity to the medical history database, the same way he did when using the anonymous credit card. The person does not have to make all of the information in this account available, but can pass the information through a personal filter that hides any information that the individual would rather keep private. Different filters can be used to provide different types of information to different health care providers. For instance, an individual need not provide his dentist with his psychiatric history. The health care provider gives the individual a double-locked box, which the individual passes along to the medical history database. The medical history database uses the health care provider's double-locked box to deliver the filtered information. These messages are the solid lines in figure 5.

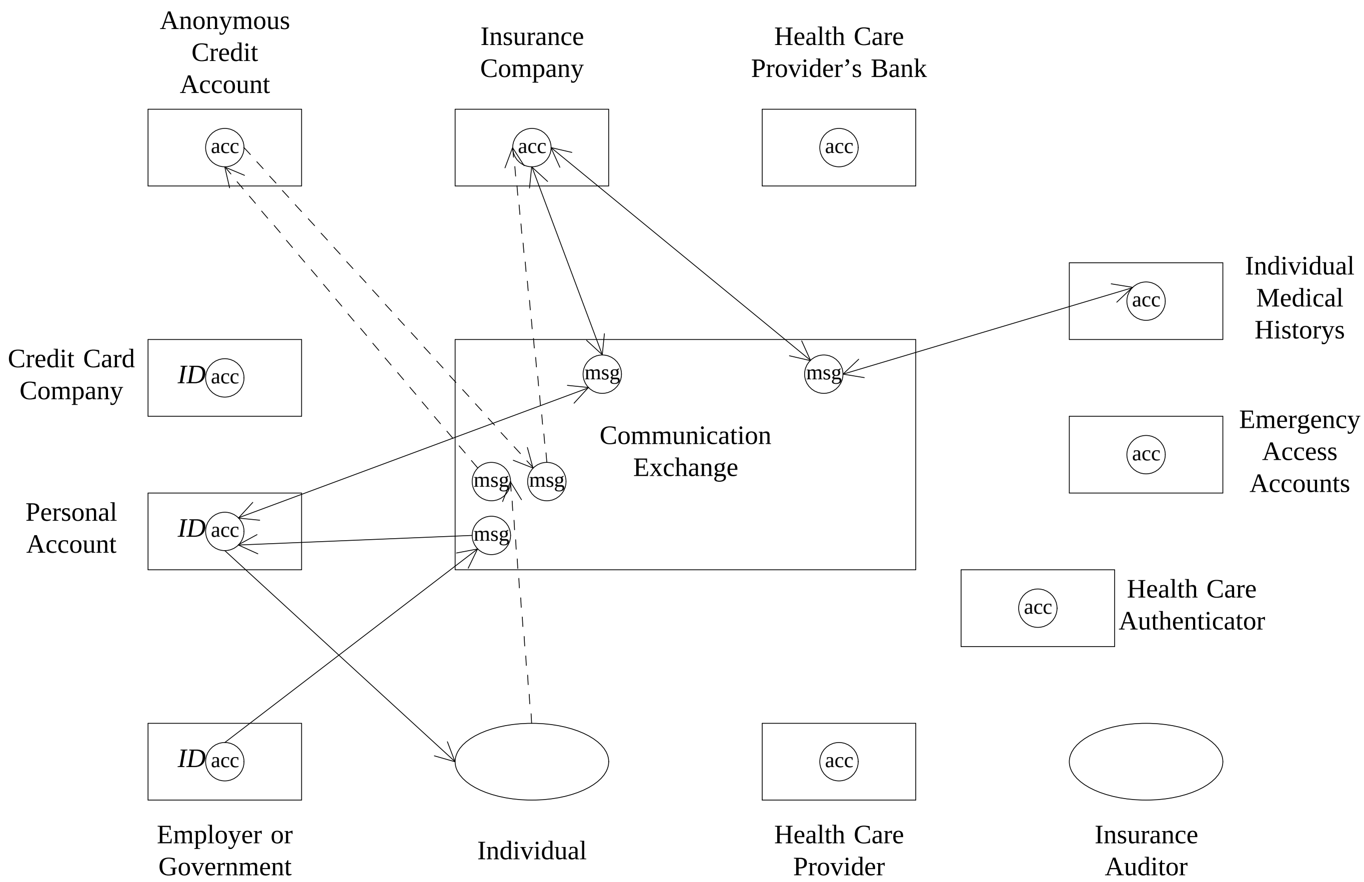

Figure 4. Paying for Insurance Coverage

In an emergency situation, if an individual does not have the smart card that enables him to access his medical history, or if he is unable to use it, his medical history is obtained through an emergency access account. The emergency access accounts associate a biometric identifier with the individual medical historys. A health care provider obtains an individual's biometric identifier, and sends it to an emergency access account. The communication exchange certifies the health care provide, using a digital signature or shared secret, and forwards the message. The emergency access account uses the biometric identifier to find the individual's medical history account and forwards the request for information to the medical history database. These communications are represented by the dashed lines in figure 5. The medical history account returns information to the health care provider through a filter that can restrict the information that is provided during emergencies, at the individual's discretion.

4.3 Paying for Medical Services

After services are rendered, the health care provider places a description of the symptoms and treatment in the individual's medical history account, along with a bill and a double-locked box that can access his bank account. When the bill is passed through the communication exchange, the exchange accesses the health care provider's account in the health care authenticator. The communication exchange verifies that the provider is approved to perform the services and that the charges are within the allowed range. The communication exchange also enters a copy of the procedures and charges in the provider's account, which can be audited. The validation and bookkeeping services performed by the communication exchange are similar to the federal reserve functions that it must perform when transferring money between banks that don't know each other's identity. The health care authenticator can be part of the communication exchange, as is the federal reserve.

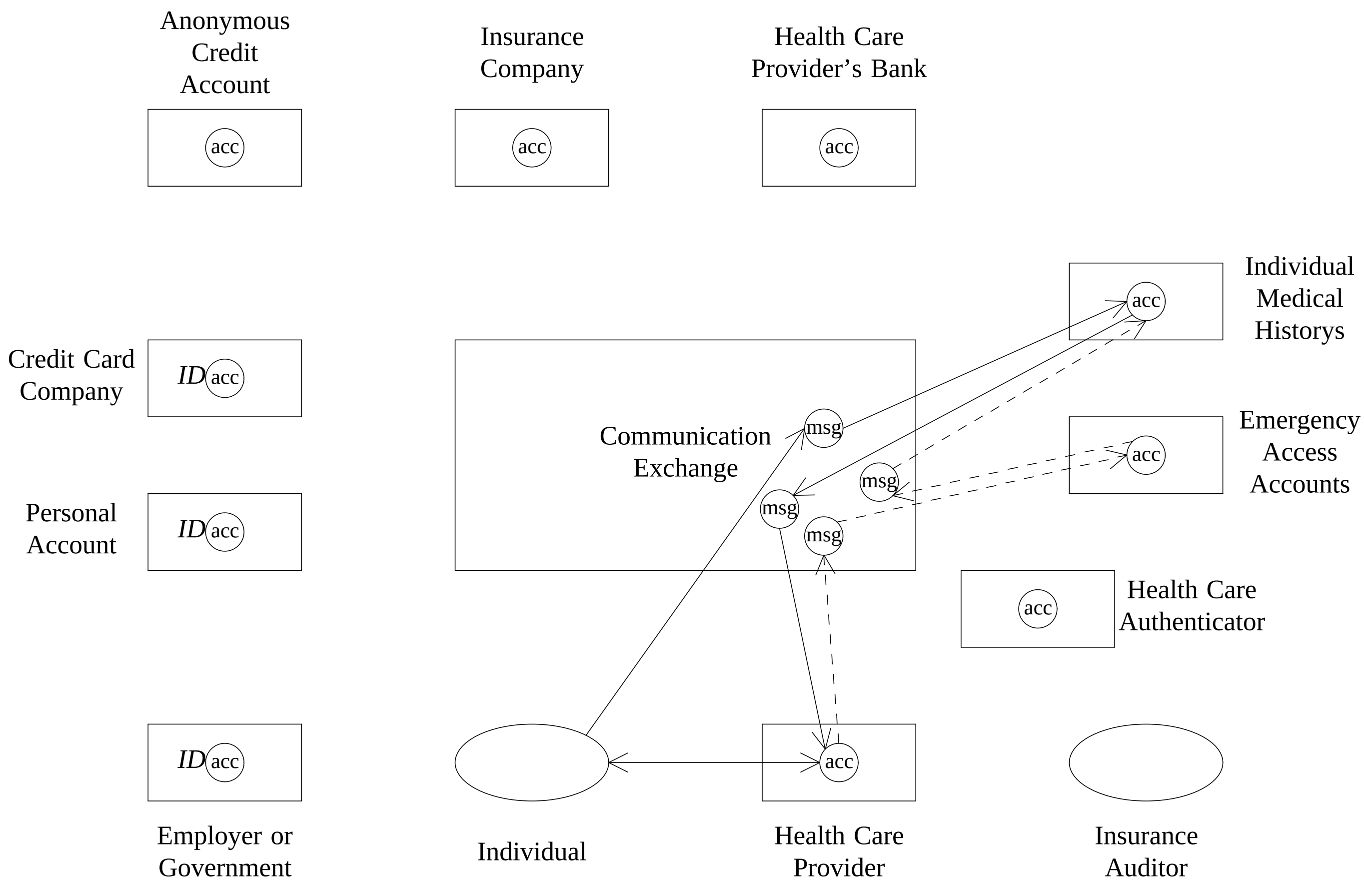

Figure 5. Receiving Medical Treatment

The medical history account presents the bill to the insurance company, using the double-locked box that the individual stored in the medical history account. The insurance company checks the individual's coverage, transfers the covered amount to the health care provider's account, and forwards the bill and a statement of payments to the individual's personal account. A notice is sent to the individual. If there is a balance due, the individual is responsible for contacting his anonymous account and transferring the additional funds. The health care provider's bank informs him that all or part of the payment is made. If payments are not made, the health care provider uses a digital subpoena to identify the patient. The messages are shown in figure 6. As in the previous figures, the solid lines always occur and the dashed lines may or may not occur.

4.4 Collusion Analysis

When two parties collude and have a unique piece of information associated with accounts, they can join the information in those accounts. However, if the no accounts in the agencies have information in common, or if the information that they have in common also appears in many other accounts, then the agencies cannot know which accounts belong to the same individual.

The functions that the parties in the system perform are allocated so that they do not need to operate on the same unique information. Several parties have access to the bills for a medical treatment or insurance coverage, however, it is assumed that in a large system there will be many bills of the same amount. The underlying protocol makes the messages in the system unique, in order to prevent replay attacks. Messages are the unique information that two parties can use to collude. For instance, if an individual's insurance company colludes with the medical history database, there is no common

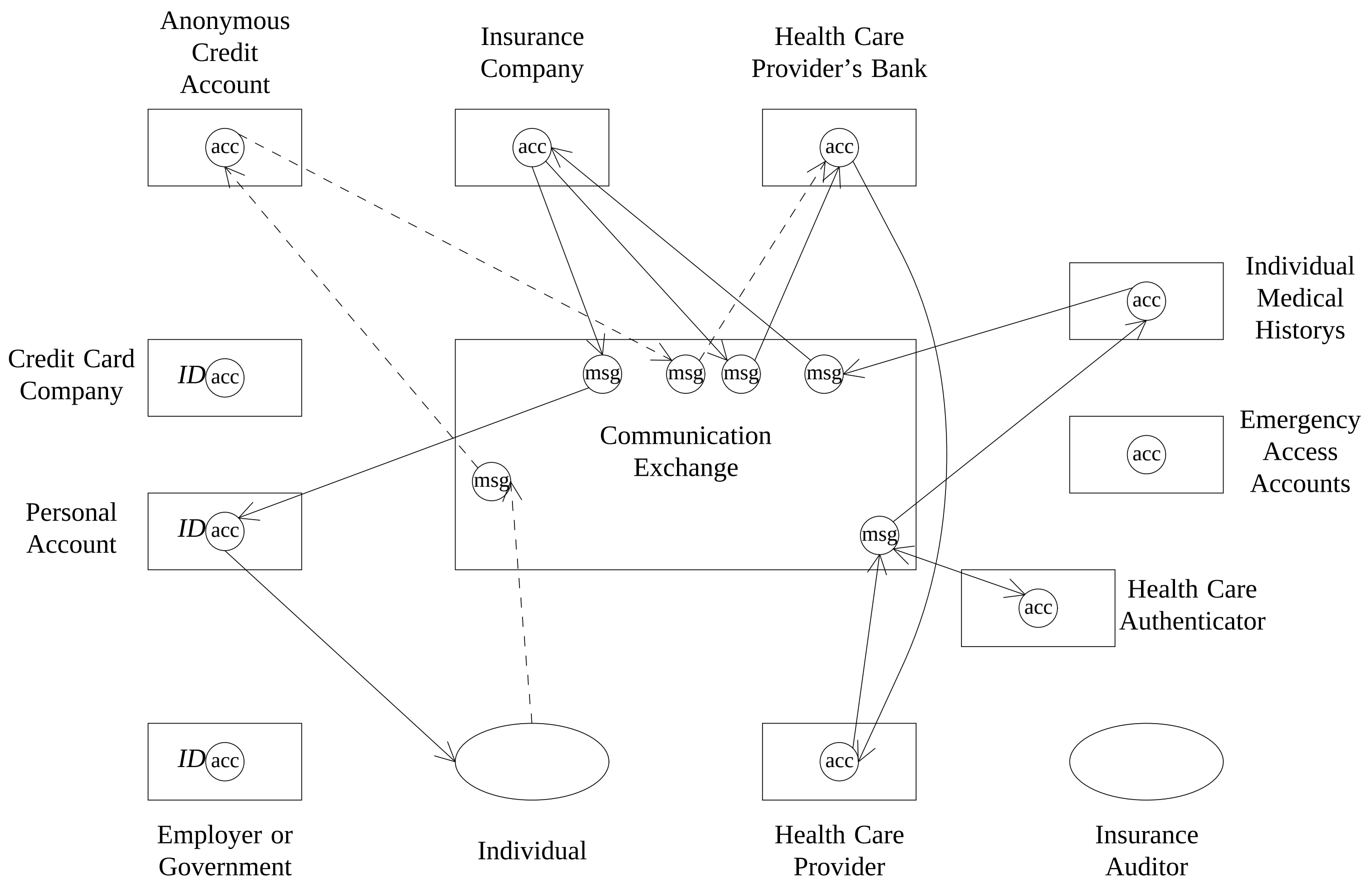

Figure 6. Paying For Medical Services

information in the two accounts belonging to a specific individual, and no way to join the information. However, if the insurance company colludes with the communication exchange, the message that it sent through the communication exchange to the medical history database can be used to determine which medical history database an individual's history is in, and the unique message that was forwarded from the communication exchange to that database. This information can be associated with the individual's account in the insurance company. If the insurance company now colludes with the medical history database, there is a unique message associated with the individual's account in the insurance company and his account in the medical history database, and the information can be joined.

All of the message communications in the system is shown in figure 7. The solid lines are messages that always occur, the dashed lines only exist if the individual makes parts of payments by credit card, or receives emergency treatment. The accounts within organization or banks, and the messages that pass through the communication exchange are the nodes in the collusion graph. The message paths are the links. Information in two accounts can be joined if the parties along any path between the accounts collude. The number of "acc" nodes, including the two with the information, is the number of organizations that must collude, and the number of "msg" nodes is the number of times the communication exchange must be traversed.

The number of times that the communication exchange must be traversed is an important consideration. It is not necessary for all of the messages to pass through the same communication exchange. In order for the system to scale to large numbers of individuals, health care providers, banks, databases, and insurance companies, there must be many communication exchanges. It is also very likely that the communication exchange that acts as a federal reserve for banks, the communication exchange that authenticates and keeps records for health care providers, and the communication exchange that simply

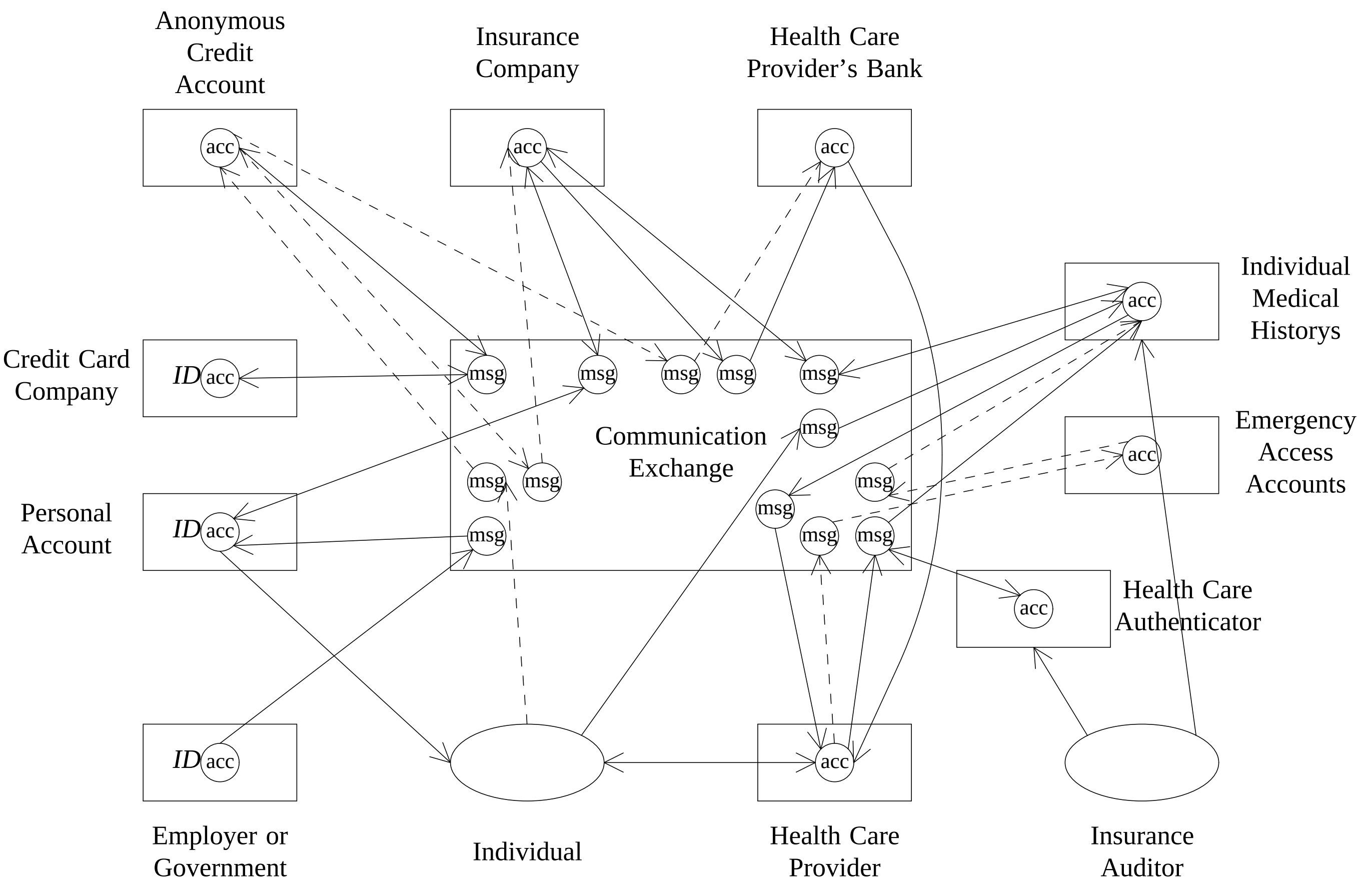

Figure 7. Complete Set of Messages

hides the source and destination of messages will not be co-located. Therefore, if one communication exchange is compromised, all of the message pairs may not become available.

The shortest path between an individual's medical history and his identity includes accounts at three organizations, the medical history database, the insurance company, and the personal account database, as well as two traversals through a communication exchange. In order to join the information three organizations and one or two communication exchanges must collude.

Keeping the number of parties that must collude high reduces the chances of joining information as the parties are compromised. However, there are valid reasons for joining information. For instance:

— When an auditor detects an unusual set of aliments or treatments associated with an individual's account. The auditor can obtain a subpoena which forces the three organizations and communication exchange to trace the message path from the medical history account to the personal account.

— When a health care provider does not receive compensation after several attempts to bill an individual through the medical history database, he can obtain a similar subpoena to determine the individual's identity.

— When the auditor finds an abnormal number of procedures being logged in the health care authenticator, it can locate the health care provider by obtaining a subpoena to force the communication exchange to identify the health care provider. When an individual receives an unexpected bill, he can access his individual medical history, as he does when receiving treatment, to learn the procedure that he is being charged for. If he questions the charge, he can obtain a subpoena to determine the identity of the health care provider that billed him.

If any of the parties following the medical history database fail to respond to a subpoena, and the individual pays for treatment partly with his credit card, then there is an alternate path. The medical history database can follow the path through the health care provider and the banking system to the individual's credit card company. This path may not always exist. If it is likely that parties will be unable or unwilling to respond to a subpoena, than a solid alternate path should be designed into the system.

5. Conclusion

An architecture for a national health insurance plan has been described. The architecture gives an individual the ability to decide what parts of his medical history will be released, under which situations. It does not preclude the possibility of a personal relationship with a health care provider, nor does it stop an individual from making his complete medical history accessible in every situation. It just leaves the decision to the individual.

The architecture improves the availability of medical information by making an individual's medical history accessible wherever he may need it, even when he is incapable of accessing it himself. The centralization of all of the information about a particular individual or a particular health care provider in a single account, in a single location makes it easier to administer and monitor a health insurance plan.

The architecture that has been described is one of a set of architectures that can be applied to this problem, and changes in the requirements and assumptions can be accommodated within this framework. For instance, in this architecture the insurance company that an individual uses is not known to the individual's employer or the government. If the insurance company is controlled by the employer, and the employer knows the individual's identity, it is easier to compromise an individual's privacy because the distance between the individual's medical history account and an account that knows his identity decreases. The distance can be restored, by inserting an independent server between the insurance company and medical history account. The added server need only forward messages, as does the communication exchange.

REFERENCES

[1] N. F. Maxemchuk, S. Low, "Increasing Personal Privacy By Communications," Submitted to Infocom '95

[2] S. Low, N. F. Maxemchuk, S. Paul, "The Anonymous Credit Card and Its Collusion Analysis," Submitted to IEEE Trans. on Networking.

[3] S. Low, N. F. Maxemchuk, S. Paul, "Anonymous Credit Cards," Proceedings of 2nd ACM Conference on Computer-Communications Security, Fairfax, Va. Nov. 2-4, 1994.

[4] N. F. Maxemchuk, "Electronic Publishing," ATT Technical Journal, Sept. 1994, pg 73-80.

[5] J. T. Brassil, S. Low, N. F. Maxemchuk, L. O'Gorman, "Electronic Marking and Identification Techniques to Discourage Document Copying," IEEE Infocom '94, June 14-16, 1994, Toronto, Canada, pg. 1278-87.

[6] D. L. Chaum, "Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms," Proceedings of Communications of the ACM, Vol. 24, No. 2, February 1981, pp. 84-88.

[7] H.-P. Konigs, "Cryptographic Identification Methods for Smart Cards in the Proces of Standardization," IEEE Commun. Mag., June 1991, pp. 42-48.

[8] B. Miller, "Vital Signs of Identification," IEEE spectrum, Feb. 1994, pp. 22-30.