

NOTICE: this is the author's version of a work that was accepted for publication in International Journal of Medical Informatics. Changes resulting from the publishing process, such as peer review, editing, corrections, structural formatting, and other quality control mechanisms may not be reflected in this document. Changes may have been made to this work since it was submitted for publication. A definitive version is published in Int. J. Med. Inform (2007), doi:10.1016/j.ijmedinf.2006.09.015.

Securing electronic health records without impeding the flow of information

Rakesh Agrawal^{a,*}, Christopher Johnson^b

^a *Microsoft Search Labs, 1065 La Avenida, Mountain View, CA 94043, United States*

^b *IBM Almaden Research Center, 650 Harry Road, San Jose, CA 95120, United States*

Abstract

Objective: We present an integrated set of technologies, known as the Hippocratic Database, that enable healthcare enterprises to comply with privacy and security laws without impeding the legitimate management, sharing, and analysis of personal health information.

Approach: The Hippocratic Database approach to securing electronic health records involves (1) active enforcement of fine-grained data disclosure policies using query modification techniques, (2) efficient auditing of past database access to verify compliance with policies and track security breaches, (3) data mining algorithms that preserve privacy by randomizing information at the individual level, (4) de-identification of personal health data using an optimal method of k-anonymization, and (5) information sharing across autonomous data sources using cryptographic protocols.

Conclusions: Our research confirms that policies concerning the disclosure of electronic health records can be reliably and efficiently enforced and audited at the database level. We further demonstrate that advanced data mining and anonymization techniques can be employed to analyze aggregate health records without revealing individual patient identities. Finally, we show that web services and commutative encryption can be used to share sensitive information selectively among autonomous entities without compromising security or privacy.

1. Introduction

The 1995 European Union Directive on Data Protection (“Directive”) [1] set forth stringent cross-industry standards regarding privacy and security of personal data. Pursuant to these standards, EU member states enacted data protection laws that obligate controllers of health data to provide all data subjects with: (1) notice of the

purposes for which they collect and use personal data; (2) choice regarding whether their data may be disclosed to third parties or used for a different purpose than it was originally collected or subsequently authorized; (3) reasonable assurance that the data will be secured and its integrity maintained; (4) access to the data and the opportunity to correct inaccuracies; (5) legal recourse to ensure compliance with data protection requirements. EU states may allow processing of health data without patient consent for purposes of preventative medicine, diagnosis, treatment, management of medical services, or otherwise under professional confidentiality obligations, only if suitable safeguards are provided.

Similar laws in the United States [2], Canada [3], Australia [4], and Japan [5] require healthcare institutions to protect the privacy and security of personal health data. Advisory reports commissioned by the United States government [6,7] also stress the importance of developing secure, interoperable electronic health records systems that preserve patient privacy. As countries around the world transition from paper-based to electronic health records infrastructures, compliance with these data protection laws will require sophisticated information management technologies. Healthcare organizations must implement privacy and security protections that do not unduly constrain proper use and dissemination of health data or inhibit scientific discovery. Technical and policy challenges concerning the widespread adoption of electronic health records systems have been discussed, for example, in [8] and [9].

The Hippocratic Database (“HDB”) [10] is an integrated set of technologies that manages disclosure of electronic health records in compliance with data protection laws without impeding the legitimate flow of information. HDB’s Active Enforcement component limits disclosure of personal health information at a fine-grained level in strict accordance with enterprise policies,

* Work done while author was at IBM Almaden Research Center.

legal regulations, and individual patient choices. Its Compliance Auditing component efficiently tracks past disclosures to verify compliance with these policies. Finally, its data mining, de-identification, and information sharing components enable organizations to derive maximum value from sensitive data without compromising privacy or security.

The remainder of this paper is organized as follows. Sections 2 and 3 describe Active Enforcement and Compliance Auditing. Sections 4 through 6 discuss HDB’s Privacy-Preserving Data Mining, Optimal k-Anonymization, and Sovereign Information Integration components. In each section, we include example scenarios demonstrating practical applications of these technologies. In Section 7, we suggest a number of opportunities for further research in securely managing electronic health records. We conclude in Section 8.

2. Active Enforcement

HDB Active Enforcement (“AE”) [11] is a disclosure management component that is transparent to enterprise applications and agnostic to database systems. It resides in a layer above the database, rewriting user queries to conform to the organization’s data disclosure policies and individual patient choices. AE enforces disclosure policies down to the cell-level in the database, allowing health organizations to comply with detailed requirements of data protection laws without recoding their applications. HDB policy controls are more fine-grained than conventional role-based access controls [12], as they account for the purpose of access, the intended recipient of the information, and patient consent rights, in addition to the user’s access privileges. The complete AE solution

is comprised of three stages—policy creation, preference negotiation, and application data retrieval (see Fig. 1).

In the *policy creation* stage, the healthcare organization specifies a data disclosure policy through the HDB control center. The policy governs the access privileges for each role within the organization according to the category of information sought, the purpose of the request, and the intended recipient of the results. It may also provide individual patients with the opportunity to express opt-in or opt-out choices regarding the disclosure of their personal information, also according to category, purpose, and intended recipient. For example, a patient may opt into sharing his medical information with universities for research purposes, but opt not to share his contact information with pharmaceutical companies for marketing purposes. Policies are expressed in a language such as P3P [13] and installed in the database in a form amenable to symbolic manipulation. The organization may update or replace policies through a one-step installation process in the HDB Control Center. The database stores multiple policies and versions of policies.

In the *preference negotiation* stage, the patient is notified of the health organization’s policies concerning data use and disclosure, advised of any conflicts with his own privacy and security preferences, and allowed to express personal opt-in or opt-out choices. This fully automated process is completed before the patient provides any personal data to the organization. The patient first uses the HDB preference interface to express his preferences concerning the use and disclosure of his personal data. This information is then specified in a preference language [14] and matched with the health organization’s privacy and security policies to identify any conflicts. The patient is advised of these conflicts and

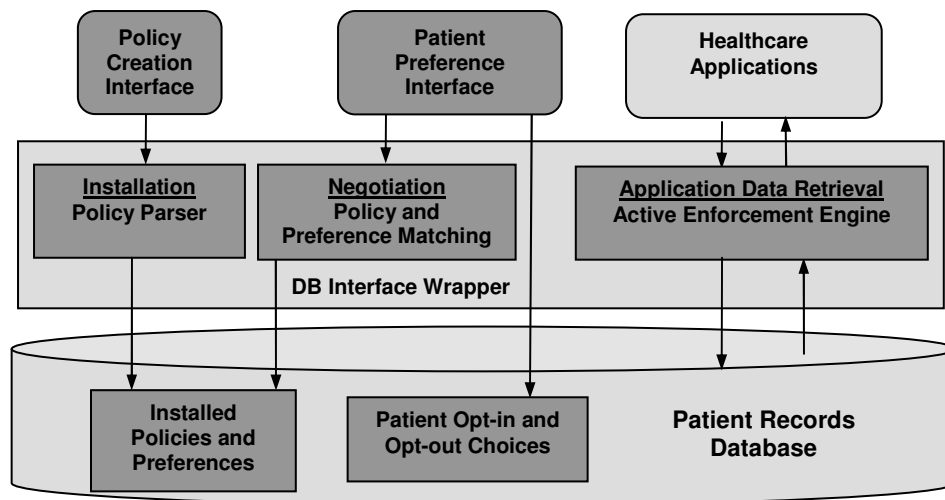


Fig. 1 - HDB Active Enforcement architecture.

given an opportunity to resolve them or terminate the process. Lastly, the patient is provided opt-in or opt-out choices regarding whether his data may be disclosed to third parties or used for a different purpose than for which it was collected. These choices are recorded in the database and factored in at the time of query processing. A successful preference negotiation confirms agreement between the patient and health organization concerning the processing of his personal data.

In the *application data retrieval* stage, the AE engine programmatically modifies all queries to be executed on the data source so that the application only retrieves results that are compliant with the organization's disclosure policies and the patient's opt-in and opt-out choices. The query rewrite process transparently enforces cell-level access controls based upon the user's role, purpose, and intended recipient. This ensures that queries from any application return all responsive data that the particular user is entitled to access, but none that he is not. HDB Active Enforcement can also be configured to support policy rules granting read-only or write access, on a cell-by-cell basis, depending on the context of the query.

AE is integrated into existing environments through a database interface, such as ODBC or JDBC. Its fine-grained enforcement capability implements cell-level restrictions, such as opt-in and opt-out choices, without requiring any changes to enterprise applications. AE is scalable to large databases and actually improves query processing speed in the many instances because rewritten queries benefit from the filtering of sensitive data [11]. It also leverages the existing optimizations and performance enhancements of the database engine. AE can also be combined with techniques that allow queries over encrypted data without significantly degrading performance [15].

2.1. Active enforcement scenario

Richard is a young professional who has recently moved to a new city and would like to select a local healthcare provider and schedule his annual physical. He is considering Continental Hospital, a well-regarded healthcare organization that is part of a large network with many patients and locations. Prior to scheduling an appointment, Richard must register for membership on Continental's website.

2.1.1. Policy creation

As part of its transition to an electronic records infrastructure, Continental has recently installed HDB Active Enforcement software. The first step in the enforcement process is for Continental to create a data disclosure policy. Hospital management starts by reviewing Continental's existing written policy to ensure

it is consistent with current data protection laws and its own business objectives. Amanda, its Chief Privacy Officer, then specifies the policy through the HDB Control Center, which installs the policy in the database in the chosen privacy language.

2.1.2. Preference negotiation

After entering the registration page on Continental's website, Richard is notified of the hospital's data disclosure policy prior to submitting any personal information. He submits a list of preferences regarding the use of his personal data through an extension of his web browser. Among other preferences, Richard indicates that he does not want to share his medical information with government agencies for any purpose and does not want to share his telephone number with third parties for marketing purposes.

HDB automatically compares Richard's preferences with Continental's data protection policy and uncovers one potential conflict. Continental's policy is to release medical information to relevant government agencies if necessary to verify an employee disability claim or comply with a court order. After being notified of this conflict, Richard decides to waive his preference regarding disclosure to government agencies and completes Continental's registration form.

Prior to submitting his completed registration, Richard is provided with two opt-in choices. These choices are intended to allow the patient and healthcare provider to reach an agreement concerning the provider's discretionary use of his personal information, in compliance with national data protection laws. Richard opts to share his medical information with academic institutions for research purposes, but opts not to share his medical information with pharmaceutical companies for research purposes.

2.1.3. Application data retrieval

Several months after his annual physical, Richard injures his ankle playing basketball. His doctor performs a series of X-rays and prescribes a new anti-inflammatory drug to reduce the swelling in Richard's ankle. One year later, a Continental marketing manager queries the hospital database seeking the medical records of all patients who were prescribed the new drug. The pharmaceutical company has offered to purchase this information for follow-up research. In the absence of HDB controls, the manager would see the personal health records of all patients in Continental's database who were prescribed the new drug. However, with HDB Active Enforcement in place, the application returns only the information that patients have consented to share with pharmaceutical companies for research purposes. Accordingly, Richard's

information is filtered from the results. This complies with Continental's data disclosure policy and Richard's privacy preferences.

2.2. Information sharing with active enforcement

A second scenario demonstrates how AE can be used to facilitate policy-compliant information sharing among multiple organizations.

Joan is a professor at Northern University Medical School with access to Continental's patient database under a joint research agreement. She is currently working on a project to evaluate whether various environmental and genetic factors contribute to high cholesterol levels. To begin her research, Joan logs into Northern's web portal and submits the following SQL query to the Continental Hospital database:

```
Select *from patients where total cholesterol ≥ 200
```

Without HDB controls, Joan would be given total access to the records of all patients with total cholesterol levels of 200 and above. This is a violation of Continental's privacy policy and EU data protection laws, because not all patients have consented to reveal their health information to third parties for research purposes. With HDB in place, the AE engine rewrites Joan's query to comply with Continental's data disclosure policy and patient opt-in and opt-out choices. Thus, AE filters out the personal data that patients did not opt to share with third parties for drug research purposes and returns the remaining data that is responsive to the query.

3. Compliance Auditing

Pursuant to the EU Directive and member state laws enacted thereunder, health organizations must be accountable to patients for all processing of their personal data. Upon request, patients are entitled to a description of the data disclosed, the recipients of the data, and the purposes of the processing. Further, member states must provide patients with a remedy for any breach of their rights under these data protection laws. In the United States, the Health Insurance Portability and Accountability Act ("HIPAA") [2] requires healthcare organizations to account for certain disclosures of protected health information upon request and provides penalties for unlawful disclosures. Accordingly, there is a critical need for auditing systems that track past disclosures of information to determine whether they complied with applicable laws and policies.

HDB Compliance Auditing [16] enables organizations to investigate past disclosures without the performance and overhead burdens of other auditing systems. HDB efficiently logs relevant database queries and updates and allows auditors to track the identities of users who have

accessed any cell in the database, the date and time of access, the purpose of the access, the recipient of the information, and the exact information disclosed. Thus, HDB auditing provides reliable and practical means for health organizations to account for its processing of personal information.

The HDB Compliance Auditing system consists of two parts—a logical logging system and an audit application. The logical logging system records all queries and contextual information (i.e., identity, time, purpose, recipient) in query logs. It also stores all data updates, insertions, and deletions in backlog tables, which are populated using database triggers. Alternatively, the backlog tables can be replaced by database replication logs or point-in-time query features.

The audit application provides a simple user interface that allows an auditor to formulate audit queries to specify the exact data she would like to audit. Upon receiving the audit query, the application generates a list of suspicious queries that may have disclosed the specified information. Using the query logs and backlog tables, the application then produces an audit report that identifies the user, time, purpose, recipient, and results returned for each suspicious query (see Fig. 2).

HDB Compliance Auditing is superior to systems that log the actual results of database queries, because it does not incur a cost for read queries or otherwise log redundant data. By logging only the queries and modifications to the database, HDB operates much more efficiently and requires far less overhead than result logging systems. HDB also has a security advantage in that it captures information revealed by a query that may not be reflected in the output. For instance, the query "Select 'yes' if patient 'Richard' has diagnosis 'ankle sprain'" would not be tracked by auditing systems that log the output of queries. The same is true for queries that aggregate values from the records accessed. In contrast, an HDB audit would show the precise information revealed in these situations.

In the following scenario, a healthcare organization uses HDB Compliance Auditing to investigate a claim that it unlawfully disclosed a patient's personal data.

3.1. Compliance Auditing scenario

Palmer is a candidate for political office and a patient of Continental Hospital. Shortly before the election, a local newspaper story discloses portions of Palmer's personal health records indicating that he has been treated for depression. He believes that Continental is responsible for this unlawful disclosure and threatens to sue the hospital under national data protection laws.

Continental's president is very concerned about this high profile accusation and requests that Amanda, the

Chief Privacy Officer, immediately provide him with an accounting of all who have accessed Palmer’s personal health data. He also demands that Amanda conduct a more specific investigation to determine who, if anyone, was actually responsible for the disclosure.

3.1.1. HDB audit specification

Amanda logs into the HDB audit interface to begin the investigation. The system is preset with many common tasks that a hospital auditor might want to perform. Examples of such tasks include “Accounting of Access and Disclosure”, “Who Accessed Medical Information”, “Which Third Parties Accessed Information”, and “Frequent Access of Information.” Alternatively, auditors can declaratively execute custom audits in a SQL-like syntax. They can also specify an exact timeframe for the disclosures they would like to audit.

3.1.2. Investigation of suspicious access

Amanda would first like to know who has accessed Palmer’s medical information within the past year. To accomplish this, Amanda selects the “Accounting of Access and Disclosure” task and restricts her search to only Palmer’s medical information, rather than all of his personal records (e.g., address, telephone number, payment information), and defines the audit timeframe as the past twelve months. The audit application identifies suspicious queries that accessed Palmer’s medical records during the last year and returns a list of users who

accessed them, the time and purpose of each access, and the exact data returned in response to each query. Amanda quickly provides a printed report of this accounting to the hospital president and proceeds with her investigation.

Amanda notices that the results show a large number of queries accessing Palmer’s medical records, but not all of those queries revealed the diagnosis of depression or his prescription for anti-depression medication. Many queries returned information about Palmer’s past diagnoses for influenza and strep throat, but these queries could not have resulted in the improper disclosure. Thus, Amanda adds a custom column to the audit based on diagnosis to sort the information further, so that she can isolate those queries which accessed information about Palmer’s diagnosis of depression. She repeats the same process for the prescription column.

Among the queries that accessed Palmer’s depression diagnosis or treatment, Amanda sorts the results by user. Comparing the user identities with her record of Palmer’s treating physicians, she notes that his primary physicians and nurses frequently accessed medical data relating to his depression. However, another physician, Dr. Roberts, who is not listed as one of Palmer’s treating physicians, also accessed this data several times over a short period, purportedly for treatment purposes.

Amanda is suspicious of this access pattern, so she runs another audit to determine the precise records that Dr. Roberts has accessed within the past year. She notices that Dr. Roberts has made only a few queries

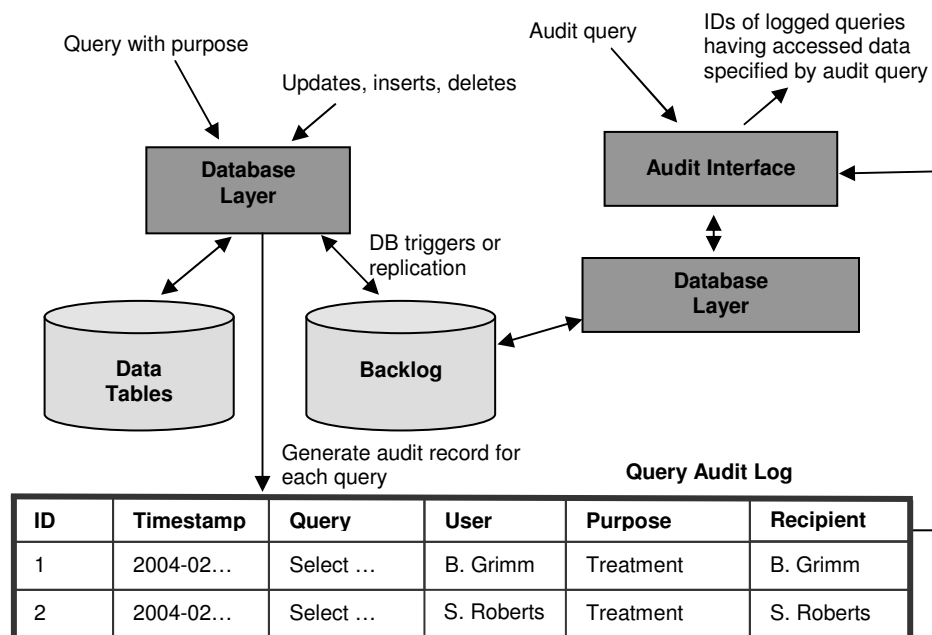


Fig. 2 - HDB Compliance Auditing.

in the system, but has accessed a large number of patient's records, all with diagnoses related to depression. Amanda then requests another audit, this time seeking physician queries that accessed over 200 patient records with a diagnosis of depression. Still, Dr. Roberts is the only physician that has run such a query. Amanda must now interview Dr. Roberts to determine whether she may have been responsible for the leak.

Without a reliable auditing system, Amanda could have spent numerous hours searching through files and notes and interviewing various hospital employees, attempting to locate the actual source of the leak, if any occurred. In contrast, HDB Compliance Auditing allows a hospital auditor to conduct a series of audits, in a matter of minutes, to isolate potential sources of the leak. In fact, Amanda could have reduced the steps above by formulating a more precise initial audit query seeking only queries that accessed Palmer's depression diagnosis or anti-depressant prescription.

An HDB audit may either reveal employee misconduct or indicate that the hospital is not responsible for the disclosure. Moving forward, Amanda can periodically conduct proactive audits to investigate the effectiveness of the hospital's disclosure controls. This type of audit capability also provides a significant deterrent to improper access and disclosure in the future.

4. Privacy-Preserving Data Mining

HDB's Privacy-Preserving Data Mining ("PPDM") [17] allows health organizations to mine aggregate data without revealing individually identifiable information. Thus, it enables analysis of large data sets for epidemiological studies and other medical research without violating patient privacy.

PPDM uses a randomizing function to perturb sensitive values in a patient's record such that they cannot be estimated with reasonable precision. From the randomized data, it reconstructs the original data distribution to allow data mining at the aggregate level, without revealing individual values. Algorithms for building classification models and discovering association rules on top of privacy-preserved data can be applied with only a small loss of accuracy [18].

4.1. Privacy-Preserving Data Mining scenario

Continental recently began a home health monitoring program in which patients measure their vital statistics at home on a daily basis. Scales, blood pressure monitors, cholesterol monitors, thermometers, and other pervasive devices wirelessly feed data into a web application on the patient's home computer that transmits this data to the hospital. The data is fed into patient medical records and used to monitor and diagnose various health conditions.

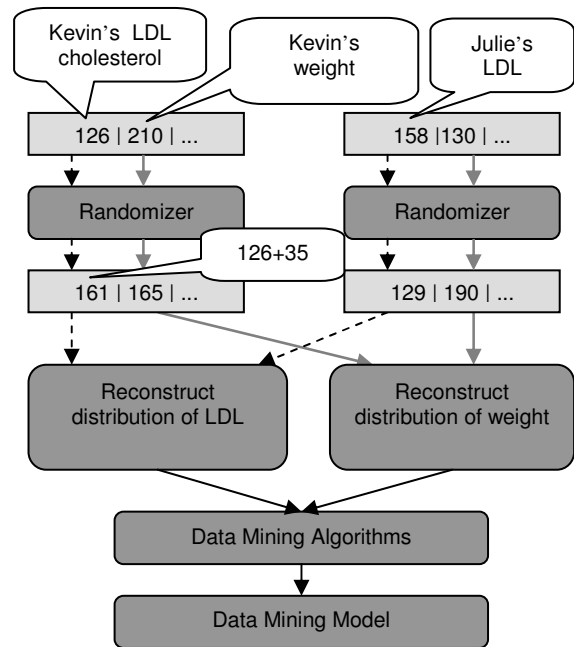


Fig. 3 - HDB Privacy-Preserving Data Mining.

Hospital management recognizes that these large sets of patient health data would also be valuable for research and other data mining purposes. They would like to share this information with third party researchers, on an ongoing basis, without revealing any private information.

Continental decides to solve this problem using PPDM. As patient data is received from the home monitoring system, one copy is sent to the standard patient database and another copy is sent to a PPDM randomizer. Upon receiving each data item, the randomizer perturbs the data and sends it onto a research database consisting of only privacy-preserved data. Continental can provide access to the research database to third party researchers, who can run PPDM algorithms to mine the privacy-preserved data without revealing any individually identifiable information about the patients (see Fig. 3).

5. Optimal k-Anonymization

The EU Directive generally prohibits the processing of personal health data without patient consent, unless required in connection with the provision of medical care. However, member states may allow exemptions to this prohibition if the data are processed under an obligation of professional secrecy or for reasons of substantial public interest, subject to suitable safeguards. In the US, HIPAA allows healthcare organizations to process personal data without patient consent if they remove all possible identifiers or use other statistically and scientifically acceptable methods to de-identify the data.

HDB's Optimal k -Anonymization [19] component offers an optimal method of de-identifying sensitive data that protects patient privacy, but maintains the value of the data for research purposes. In a k -anonymized data set, each record is indistinguishable from at least $k-1$ other records [20]. The process of k -anonymization involves data suppression (deleting cell values or entire tuples) and cell-value generalization (replacing specific values with more general ones). A larger value of k provides greater the privacy protection, but less specific data.

k -anonymization is superior to naive de-identification approaches that simply remove certain identifiers. These naive approaches are prone to linkage attacks that combine the subject data with other publicly available information to re-identify the data subjects. k -anonymization was designed to avoid such linkage attacks, while preserving the integrity of the de-identified data. Unlike anonymization techniques that involve condensation, data scrambling and swapping, or adding noise, the records that remain in a k -anonymized data set are entirely truthful [20].

Since even simple restrictions of optimized k -anonymity are NP-hard [21], Bayardo and Agrawal advanced a new approach [19] that reduces the computational complexity of exploring the array of possible anonymizations. Experiments on a real data set show that “the resulting algorithm can find optimal k -anonymizations under two representative cost measures and a wide range of k ” and “produce good anonymizations in circumstances where the input data or input parameters preclude finding an optimal solution in reasonable time” [19].

Name	Address	City	Age	Diagnosis
Eric	7, rue du Mont Dore	Paris	26	Influenza
Paul	13, rue des Canettes	Paris	42	Hypertension
Marc	48, rue du Four	Paris	47	Diabetes
Henri	21, rue du Mont Dore	Paris	28	Asthma



k -anonymization
($k=2$, on name,
address, city, age)

Name	Address	City	Age	Diagnosis
*	17 th Arrondissement	Paris	20-29	Influenza
*	6 th Arrondissement	Paris	40-49	Hypertension
*	6 th Arrondissement	Paris	40-49	Diabetes
*	17 th Arrondissement	Paris	20-29	Asthma

Fig. 4 - HDB Optimal k -Anonymization.

5.1. Optimal k -Anonymization scenario

Continental Hospital would like to share de-identified data sets with Northern University for medical research purposes. However, removing personal identifiers such as name, street address, telephone number, is insufficient because it leaves the data set prone to data linkage attacks. While no records in the de-identified data set contain a single identifying value, many of them may contain unique value combinations. An individual who is the only Caucasian male born in 1925 living in a sparsely populated area could have his age, race, gender, and zip code joined with a voter registry from the area to obtain his name and mailing address. This would reveal all of the individual's private medical information. However, removing all information that could possibly be used for data linkage attacks would render the data useless for research purposes.

Optimal k -anonymization strikes a balance between protecting the individual privacy and maintaining useful data for analysis. Rather, than categorically removing or revealing columns of information, k -anonymization removes certain cells of data and generalizes others so that every record is indistinguishable from $k-1$ records. In Fig. 4 below, the records in the top table are de-identified such that $k = 2$ with respect to name, address, city, and age. Accordingly, names are suppressed and addresses and ages are generalized to the extent that each record is indistinguishable from at least one other record. The remaining data is truthful and valuable for research.

6. Sovereign Information Integration

HDB's Sovereign Information Integration (“SII”) [22] component enables two or more autonomous entities to run queries across their databases in such a way that the results of the query are revealed, but no other data is exposed among the databases. SII uses a web services infrastructure to apply a set of commutative encryption functions to uniquely identifiable data in different orders and at different locations. The multiply encrypted values are then compared, and the query results provided, without compromising the security of either data set.

Unlike other data integration approaches, such as centralized data warehouses and mediator-based data federations, which reveal all data among the databases, SII only reveals results of the query. This allows collaborating parties to perform a variety of joins and other operations across their databases without revealing unnecessary information. SII is a scalable software solution that can be integrated seamlessly into existing data environments without the need for a trusted third party or any anonymization of the original data. In the following scenario (depicted in Fig. 5), SII presents an

ideal solution to a research problem requiring secure sharing of information among autonomous entities.

6.1. SII clinical genomics scenario

Walter is a medical researcher at Northern University who would like to test his hypotheses concerning correlations between certain genetic expressions and efficacy of a new diabetes drug, Glucotin. Specifically, Walter believes that Glucotin is ineffective in patients with a specific DNA sequence and highly effective in patients with another specific DNA sequence.

To test these hypotheses, Walter must have access to the medical records of patients who are taking Glucotin as well as genetic information about these same patients. Walter is aware that Continental and GeneBank have a number of common patients, many of whom have been prescribed Glucotin. However, national data protection laws prohibit Continental and GeneBank from revealing personally identifiable information without patient consent. Thus, Walter would like to investigate the correlation between the two specific DNA sequences and the efficacy Glucotin, without revealing any other information among the three organizations.

Continental, GeneBank, and Northern have installed SII to facilitate secure, privacy-preserving information sharing. Fig. 5 illustrates the process of Walter’s join operation in the following five steps. (1) To determine whether the first DNA sequence correlates with

ineffective Glucotin treatment, Walter sends an intersection query to the Continental SII service via Northern’s client application. (2) Continental then encrypts the patient table with its own key and sends the table to GeneBank’s SII service. (3) Next, GeneBank encrypts Continental’s singly encrypted patient table and its own DNA table with its own key and sends both tables back to Continental’s SII service. (4) Continental then encrypts GeneBank’s singly encrypted table so that both data sets are now doubly encrypted. (5) Finally, SII joins both doubly encrypted tables and sends the number of matching results to Northern’s application.

7. Research challenges

7.1. Policy specification

Effective HDB Active Enforcement controls rely on the ability of policies to capture the intent of the policy maker accurately. At the same time, the policy specification should be clear enough that the patient can easily understand the policy and the implications of his choices. While privacy policy specification languages such as P3P offer vast improvement over long legal texts of privacy policies and make polices amenable to symbolic manipulation, they fall short on readability and understandability. Thus, there is a major challenge in designing a policy language that reconciles the goals of understandability and efficient computation.

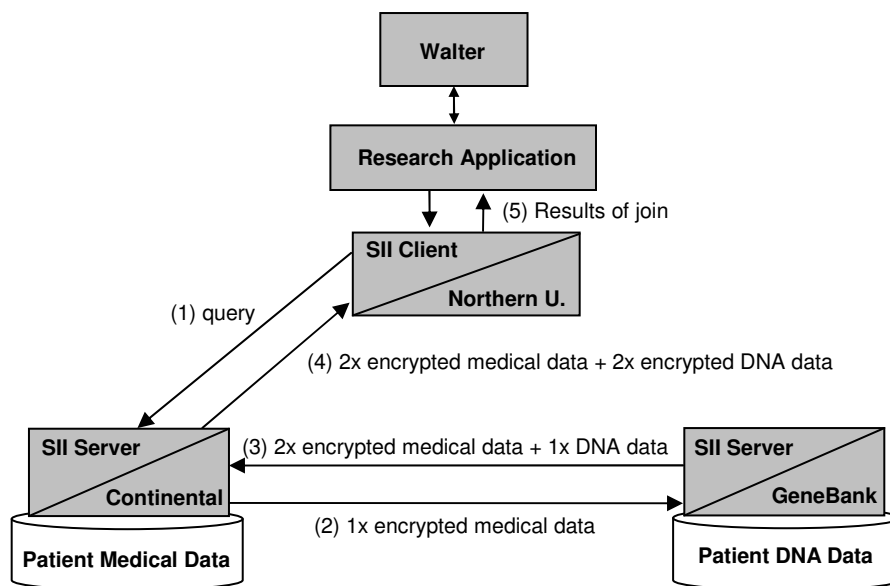


Fig. 5 - HDB Sovereign Information Integration.

7.2. Sticky policies

As healthcare organizations share personal health data with multiple entities, they should be assured that the original policy controls will be enforced over that data as a condition of transfer. When a patient agrees to provide personal data to a healthcare organization under a set of policies and preferences, he is entering into a contract regarding the handling of his data. If the policies allow the data to be transferred to another entity, the patient should be assured that the same disclosure rules will apply to the data after transfer. Thus, it is necessary to have “sticky policies” that transfer with the data and remain with it after consolidation. The transferee should be capable of applying the source disclosure policies to any information in its database. Currently, such data sharing is governed by contracts that require transferees to apply appropriate privacy and security controls to data they receive. Assuming interoperable enforcement systems, sticky policies would be much more effective in ensuring that personal data is always processed in accordance with the patient’s expectations.

7.3. Data pointillism

As electronic health records become more prevalent, patients are likely to have personal health data stored in a variety of distributed data sources. Physicians with assorted specialties may be located in different areas, and patients may change healthcare providers when they relocate, change jobs, or switch insurance companies. To provide physicians with a complete health history for each patient, there is an important need for technologies that unambiguously identify patients and link their information from multiple sources. Such consolidation greatly assists physicians in diagnosis and treatment decisions and reduces the cost of duplicative and unnecessary procedures. Healthcare providers should be able to integrate patient information coherently by combining small, continuously arriving “points” of data. Several techniques exist for this type of data integration [23,24], but further research is needed to accommodate and correct errors in the data, incorporate different data types, and limit false positives. Mechanisms are also needed to enable patients to check the accuracy of their data and make corrections in case errors are found.

7.4. Management of massively distributed data

There are many questions raised by the growing amounts of personal health data stored on inexpensive personal devices such as memory keys, portable disks, and smart cards. In addition, pervasive devices such as wireless monitoring devices are becoming increasingly important for modern healthcare. Accordingly, new technologies are needed to protect the security of the information on these

devices, enable selective sharing of this information, and create back-up mechanisms to prevent data loss.

7.5. User authentication and authorization

Secure access to health information requires mechanisms for accurately identifying those accessing and modifying patient records and ensuring that they have proper authorization. Currently, there are not defined standards for electronic authentication of users and transmitting instant authorizations. To enable information sharing among a network of unaffiliated healthcare organizations, research should define extensible trust hierarchies and authentication standards [6]. Adequate data protection can be assured only if there are accepted and reliable methods for verifying the identities of users accessing sensitive data.

7.6. Data lifecycle management

As electronic health records are stored in databases, technologies that facilitate data life cycle management will become crucial. Data controllers should be able to define retention periods for data based upon legal requirements and patient specifications. At the end of the retention period, storage systems should have methods to remove expired data and forget any persistent data that would allow recreation. Because healthcare organizations require superior availability and reliability of data, storage systems must be secure from data contamination, loss, and leakage and provide methods for establishing the truthfulness of data.

7.7. Interoperability

Another technical challenge facing the healthcare industry is interoperability. Effective sharing of health information requires the ability to communicate among sovereign systems, using standard data formats and clinical vocabularies. While there has been progress toward developing messaging standards such as HL-7, standard vocabularies such as SNOMED-CT, and document standards such as CDA and CCR, much further work remains to be done to ensure that patient health records are complete and healthcare organizations have access to all information necessary for diagnostics, treatment, and medical research [25]. An intriguing research direction worth exploring is the use of mass collaboration [26] to define clinical vocabularies and taxonomies.

Conclusion

We have shown how Hippocratic Database technologies protect the security of personal health records without sacrificing the value of information for diagnosis, treatment, or research purposes. Our example scenarios demonstrate how each of these technologies enables

efficient management, sharing, and processing of sensitive data in compliance with the principles of the EU Directive and other data protection laws. We have also identified a number of significant technical challenges that remain in this area. We hope that the technologies outlined herein serve as a foundation for modern health records infrastructures and inspire productive research in secure information management.

References

-
- [1] European Union Directive on Data Protection, Off. J. Eur. Commun. (1995) 31, No L. 281.
- [2] Health Insurance Portability and Accountability Act of 1996, United States Public Law, pp. 104–191.
- [3] Personal Information Protection and Electronic Documents Act, Second Session, Thirty-sixth Parliament, 48–49 Elizabeth II, 1999–2000, Statutes of Canada 2000.
- [4] Privacy Act of 1988, Commonwealth of Australia, Act No. 119 of 1988 as amended.
- [5] Law on the Protection of Personal Information, promulgated by the Diet of Japan on May 30, 2003.
- [6] President’s Information Technology Advisory Committee, Revolutionizing Health Care through Information Technology, Report to the President of the United States, June 2004.
- [7] Commission on Systemic Interoperability, Ending the Document Game: Connecting and Transforming Your Healthcare through Information Technology, United States Government Printing Office, October 2005.
- [8] B. Humphreys, Electronic health record meets digital library, *J. Am. Med. Inform. Assoc.* 7 (5) (2000) 444–452.
- [9] I. Iakovidis, Towards personal health record: current situation, obstacles and trends in implementation of electronic healthcare record in Europe, *Int. J. Med. Inform.* 52 (1) (1998) 105–115.
- [10] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu, Hippocratic databases, in: Proceedings of the 28th International Conference on Very Large Databases, Hong Kong, China, August 2002.
- [11] K. Lefevre, R. Agrawal, V. Ercegovic, R. Ramakrishnan, Y. Xu, D. DeWitt, Limiting disclosure in Hippocratic databases, in: Proceedings of the 30th International Conference on Very Large Databases, Toronto, Canada, August 2004.
- [12] R. Sandhu, E. Coyne, H. Feinstein, C. Youman, Role-based access control models, *IEEE Comput.* 29 (2) (1996) 38–47.
- [13] L. Cranor, M. Langheinrich, M. Manchiori, M. Presler-Marshall, J. Reagle, Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Recommendation (2002).
- [14] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu, An XPath-based Preference Language for P3P, in: Proceedings of the 12th International World Wide Web Conference, Budapest, Hungary, May 2003.
- [15] R. Agrawal, J. Kiernan, R. Srikant, Y. Xu, Order-preserving encryption for numeric data, in: Proceedings of the ACM SIGMOD Conference on Management of Data, Paris, France, June 2004.
- [16] R. Agrawal, R. Bayardo, C. Faloutsos, J. Kiernan, R. Rantzau, R. Srikant, Auditing compliance with a Hippocratic database, in: Proceedings of the 30th International Conference on Very Large Databases, Toronto, Canada, August 2004.
- [17] R. Agrawal, R. Srikant, Privacy-Preserving Data Mining, in: Proceedings of the ACM SIGMOD Conference on Management of Data, Dallas, Texas, USA, May 2000.
- [18] A. Evfimievski, Randomization in Privacy-Preserving Data Mining, in: Proceedings of the SIGKDD Explorations: Newsletter of the ACM Special Interest Group on Knowledge Discovery and Data Mining, vol. 4 (2), December 2002, pp. 43–48.
- [19] R. Bayardo, R. Agrawal, Data privacy through Optimal k-Anonymization, in: Proceedings of the 21st International Conference on Data Engineering, Tokyo, Japan, April 2005.
- [20] P. Samarati, L. Sweeney, Generalizing data to provide anonymity when disclosing information, in: Proceedings of the 17th ACM SIGMOD-SIGACT-SIGART Symposium on the Principles of Database Systems, vol. 188, 1998.
- [21] H. Lewis, C. Papadimitriou, Elements of the Theory of Computation, 2nd ed., Prentice Hall, 1998, pp. 293–298.
- [22] R. Agrawal, A. Evfimievski, R. Srikant, Information sharing across private databases, in: Proceedings of the ACM SIGMOD Conference on Management of Data, San Diego, California, June 2003.
- [23] O. Benjelloun, H. Garcia-Molina, J. Jonas, Q. Su, J. Widom, Swoosh: a generic approach to entity resolution, Stanford University Technical Report, March 2005.
- [24] S. Ellard, System and method for indexing information about entities from different information sources, United States Patent No. 5,991,758, Issued November 23, 1999.
- [25] California Healthcare Foundation, Clinical Data Standards Explained, November 2004.
- [26] M. Richardson, R. Agrawal, P. Domingos, Trust management for the semantic web, in: Proceedings of the Second International Semantic Web Conference, Sanibel Island, Florida, October 2003.