

Preserving Privacy in Web-based e-Health Systems

Yuan Hong

A Thesis  
in  
The Department  
of  
Computer Science

Presented in Partial Fulfilment of the Requirements  
for the Degree of Master of Computer Science at  
Concordia University  
Montreal, Quebec, Canada

March 2008

© Yuan Hong, 2008



Library and  
Archives Canada

Published Heritage  
Branch

395 Wellington Street  
Ottawa ON K1A 0N4  
Canada

Bibliothèque et  
Archives Canada

Direction du  
Patrimoine de l'édition

395, rue Wellington  
Ottawa ON K1A 0N4  
Canada

*Your file    Votre référence*  
*ISBN: 978-0-494-40940-4*  
*Our file    Notre référence*  
*ISBN: 978-0-494-40940-4*

#### NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

#### AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

---

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.

  
**Canada**

## ABSTRACT

### Preserving Privacy in Web-based e-Health Systems

Yuan Hong

Safeguarding patients' private information is one of the most challenging issues in the design and implementation of modern e-Health systems. A patient's consent should usually be obtained before his/her private information can be disclosed. Recent advances in Hippocratic Databases (HDB) show a promising direction towards the enforcement of privacy policies in e-Health systems. With HDB, patients need to specify their privacy preferences about what data to be disclosed to which recipients for what purposes. However, this may become a daunting task in a complicated application that involves potentially a large number of combinations of data recipients, purposes, and granularities of data. This thesis tackles issues in applying the HDB design to e-Health systems. More specifically, I design an architecture for integrating APPEL preferences with HDB, extend the original HDB design to support fine-grained privacy authorizations demanded by patients and adapt the design to a multi-dimensional model; I propose a series of methods for patients to more conveniently specify their privacy preferences based on the hierarchies naturally existing in each dimension of a privacy preference, define meta-policies to resolve potential conflicts between preferences specified over time, and discuss how to represent such preferences with a snow-flake schema in backend databases. Finally, I illustrate the implementation issues and justify my designs with experimental results.

## ACKNOWLEDGEMENTS

I would like to appreciate lots of people who provide help for this thesis.

First of all, I would like to express my sincerely thanks to my supervisors, Dr. Rachida Dssouli and Dr. Lingyu Wang, for their help, guidance and advices to my study at Concordia University, especially to the published papers and this manuscript.

Meanwhile, I am grateful to the professors in the thesis examining committee, Dr. Olga Ormandjieva, Dr. Nematollaah Shiri and Dr. Bo Zhu. With their valuable comments, I improved this thesis so much.

Finally, I would like to thank all group members in the e-Health research. We work together and implement the system very well as a team.

*To My Parents, With Love*

# Table of Contents

<b>List of Tables</b>	<b>ix</b>
<b>List of Figures</b>	<b>x</b>
<b>1 INTRODUCTION</b>	<b>1</b>
1.1 Introduction on e-Health Systems . . . . .	1
1.2 Privacy in e-Health Systems . . . . .	2
1.3 Objective and Motivation . . . . .	3
1.4 Research Contributions and Organizations . . . . .	5
<b>2 LITERATURE REVIEW</b>	<b>8</b>
2.1 Background . . . . .	8
2.1.1 Service-Oriented Architecture . . . . .	8
2.1.2 HIPAA . . . . .	9
2.1.3 Multi-Dimensional Data Warehouse/OLAP . . . . .	10
2.1.4 Hippocratic Databases . . . . .	11
2.1.5 P3P and APPEL . . . . .	12
2.1.6 Fine-Grained Access Control . . . . .	14
2.2 Related Work . . . . .	15
2.2.1 E-Health and Privacy . . . . .	15
2.2.2 Access Control . . . . .	16
2.2.3 Other Related Work . . . . .	18
<b>3 SCHEMA DESIGN FOR HIPPOCRATIC DATABASES</b>	<b>20</b>
3.1 System Architecture . . . . .	20
3.1.1 Service-Oriented System Architecture . . . . .	20

3.1.2	Preference Negotiation and Data Access . . . . .	23
3.1.3	Electronic Patient Records Database Schema . . . . .	26
3.2	Fine-Grained Authorization in HDB . . . . .	28
3.3	Multi-Dimensional Hippocratic Databases . . . . .	32
<b>4</b>	<b>HIERARCHICAL SPECIFICATION OF PRIVACY PREFERENCES</b>	<b>36</b>
4.1	Hierarchies in Privacy Preferences . . . . .	36
4.2	Specifying Preferences Using Hierarchies . . . . .	39
4.2.1	Option 1 . . . . .	40
4.2.2	Option 2 . . . . .	41
4.2.3	Option 3 . . . . .	42
4.3	Conflict Resolution . . . . .	42
4.3.1	Latest Take Precedence . . . . .	44
4.3.2	Disclosure Take Precedence . . . . .	45
4.3.3	Denial Take Precedence . . . . .	45
4.4	Implementation Option . . . . .	46
4.4.1	Using HDB Metadata . . . . .	47
4.4.2	Using Snowflake Metadata . . . . .	48
<b>5</b>	<b>IMPLEMENTATION</b>	<b>54</b>
5.1	Implementation Environment . . . . .	54
5.1.1	SOA with BEA Weblogic 8.1.2 . . . . .	54
5.1.2	Oracle 9i Database Support . . . . .	57
5.2	Hippocratic Databases Schemata Implementation . . . . .	58
5.2.1	Oracle Schemata Implementation . . . . .	59
5.2.2	Experiments . . . . .	59
5.3	Preferences Negotiation Implementation . . . . .	61
5.3.1	Generating Disclosure Set . . . . .	62
5.3.2	Meta-Policy for Conflict Resolution . . . . .	63
5.3.3	APPEL Integration . . . . .	65
5.4	A Scenario for Electronic Patient Record Data Access . . . . .	67

<b>6 CONCLUSION AND FUTURE WORK</b>	<b>70</b>
6.1 Conclusion . . . . .	71
6.2 Future Work . . . . .	71
<b>Bibliography</b>	<b>73</b>
<b>Glossary</b>	<b>81</b>



# List of Tables

3.1	Schemata of the Original HDB Design . . . . .	28
3.2	Relational Schemata for Fine-Grained Authorizations . . . . .	30
3.3	Signed Attribute <i>Authorized-Users</i> . . . . .	32
3.4	HDB in Star Schema . . . . .	33
3.5	Schema for Multi-Dimensional HDB . . . . .	33
3.6	Access Control for Dimensional Tables . . . . .	34
4.1	Authorization under Original HDB Design . . . . .	47
4.2	Snowflake Schema for EPR Data and Metadata . . . . .	49
4.3	Redesigned Schema for Hierarchical Authorizations . . . . .	50
4.4	Authorization under the Redesigned Schema . . . . .	50
5.1	Web Service Deployment Platforms . . . . .	56
5.2	Main Security Features of Database Products . . . . .	57
5.3	Storage Requirements for Fine-Grained Authorization . . . . .	60
5.4	Original HDB Preferences and Hierarchical Preferences . . . . .	61

# List of Figures

3.1	A Service-Oriented Architecture for e-Health Portal . . . . .	21
3.2	Preference Negotiation and Data Access Architecture . . . . .	23
3.3	Define Policies and Make Preferences . . . . .	24
3.4	Star Schema for Electronic Patient Records . . . . .	27
4.1	Hierarchy on Private Data . . . . .	37
4.2	Hierarchy on Data Recipients . . . . .	38
4.3	Hierarchy on Purposes . . . . .	39
4.4	Hasse Diagram for Hierarchy on Private Data . . . . .	40
4.5	Disclosure Inheritance for Option 2 . . . . .	41
4.6	Disclosure Set Example for Option 2 . . . . .	42
4.7	Disclosure Inheritance for Option 3 . . . . .	43
4.8	Disclosure Set Example for Option 3 . . . . .	43
4.9	Snowflake Metadata for Hierarchical Authorizations . . . . .	52
5.1	The Hierarchy Used in Implementation . . . . .	62
5.2	Specifying Preferences in Web Interface . . . . .	63
5.3	Privacy-Authorization Table before Specifying Preferences . . . . .	64
5.4	Privacy-Authorization Table after Specifying Preferences . . . . .	64
5.5	Detailed Preference Negotiation Architecture . . . . .	66
5.6	Web Interface for EPR Service to Use Private Data . . . . .	67
5.7	Electronic Patient Records Data Access Based on Authorization . . . . .	68

# Chapter 1

## INTRODUCTION

### 1.1 Introduction on e-Health Systems

“e-Health is an emerging field in the intersection of medical informatics, public health and business, referring to health services and information delivered or enhanced through the Internet and related technologies” [Eys01]

The healthcare industry is undergoing fundamental changes. Examples of such changes include a shift from hospital-centric services to a more ambulatory system (with home care, day care clinics, and so on) and the treatment of chronic diseases that actively involves the patient himself/herself [Bis05]. The emergence of Web-based e-Health portals is a natural result of such changes because such portals provide patients and healthcare professionals easy accesses to information no matter where they are. According to a recent survey, most patients say they are very interested in and capable of accessing healthcare information and services via a Web-based portal system [Eys01]. However, my research work is based on the privacy-preserving issues in such a Web-based e-Health systems.

Furthermore, e-Health can encompass a range of services which are related to medicine or healthcare and information technology: Electronic medical records (also called Electronic patient records), Telemedicine, Evidence based medicine, Consumer health informatics, Health knowledge management, Virtual healthcare team and Medical research uses e-Health grids

[Wik]. Among those research, Electronic Patient Records (EPR) enables healthcare professionals to establish easy communication of patient data and various accesses on patients' personal records. My research work concentrates on those processes and provides several solutions on how to preserve privacy in such processes.

## 1.2 Privacy in e-Health Systems

Electronic healthcare systems are playing a critical role in today's medical organizations. How to safeguard patients' private information is an important and challenging issue faced by the designer and administrator of e-Health systems. The privacy issue is critical to such systems because most medical data are about individual patients and highly sensitive [AJ07]. Inappropriate disclosures of those data cause privacy breaches to patients, which in turn lead to serious legal and financial consequences to the organization. At the same time, the privacy issue is particularly challenging in e-Health systems due to the usually complex design and implementation of such systems. There exist standards and solutions for addressing the privacy issue in general-purpose applications. The Platform for Privacy Preferences (P3P) developed by the World Wide Web Consortium (W3C) allows users and websites to declare privacy preferences and policies in a machine-readable format [Con]. On the other hand, the Hippocratic Databases is a framework for enforcing privacy policies based on database technologies [AKSX02a]. The integration of P3P and HDB is a natural solution for preserving privacy in e-Health systems, which forms the basis of this research work.

Meanwhile, the privacy issue in online systems has drawn considerable attentions lately. Users may abandon an electronic service if they feel being requested for too much private

information. This fact calls for a more flexible way for users to specify their privacy preferences rather than simply accepting or rejecting a fixed set of policies determined by the service provider. For example, in an e-Health system, patients should be allowed to control aspects like *“which parts of their personal data will be disclosed”*, *“to whom the data will be disclosed”*, and *“for what purposes the data will be disclosed”* according to privacy legislation and regulations [10496]. However, specifying such preferences becomes a daunting task in a large and complicated application where personal data may exist at different granularity levels, and such data may potentially be disclosed to a large number of recipients and for many different purposes.

### 1.3 Objective and Motivation

The main objective of this research is to address following two pertinent issues in the collection phase of patients’ private data in an e-Health portal system:

- Demonstrate a model for representing privacy requirements of patients’ data based on relational data model and multi-dimensional data model. The solution will extend the well known Hippocratic Database (HDB) model to support fine-grained access control and multi-dimensional database applications.
- Propose a method to simplify the tasks of patients in specifying privacy requirements. The method will take advantage of hierarchies naturally existing in various dimensions of privacy requirements.

The main motivation of this research comes from the unique requirements and specifications of e-Health systems. As described in Section 1.1, e-Health is a process to provide healthcare via electronic approaches, especially over the Internet. E-Health systems share

lots of characteristics with other systems, such as e-Commerce systems or other e-Services related systems. These characteristics include general security features (e.g. SSL), system architecture, communication protocols etc. However, e-Health systems also exhibit some distinguishable features, in particular with respect to preserving privacy, that are rather different from general e-Services systems:

- First, the data in e-Health systems are more sensitive than data in general e-Commerce systems. This is evidenced by the fact that e-Health systems are usually under more scrutinies from governments and organizations through regulations and laws, such as HIPAA.
- Second, in contrast to general e-Commerce systems, e-Health systems require more non-traditional data types, such as multimedia data (for example, Electrocardiogram (ECG) or Magnetic Resonance Imaging (MRI)). The optimization of database schema thus becomes critical since redundancy in such multimedia data can lead to significant overhead in storage.
- Thirdly, e-Health systems have a broader range of services than general e-Commerce systems. In e-Commerce systems, most applications are based on similar processes (such as online transactions), whereas applications in e-Health systems cover online operations on patients' records, tele-consultant, tele-monitoring, and so on. This diversified collection of services will significantly complicate the security and privacy requirements.
- Finally, privacy requirements in e-Health pose more technical challenges than in general e-Commerce systems. In e-Commerce systems, personal data are generally submitted only during the registration time. In e-Health systems, diagnosis results and other

sensitive information are constantly collected during actual services. These information may have very different requirements in terms of privacy. Fine-grained control must be supported to meet patients' diversified needs.

Based on the above analysis of the comparison between e-Health systems and other e-Services systems, this thesis presents several research outcomes on how to preserve privacy in a web-based e-Health systems. More specifically, it addresses some new privacy-preserving approaches to solve problems in analytical data warehouses. And therefore, my research work also concentrates on a privacy-preserving model for e-Health systems with multi-dimensional medical data warehouses, in particular over the web-based e-Health systems, which includes a group of privacy-preserving approaches in such a model.

## 1.4 Research Contributions and Organizations

In this thesis, we tackle several issues around the integration of P3P and HDB technologies in e-Health systems, and propose a hierarchical approach to address these issues.

- First, I design an architecture for integrating P3P preferences with HDB policies. Patients' preferences specified in APPEL (P3P's Language for Privacy Preferences) are mapped to privacy metadata tables stored in HDB. Doctors requesting for private data are authorized against the privacy metadata.
- Second, we study how to support fine-grained privacy authorizations in HDB. The normally used attribute-level access control turns out to be insufficient for e-Health systems. We provide a solution based on a redesigned schema.
- Third, we extend HDB to support the multi-dimensional model. The original design of HDB based on relational database model is not suitable for multidimensional models

used in medical data warehouses.

- Moreover, the hierarchical approach is based on the observation that hierarchies naturally exist in each of the three dimensions of privacy preferences, namely, private data, recipients and purposes. Specifically, private data can be disclosed in different levels of generalization, to recipients with different privileges, and for purposes dominating each other. We propose a series of methods for patients to more easily specify their privacy preferences by taking advantage of the hierarchies. We then study meta-policies for resolving potential conflicts that may arise between preferences specified over time.
- Meanwhile, we discuss how to encode the specified privacy preferences in backend databases and provide a snowflake schema-based design for simplifying the representation of hierarchical authorizations.
- Finally, implementation issues and some experimental results are described.

The rest of this thesis is organized as follows: Chapter 2 reviews the literature; Chapter 3 discusses issues in multi-dimensional database schema using Hippocratic Databases; Chapter 4 proposes a hierarchical approach to the specification of privacy preferences and the related issues of such an approach; Chapter 5 discusses implementation issues. Finally, Chapter 6 concludes the thesis.

Furthermore, here are some publications related to this thesis:

#### Refereed Conference Papers

1. Yuan Hong, Shuo Lu, Qian Liu, Lingyu Wang and Rachida Dssouli, "A Hierarchical Approach to the Specification of Privacy Preferences," *Proc. 4th IEEE International Conference on Innovations in Information Technology (Innovations'07)*, Nov. 18-20, 2007, pages



660-664.

2. Shuo Lu, Yuan Hong, Qian Liu, Lingyu Wang and Rachida Dssouli, "Access Control in e-Health Portal Systems," *Proc. 4th IEEE International Conference on Innovations in Information Technology (Innovations'07)*, Nov. 18-20, 2007, pages 88-92.

3. Qian Liu, Shuo Lu, Yuan Hong, Lingyu Wang and Rachida Dssouli, "Securing Telehealth Applications in a Web-based e-Health Portal," *Proc. ARES 2008 - International Conference on Availability, Reliability and Security*, IEEE Press, To appear.

#### CIISE Technical Reports

4. Yuan Hong, Shuo Lu, Qian Liu, Lingyu Wang and Rachida Dssouli, "Preserving Privacy in e-Health Systems Using Hippocratic Databases".

5. Shuo Lu, Yuan Hong, Qian Liu, Lingyu Wang and Rachida Dssouli, "The Design and Implementation of a Secure Web-based e-Health Portal".

# Chapter 2

## LITERATURE REVIEW

In this chapter, necessary background related to my research is introduced in Section 2.1.

Related work is reviewed in Section 2.2.

### 2.1 Background

#### 2.1.1 Service-Oriented Architecture

A Service-oriented Architecture (SOA), is a software architecture that uses loosely coupled software services to support the requirements of business processes and software users. Resources on a network in an SOA environment are made available as independent services that can be accessed without knowledge of their underlying platform implementation. SOA can be implemented using one of the protocols (REST, RPC, DCOM, CORBA and Web Services) and might use, for example, a file system mechanism to communicate data to a defined interface specification between processes conforming to the SOA concept. SOA can also be regarded as a style of information system architecture that enables the creation of applications that are built by combining loosely coupled and interoperable services. These services interoperate based on a formal definition (e.g. WSDL) that is independent of the underlying platform and programming language.

The main drivers for SOA adoption are that they link computational resources and promotes their reuse. Enterprise architects believe that SOA can help businesses respond more quickly and cost-effectively to changing market conditions. Such an architecture promotes reuse at the macro level (service) rather than micro level (objects). It can also simplify interconnection to - and use of - existing IT (legacy) assets. Hereby, according to the macro viewpoint of an object oriented system, service level partition would make the use of resources much more effective.

### **2.1.2 HIPAA**

After too many anguished victims let their issues be known on the health privacy protection, the lawmakers finally noticed. In 1996 the US Congress addressed these issues head on, and acted to pass a landmark law to eliminate all those problems by setting specific mandates in all aspects of the transactions between healthcare companies, providers, and carriers. This law is known as Public Law 104-191, the Kennedy-Kassenbaum Bill, more popularly known as the Health Insurance Portability and Accountability Act (HIPAA). [NB03]

**Main Significations** of HIPAA are as follows:

- **Patient's Access to Information**

In order to eliminate the potential problem and errors due to unavailable records, the law states that patients can access their information any time and in a standard format. This information should be immediately available to the patient and, on request, to parties acting on their behalf, such as their new doctors and other healthcare professionals.

- **Standardized Information Exchange**

HIPAA also mandates that the information related to health insurance must be exchanged in a standard, predefined way.

- Privacy of Information

HIPAA places a great deal of emphasis on the privacy issue that affects all of us in some way or other. Medical records are considered private and should be protected, just like any other tangible property, such as money.

The law mandates that organizations must establish a clear security policy that can be verifiable and, more importantly, auditable.

### **2.1.3 Multi-Dimensional Data Warehouse/OLAP**

A data warehouse is the main repository of an organization's historical data. The most important factor leading to the use of data warehouse is that data analysts can perform complex queries and analysis on the information without slowing down the operating systems.

Data warehouses are the basis of important analytical applications, such as Online Analytical Processing (OLAP). OLAP is an approach to quickly provide answers to analytical queries that are multi-dimensional in nature. OLAP is part of a broader category of business intelligence, which also includes extract, transform and load (ETL), relational reporting and data mining. The typical applications of OLAP are in business reporting for sales, marketing, management reporting, business process management (BPM), budgeting and forecasting, financial reporting and similar areas.

Data warehouses configured for OLAP employ a multi-dimensional data model, allowing for complex analytical and ad-hoc queries with a rapid execution time. N. Pendse has suggested an alternative and perhaps more descriptive term to describe the concept of OLAP: Fast Analysis of Shared Multi-dimensional Information (FASMI) [Pen]. They borrow aspects

of navigational databases and hierarchical databases that are faster than their relational kin.

#### 2.1.4 Hippocratic Databases

The Hippocratic Database (“HDB”) is a 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, legal regulations, and individual patient choices. These technologies [AJ07] include (1) active enforcement of fine-grained data disclosure policies, (2) efficient auditing of past database access to verify compliance with policies, (3) privacy-preserving data mining, (4) de-identification of personal data using an optimal method of k-anonymity, and (5) secure information sharing among autonomous data sources.

Hippocratic databases include a group of technologies for securing electronic health records [AJ07], such as privacy metadata, compliance auditing [ABF<sup>+</sup>04], privacy preserving data mining [AS00], optimal k-anonymization [BA05], .etc.

The privacy metadata in the Strawman architecture of Hippocratic database [AKSX02a] use *Purpose* as a special identical attribute for authorizations in each table. The purpose will be used for authorizations and policies in other two tables of the metadata called Privacy-Authorization Table (PAT) and Privacy-Policies Table (PPT).

In PAT, four attributes are used for privacy authorization:

- Purpose: the purposes for the authorized users to use the private data
- Table: the table which the private attribute belongs to
- Attribute: the private attribute in the original database schema

- Authorized-users: the users authorized to access the data

Meanwhile, in PPT, there exist five attributes for policy recording:

- Purpose: the same as in PAT
- Table: the same as in PAT
- Attribute: the same as in PAT
- External-recipients: the actors to whom the data will be disclosed
- Retention: the period during which the data item should be maintained

The privacy metadata in the architecture will be used to store corresponding special data for attribute access control.

### **2.1.5 P3P and APPEL**

The Platform for Privacy Preferences (P3P) is a standard for assisting web applications to collect users' data and corresponding disclosure elements in a machine-readable XML format, and the specification defines the base data schema for data to be collected together with a standard set of uses, recipients, data categories and the association of policies and websites and so on.

In the P3P specification [CLM<sup>+</sup>02b], two types of standards are defined: Privacy Policies which describes how to use the collected personal data by the websites (enterprises or medical organizations) and Privacy Preferences which will be undergone a conformance checking against the policies when recipients plan to access data relevant to the preferences.

P3P has provided a standard language to express users' preferences on their personal data, which is A P3P Preference Exchange Language (APPEL) [CLM02a]. APPEL files

specify which data to be disclosed, with what purpose, and to whom, and it is not designed to be read by end users but for applications which request policies to use the data. The architectures for negotiating preferences and checking preferences will be presented later in this thesis.

Privacy preferences will be expressed in APPEL rules [CLM02a]. Users can describe their preferences in a set of preference rules with APPEL. When recipients request the system to access the disclosed data in the presence of some policies, whether to accept the request for data access will be determined by those preferences rules in APPEL and the requested policies in P3P. Since the set of rules are bound to those policies, the contents of the APPEL specifications include some elements of the P3P policies as well. There are two components in a rule: Rule Behavior which specifies the triggered actions once a rule takes effect and Rule Body which provides the pattern which is matched against a policy. For the behaviors, when the requested policy conforms to the preferences, the request will be acceptable and the behavior can be valued *“Request”* to accept the request policy. However, if the requested policy does not comply with the preferences, the request will be rejected and the behavior can be specified *“Block”* to produce such actions. The requests could be somewhat acceptable if the behavior is *“Limited”*.

The connective attribute within APPEL could define some logical operations for the preferences. There are six APPEL connectives: or, and, non-or, non-and, or-exact, and-exact. Two unusual connectives: or-exact and and-exact should be indicated. Or-exact matches if at least one of the contained preference expressions exist in the policy and no other preference expressions contained in that part of the policy, while and-exact matches if all of the contained preference expressions exist at the correct positions in the policy and no other preference expressions included in that part of the policy.

### 2.1.6 Fine-Grained Access Control

Privacy legislation, such as the Health Insurance Portability and Accountability Act (HIPAA) [10496], usually requires accesses to be granted to only users with appropriate privileges. A typical requirement is that a patient's private data can only be accessed by authorized persons for intended purposes during a given time period. For example, a doctor can only access data of those patients who are under his/her treatment [10496]. Such fine-grained access control (FGAC) is traditionally handled at the application layer using filtering predicates to modify database queries. This approach is not scalable for complex applications and its security relies on the correct coding by application developers who are usually not familiar with security requirements [Kyt99]. In case of third party canned applications, the approach becomes infeasible since the source code is normally out of the control of the medical organization.

A common approach to FGAC supported by most modern database management systems is the view-based security. That is to allow users to access only certain predefined views instead of the tables. For example, Oracle's implementation of FGAC, known as Row Level Security (RLS) or Virtual Private Database (VPD), allows policies to be attached to tables and triggered by accesses [Kyt99]; Sybase's similar solution, named Row Level Access Control [Ince], allows users to define access rules that apply restrictions when retrieving data from the system; Microsoft SQL Server primarily supports traditional view-based access control with an optional feature called row level permissions usable only with table hierarchies [Incd]; IBM DB2 supports FGAC through a combined mechanism based on views, triggers and special registers [Incb].



## 2.2 Related Work

### 2.2.1 E-Health and Privacy

E-Health systems have received significant attentions in recent years. Over thousands of health care sites exist online, and a large number of them are web-based. A survey of research topics and trends on e-Health systems can be found in [KG06]. It reported that most of current research is in Clinical Information Management (CIM) and Electronic Patient Record (EPR), and research on Consumer Health Informatics (CHI) and Bio Medical Cognitive Science (BCS) need to be developed, while efforts on Telemedicine (TEL) and Medical Language Processing (MLP) are considerable. Protecting patients' privacy is a mandatory requirement in most e-Health systems according to privacy legislation and regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) [10496]. Threats to patients' privacy may arise from many aspects of a medical organization. For example, published medical data can lead to attacks on patients' privacy even though the data are sanitized. The concept of k-anonymity requires identifying attributes to be generalized such that any real-world individual can be linked to at least k records in the published data, which is considered a tolerable privacy threat [SS98].

Since most of the private information in e-Health systems are recorded in Electronic Patient Records (EPR), we can initiate the research from EPR to find out those privacy issues that we are concerned about. Many EPR research establish scientific basis or new technologies to improve the performance of clinical information systems. [BMA<sup>+</sup>01] supports viewing of medical imaging exams and graphically-rich clinical reporting tools and conforms to the latest emerging XML standard for digital documents. However, my research provides security and privacy enforcement on EPR in the whole e-Health system, but not only improving

the presentations of EPR system.

Section 1.3 has illustrated the distinguishable features of e-Health system, compared with other e-Services system. Most of the differences result from the sensitive data (privacy) and security requirements. Not only the sensitive data but also the data recipients roles and authorization purposes in the e-Health systems are more complicated than those items corresponding to other systems. Some contributions of my research work are based on this feature for e-Health systems, such as the hierarchical preferences (sensitive data, data recipients and authorization purposes). My research work also presents some improvements on specifying preferences based on the details of those characteristics of e-Health systems to enforce privacy protection. Some research also include some work on hierarchies in e-Services systems. In [MMZ], the authors propose a flexible framework for preference negotiation between customers and enterprises that adopt different processes for the same service. The framework can help users to use the service with minimal disclosures of their private data. In [IY05], a hierarchical approach is presented to pose the least number of questions to users about their privacy preferences. This can be considered as “asking” users for their privacy preferences, whereas our approach lets users to “tell” the preferences. Furthermore, Role-Based Access Control (RBAC) [SCFY96] [San95] supports hierarchical roles management in all kinds of systems.

### **2.2.2 Access Control**

Privacy threats within a medical organization may come from unauthorized accesses to sensitive data. A basic requirement found in most privacy regulations is that accesses to patient records should only be granted to users with appropriate privileges, for intended purposes, and during a given time period [10496]. Such a requirement for fine-grained access

control (FGAC) can be handled by the application layer or by database systems through view-based security. For example in subsection 2.1.6, Oracle’s implementation of FGAC, known as Virtual Private Database (VPD), allows policies to be attached to tables and triggered by accesses [Ky+99]. Other popular commercial products like Sybase, Microsoft SQL Server, and IBM DB2, all have different degree of support for FGAC.

Traditional policy-dependent access control models include discretionary access control (DAC) in which accesses are assigned by the owner of objects, mandatory access control (MAC) in which accesses are based on security clearances and levels. The controls in DAC are discretionary that an object with access permissions is able to pass the permissions onto other objects (unless restrained by MAC), while the security policies are centrally controlled by policy administrators (users cannot override the policy) with MAC. Policy-neutral access control models and mechanisms also exist, such as the role-based access control (RBAC) that leverages roles to simplify access control administration [SCFY96] and the Flexible Authorization Framework (FAF) that uses meta-policies for conflict resolution and policy derivation [JSSS01]. [LHL<sup>+</sup>07] addresses a two-tier approach to control access for e-Health portals, which supplements existing RBAC capabilities with a rule-based access control module based on the classical FAF model. With RBAC in e-Health systems, roles are defined based on their different responsibilities, and we can simple control the permissions on various roles, especially those data recipients in the privacy model.

This thesis focuses on the modeling and representation of privacy requirements in backend databases. In terms of access control models such as RBAC, our focus is on internal details of the permission component of such models. That is, how can we devise the permissions to satisfy patients’ privacy requirements. Our work can thus be seamlessly integrated with most existing access control model, including RBAC. As to our contribution of hierarchical

specification of privacy requirements, it shares similar ideas with the role hierarchies in RBAC, but these are for completely different purposes.

### 2.2.3 Other Related Work

As mentioned in Section 2.1, the Platform for Privacy Preferences (P3P) is a standard for encoding a user's privacy preferences and an organization's privacy policies in a machine-readable format, such that a user's browser can interact with the organization's website to determine whether the former's privacy preferences matches the latter's privacy policies [Con]. P3P provides a standard language for specifying privacy preferences about disclosing private data, namely, A P3P Preference Exchange Language (APPEL) [CLM02a]. APPEL developed by the World Wide Web Consortium (W3C) [CLM02a] allows a user's privacy preferences to be specified in a machine-readable format. Tools exist for creating APPEL specifications, such as the Java-based editor developed by the European Union's Joint Research Center (JRC) [Cen]. Using the tool, an APPEL rule can be created either by selecting from predefined rules or by using the advanced mode that allows creating new rules. My work can certainly employ APPEL as a standard interface between the client and server, although it focuses on how to simplify a specification using hierarchies. [BCK03] presented a software to query a set of web sites for P3P policies, check the validity of each policy, and analyze the information practices it describes. Meanwhile, totally 588 P3P-enabled web sites have been analyzed by such a software and they presented the first major analysis of the data practices of P3P-enabled web sites. Furthermore, [AJSX04] discussed how to implement P3P, which enables web users to gain control over the private information using database technology. Finally, [JZY<sup>+</sup>06] proposed a privacy policy of P3P by which current relational database management systems can be transformed into their privacy-preserving equivalents.

At the backend, Hippocratic Databases (HDB) is designed for preserving privacy in database applications [AKSX02a], which is a promising solution to managing privacy policies and preferences using database technology [AJ07] [AKSX02a]. Privacy-authorization tables in an HDB encode privacy preferences using attributes named “Authorized Users” or “Purpose”. The “Active Enforcement” (“AE”) is a framework for deploying an HDB [Incc], which supports a preference negotiation stage that ensures any access to private data adheres to privacy policies. My implementation is based on HDB and the hierarchical approach complements the existing negotiation stage found in the HDB AE. Furthermore, some other research also focus on the extension of Hippocratic Database model, such as extending HDB for XML databases with tree-like hierarchies [LWHS06] compared to existing HDB privacy metadata which integrates privacy protection into relational databases. The database schemata and tables in my implementation are multi-dimensional based. Related to data warehouse/OLAP, [SLXN04] proposed an effective solution to preserve confidential information while still providing range query accuracy in OLAP data cube. [KKST97] defined a security model for data warehouse which describes security constraints for roles in the data warehouse. In 1996, Anderson proposed a security policy model for clinical information systems [And96]. The model applies traditional confidentiality and integrity models, such as the BLP model and Clark-Wilson model, to medical data. In contrast to this high-level policy model, Chapter 4 of this thesis will address a specific issue in privacy preference specification.

## Chapter 3

# SCHEMA DESIGN FOR HIPPOCRATIC DATABASES

In this chapter, a service-oriented system architecture and its components for privacy protection will be introduced in Section 3.1. Section 3.2 will describe how to integrate finer-grained authorizations using Hippocratic databases and how to establish a multi-dimensional HDB schema designed for Electronic Patient Records (EPR) with fine-grained authorizations support.

### 3.1 System Architecture

#### 3.1.1 Service-Oriented System Architecture

Our e-Health system includes a portal and various services provided through the portal. We adopt a service-oriented architecture for the e-health portal. Traditional designs of software system usually have difficulties in meeting the functional requirements. The interoperability among heterogeneous components requires a complicated integration process, which then implies unacceptable downtime and efforts for the integration. In contrast, the service oriented architecture (SOA) exposes resources of each software component as standard-conforming services that can be accessed without understanding the underlying implementation details.

The main components of our design are portals interconnected via the Web Services for

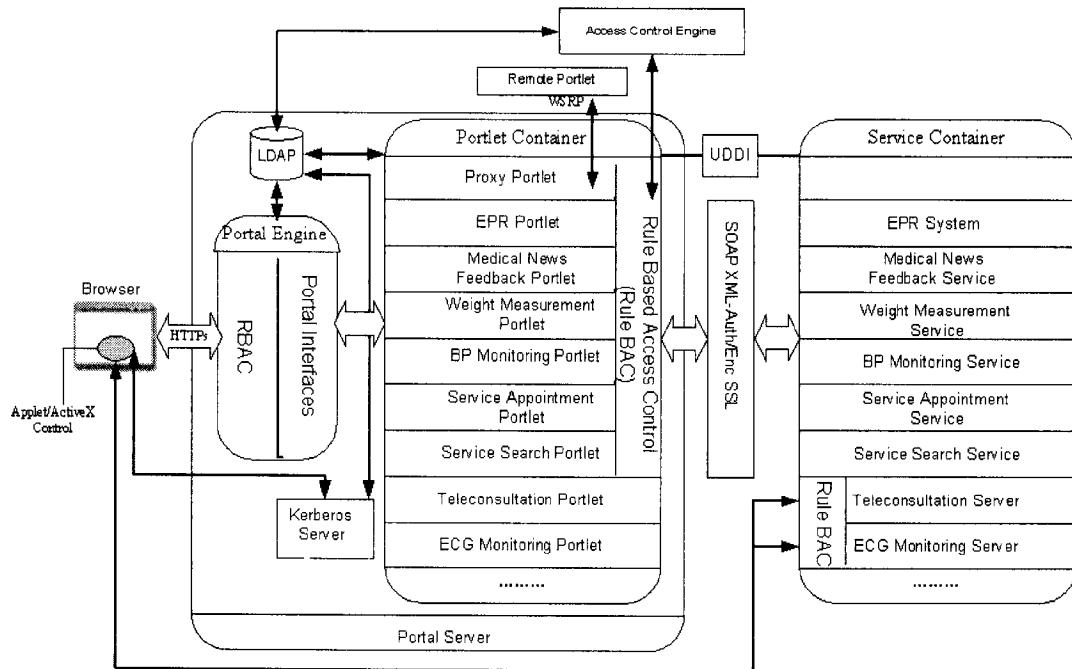


Figure 3.1: A Service-Oriented Architecture for e-Health Portal

Remote Portlets (WSRP) protocol. Each portal is a Web-based application that provides users a unified interface to all the services provided by multiple organizations. The portal allows its users to personalize desired content through creating customized webpage, namely, portal interface. Each portal interface also plays the role of a content aggregator by including a collection of related services. Navigation elements inside each portal interface allow users to easily navigate among different collections of services just like surfing the Web. In our design, all client-side processing is supported either by the built-in functionalities of a standard Web browser or through applets that are automatically downloaded and installed into the browser. Users are not required to possess sophisticated computer skills in order to use services provided by the portal.

More details of the proposed architecture are depicted in Figure 3.1. The e-Health portal

comprises the presentation layer of the architecture, which is decoupled from the implementation layer at service providers to allow independent development of services and to ease the integration of such services. The separation of portal interfaces from portlets allows portal administrators to easily customize the source of services using a content management system. Interoperability and extensibility are both inherent to the architecture because the portlets communicate with backend services through standard Simple Object Access Protocol (SOAP). As an exception, real-time services are allowed to bypass the portal since they usually demand better performance than SOAP can provide. To access a service, a user connects to the portal server via a standard Web browser. According to the user's personalized settings, a portal interface is displayed with a collection of portlets inside it. When the user clicks on a button encapsulated in a portlet, the corresponding action of the service will be performed at backend service providers (which may be governed by a remote portal). For real-time services that bypass the portal, an applet downloaded to the browser will perform the required actions on behalf of the users.

Meanwhile, the privacy protection components include preference negotiation, Electronic Patient Record (EPR) services for data users, database and metadata schema and so on. In the whole e-Health system, service container stores some EPR services that need to access the private data in the backend database. Patients make preferences through the process of preference negotiation which is integrated into the web interface of the architecture (portlets), while the private data and metadata schemata exist separately from the components in Figure 3.1. At the time users (such as doctors) want to access those data by some EPR services, my privacy protection engine will be triggered and the access control supported by privacy metadata will be invoked to limit users in accessing those privacy. The following two subsections will introduce the details of the privacy protection components within a specific



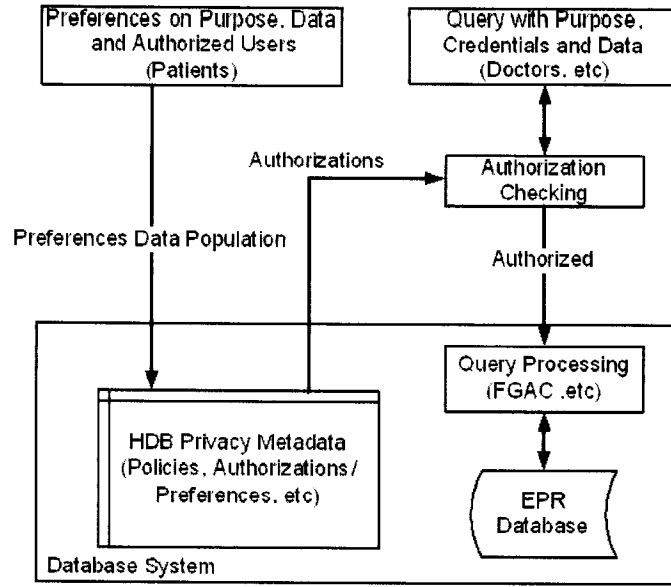


Figure 3.2: Preference Negotiation and Data Access Architecture

architecture.

### 3.1.2 Preference Negotiation and Data Access

The architecture of the privacy-protection subsystem of an e-Health system is illustrated in Figure 3.2. I consider two types of users accessing the e-Health system. First, patients state their opt-in and opt-out preferences through a web interface. Second, doctors need to access the personal information of their patients for specific purposes (such as *“Treatment”* and *“Prescription”*). Notice here the patient and doctor only refer to their roles in either providing or requesting the private data. Other users of the e-Health system, such as a nurse, may be considered as a patient, a doctor, or both depending on their roles with respect to the private data.

In Figure 3.2, patients specify their preferences about disclosing private data in the AP-PEL language through a web interface. The preferences are checked against the mandatory

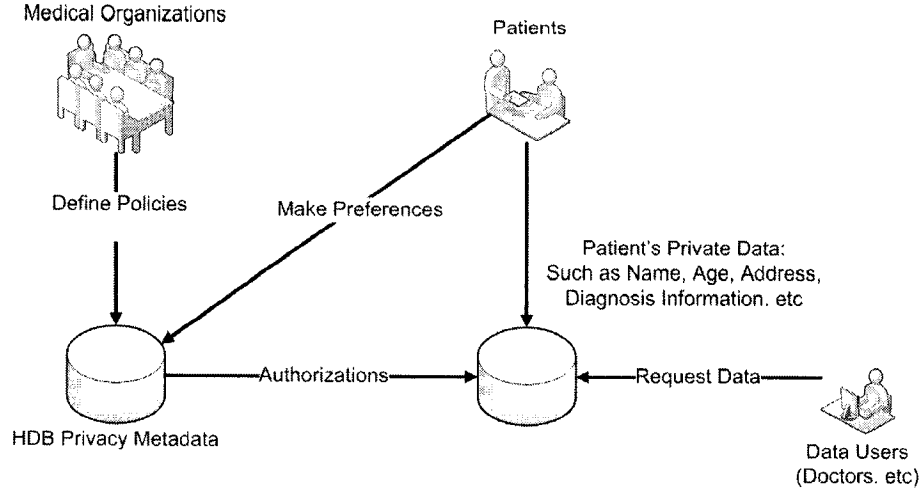


Figure 3.3: Define Policies and Make Preferences

part of P3P policies for conformance between the client's browser and the web server (this can be implemented in either a client-centric or server-centric way [AKSX02b] [AJSX04]). If the preferences match the policies, the preferences are mapped to and stored in the attributes and records of privacy metadata tables in backend HDB for later references. When doctors request for accesses to private data, the application will provide authentication credentials and associated purposes together with queries. Based on the purposes and the requested resources, such as records and attributes, the system checks corresponding metadata and determines whether the doctor is a legitimate recipient of the requested private data. If the doctor haven't been authorized before the checking, the private data would be prohibited for the services requested by the doctor.

Figure 3.3 shows a more detailed presentation on how to define policies and how to specify preferences in my privacy-preserving model. Since my model adopts the metadata of Hippocratic Databases to provide authorizations on EPR private data access, there exist two different kinds of metadata, privacy policies and privacy preferences. Policies are defined by

the hospital officials or some medical organizations. They define them to standardize limited disclosure for data use and preference negotiation in the e-Health systems. It is applicable to every role of the system, and all the authorizations should conform to the privacy policy of that system. More specifically, patient make preferences for their specific authorizations. For example, policy defines that patients should disclose their diagnosis information for treatment purpose, whereas preference will be specified to let one patient Bob disclose his diagnosis information for treatment purpose. Policy is a general concept, while preference is more specific. When the patients specify preferences, the conformance checking between the specified preferences and privacy policies will be executed, and then, the preferences will be recorded into the metadata if they are consistent. On the other hand, when data users (e.g. Doctors) initiate accesses on the EPR data (e.g. execute a query on the private data), the authentication process will be triggered. The process can be provided by Java Programming or DB Programming (triggers or procedures). Moreover, all those processes in preference negotiation architecture and limited disclosure for data use (including the policy and preference conformance checking, query processing, conflict resolution for preferences and disclosure sets generating that will be introduced in Chapter 4) can be built as middleware between users and Database Management System (DBMS), as shown in Figure 3.2 and 5.5.

Notice that our first contribution is an effort in extending the HDB model to address its limitations in supporting fine-grained access control and multi-dimensional data model. Like the original HDB model, our work focuses on the modeling of privacy requirements using either the relational model or the multi-dimensional model. The implementation of those models can be supported by existing processing techniques. However, optimization is certainly desired for improving the efficiency of privacy-enabled database applications, as

indicated in [AJ07]. Such optimization of is beyond the scope of this thesis. Our second contribution is part of the Web-based e-Health portal, and it focuses on the interaction with users in collecting data. This requires additional development efforts (in our implementation we used Java). The collected data will be stored in backend databases according to the model in our first contribution. No additional implementation efforts are thus required in the database.

### 3.1.3 Electronic Patient Records Database Schema

Our e-Health system adopts web interfaces (with APPEL integrated) to specify preferences and employs privacy metadata of Hippocratic databases to record preferences for patients. The database system includes schemata for privacy metadata and Electronic Patient Record (EPR) private data, which forms one of the most important concept in this research. Since EPR data will be used for many purposes, such as “*Research*” or “*Statistics*”, an analytical data warehouse may be required to improve the performance of analytical queries. Our research focuses on multi-dimensional medical records and studies how to preserve privacy in such a model. Therefore, a star schema for electronic patient records is given in Figure 3.4 based on the requirements and characteristics of private medical data warehouses.

As shown in Figure 3.4, a fact table (EPR) records the dimensional attributes linked to other four dimensional tables, measurement attributes such as “*Diagnosis*”, “*Prescription*”, .etc, and some other degenerate attributes. Besides the fact table, other dimensional tables provide the specific information of Patient, Doctor, Date and Department for analytical usage. For demo purpose, I have shown only a part of the attributes in each table. In reality, other useful information may be required. Some tables in this simple schema will be referred to in later paragraphs.

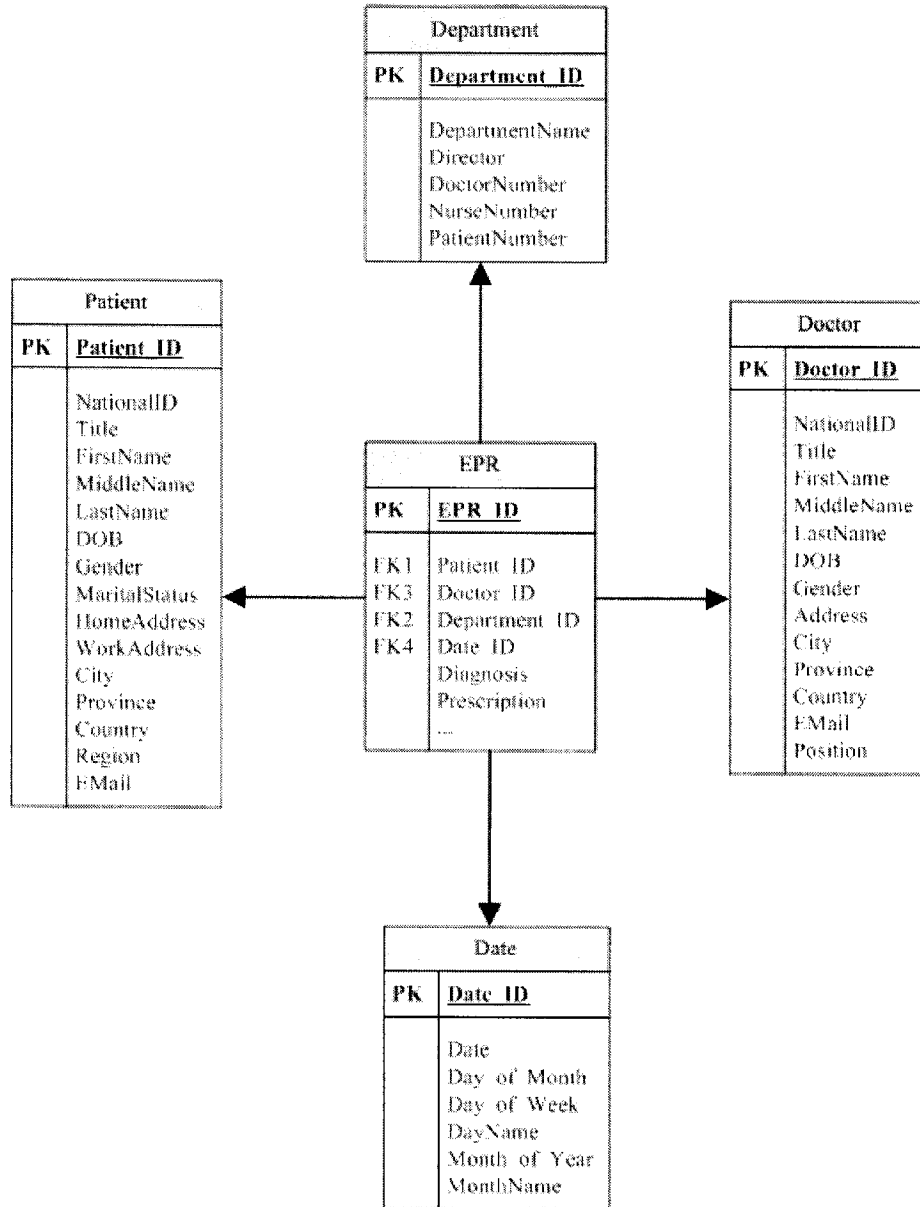


Figure 3.4: Star Schema for Electronic Patient Records

Table 3.1: Schemata of the Original HDB Design

HDB based on Relational Database Schema			
Table	Attributes		
<i>EPR</i>	EPR-ID, Patient, Diagnosis, Purpose		
<i>Patient</i>	Patient, Gender, Age, Address, City, Purpose		

Authorizations in HDB Privacy Metadata	
Table	Attributes
<i>Privacy-Authorization</i>	Purpose, Table, Attribute, Authorized-Users

Table <i>Privacy-Authorization</i>			
<i>Purpose</i>	<i>Table</i>	<i>Attribute</i>	<i>Authorized-Users</i>
Treatment	EPR	Diagnosis	Doctor
Treatment	Patient	Age	Analyst
Statistics	EPR	Diagnosis	Analyst

### 3.2 Fine-Grained Authorization in HDB

R. Agrawal et. al proposed Hippocratic databases for preserving privacy using privacy metadata[AKSX02a]. Table 3.1 shows schemata of the original HDB design when applied to an e-Health system. Table *Patient* and *EPR* store private data about patients' information and electronic records. Table *Privacy-Authorization* records which users are authorized to access each combination of table, attribute and purpose (I shall only include a few attributes for simplicity). In this example, attribute “*Diagnosis*” of table *EPR* may be accessible to doctors for treatment purpose and to analysts for statistics purpose, whereas attribute “*Age*” of table *Patient* is only for treatment purpose.

However, the above design has a limitation in that it does not support fine-grained authorizations. In a typical e-Commerce website, a transaction will not even begin unless the user's preferences match all of the website's policies. Each attribute can thus have a fixed set of authorized users no matter in which record it appears. This all-or-none approach is not suitable for medical organizations where, for example, a patient should not be refused of

treatment simply because he/she disagrees to provide age or gender for statistics purposes. Any patient should be allowed to opt out optional privacy policies or opt in with conditions. For example, Bob may decide that his diagnosis information should be disclosed for statistics purposes only if the symptom is *Flu*, and the information should be kept private for all other symptoms. As another example, he may choose to allow or disallow a specific doctor in accessing his diagnosis information for certain symptoms. Another patient Eve may have a completely different preference about these private data. None of those requirements can be represented in table *Privacy-Authorization* in Table 3.1.

Table 3.2 shows my redesigned schemata for supporting fine-grained access control when the protected data are recorded in a data warehouse. I add an additional attribute “*Authorization-ID*” to the tables and a foreign key constraint on the attributes from table *EPR* to table *Patient*. In table *EPR*, for symptom *Flu*, Bob is willing to disclose his diagnosis information for both statistics and treatment purposes and his name for treatment purpose but only to Alice. For the symptom of *Diabetes*, Bob does not want to disclose any information for any purpose. Also for symptom *Flu*, Eve has a preference different from Bob’s, which is to disclose the diagnosis data to doctors. In addition, the fine-grained authorization can be used in other tables, such as *Patient*. We can see that both of Bob and Eve authorized the access on attributes (“*Age*”, “*Gender*” .etc) for each tuples in the patient information table.

In summary, my redesigned schemata provide all necessary fine-grained authorization described in the above examples. Other special authorizations will be illustrated later on.

The above design of table *Privacy-Authorization* assumes a closed policy. That is, unless a user explicitly appears in attribute *Authorized-Users* in table *Privacy-Authorization*, the user will by default be prohibited from accessing the attribute. Such a closed policy [SdCdV04] is inconvenient when a patient chooses to prohibit certain users while allowing all others to

Table 3.2: Relational Schemata for Fine-Grained Authorizations

HDB Schemata	
Table	Attributes
<i>EPR</i>	EPR-ID, Patient, Diagnosis, Purpose, Authorization-ID
<i>Patient</i>	Patient, Gender, Age, Address, City, Purpose, Authorization-ID

HDB Privacy Metadata Schema	
Table	Attributes
<i>Privacy-Authorization</i>	Authorization-ID, Purpose, Table, Attribute, Authorized-Users

Table <i>EPR</i>				
<i>EPR-ID</i>	<i>Patient</i>	<i>Diagnosis</i>	<i>Purpose</i>	<i>Authorization-ID</i>
000001	Bob	Flu	Treatment	1,2
000001	Bob	Flu	Statistics	3
000002	Bob	Diabetes		
000003	Eve	Flu	Treatment	1

Table <i>Patient</i>					
<i>Patient</i>	<i>Gender</i>	<i>Age</i>	...	<i>Purpose</i>	<i>Authorization-ID</i>
Bob	Male	28	...	Treatment	4
Bob	Male	28	...	Research	5,6
Eve	Female	23	...	Treatment	4

Table <i>Privacy-Authorization</i>				
<i>Authorization-ID</i>	<i>Purpose</i>	<i>Table</i>	<i>Attribute</i>	<i>Authorized-Users</i>
1	Treatment	EPR	Diagnosis	Doctor
2	Treatment	EPR	Patient	Analyst
3	Statistics	EPR	Diagnosis	Analyst
4	Treatment	Patient	Age	Doctor
5	Research	Patient	Age	Researcher
6	Research	Patient	Gender	Researcher



access his/her private data. For example, Bob does not want Alice to see his name but does not care about other doctors accessing that data. It would be prohibitive to explicitly list all other doctors' names in attribute "*Authorized-Users*".

I address this issue by making attribute Authorized-Users a signed attribute. That is, each user appearing in the attribute is either to be allowed or disallowed, as denoted by a preceding sign + or -. Table 3.3 shows an example where the second authorization prohibits a specific user Alice from accessing the patient's name while allowing all other doctors to access the same data. An additional issue here is the potential conflicts between signed values. For example, if Alice is a doctor, then the first signed Value +Doctor will grant Alice the access (since she is a doctor) but the second value -Alice will prohibit this access. The solution is to define meta-policies for resolving conflicts, such as the most-specific-take-precedence meta-policy by which Alice will be prohibited (since Alice is more specific than Doctor).

Furthermore, my mechanism supports the union of authorizations as well as the open policy. Take the authorization 3 in Table 3.3 as an instance, the information of "*EPR.Patient*" can be disclosed to Nurse and Chris using the same tuple of Table *Privacy-Authorization*. In that example, suppose that Chris is a doctor but not the nurse. However, if a patient makes his/her preference, and he/she thinks that the private data should be disclosed to another doctor, Chris, for the same purpose as well as the preference he/she is making for nurse at the same time. The authorizations for this case can be represented as "*Authorization 3*" (a union) in Table 3.3.

Table 3.3: Signed Attribute *Authorized-Users*

Table <i>Privacy-Authorization</i>				
<i>Authorization-ID</i>	<i>Purpose</i>	<i>Table</i>	<i>Attribute</i>	<i>Authorized-Users</i>
1	Treatment	EPR	Diagnosis	+Doctor
2	Treatment	EPR	Patient	+Doctor, -Alice
3	Treatment	EPR	Patient	+Nurse, +Chris

### 3.3 Multi-Dimensional Hippocratic Databases

The original design of HDB is based on the relational database model, which is suitable for small to medium size operational applications. However, analytical applications dealing with a large amount of data, such as a medical data warehouse, will typically adopt a multi-dimensional model instead of the relational model [Pen]. I now extend the schema design introduced in subsection 3.2.1 to a multi-dimensional model. I also make optimizations to reduce the storage requirements of privacy metadata.

A multi-dimensional model such as a star schema, consists of a number of dimensional tables and a fact table, with foreign key constraints linking the two (a simple multi-dimensional EPR database schema has been given in subsection 3.1.3). Table 3.4 shows a star schema augmented with attributes “*Purpose*” and “*Authorization-ID*”. Table *EPR* is the fact table. It can be observed that the fact table includes much redundancy due to various purposes associated with the same record. Such redundancy may be acceptable in an operational database, but it is usually prohibitive for analytical applications based on a multi-dimensional model where the cardinality of the fact table can easily go beyond millions. This situation will be exasperated when table *EPR* includes attributes of object data types, such as laboratory test measurements in multimedia format (for example, X-ray, CT scan or MRI images).

To avoid the extra storage overhead introduced by multiple purposes, I decompose the

Table 3.4: HDB in Star Schema

HDB based on Star Schema  
with Fine-Grained Authorizations

Table	Attributes
<i>EPR</i> (fact table)	Patient-ID, Doctor-ID, Date-ID, Diagnosis, , Prescription, Purpose, Authorization-ID
<i>Patient</i>	Patient-ID, Name, Gender, Age, Address, City, Region, Province, Country, Purpose, Authorization-ID
.....	.....

Table Fact Table *EPR*

<i>Patient-ID</i>	...	<i>Diagnosis</i>	...	<i>Purpose</i>	<i>Authorization-ID</i>
001	...	Flu	...	Treatment	1,2
001	...	Flu	...	Statistics	3
001	...	Flu	...	Prescription	4
001	...	Flu	...	Research	5
001	...	Diabetes	...	Treatment	1,2
...	...	...	...	...	...

Table 3.5: Schema for Multi-Dimensional HDB

Table	Attributes
<i>EPR</i> (fact table)	EPR-ID, Patient-ID, Doctor-ID, Date-ID Diagnosis, Prescription
<i>Patient</i>	Patient-ID, Name, Gender, Age, Address, City, Region, Province, Country, Purpose, Authorization-ID
<i>Doctor</i>	Doctor-ID, Name, Gender, Age, Position, Address, City, Region, Province, Country, Purpose, Authorization-ID
.....	.....
<i>Purpose-Authorization</i>	EPR-ID, Purpose, Authorization-ID

fact table through schema normalization. As shown in Table 3.5, we add attribute “*EPR-ID*” as a primary key to table *EPR* and at the same time remove attributes “*Purpose*” and “*Authorization-ID*” from the table. I then introduce a new table *Purpose-Authorization* to record the purpose and authorizations for each EPR record. I do not normalize the dimensional tables since they typically have a much smaller cardinality than that of the fact table and the redundancy introduced by multiple purposes is thus usually acceptable.

Another issue relevant to a multi-dimensional HDB is that there are two different levels of authorizations for an attribute in the dimensional tables. For example, Bob may want his

Table 3.6: Access Control for Dimensional Tables

Fact Table <i>EPR</i>				
<i>EPR-ID</i>	<i>Patient-ID</i>	...	<i>Diagnosis</i>	...
000001	001	...	Flu	...
000002	001	...	AIDS	...

Table <i>Patient</i>				
<i>Patient-ID</i>	<i>Name</i>	<i>Gender</i>	<i>Address</i>	...
001	Bob	...	...	...

Table <i>Purpose-Authorization</i>		
<i>EPR-ID</i>	<i>Purpose</i>	<i>Authorization-ID</i>
000001	Treatment	1

<i>Privacy-Authorization</i>				
<i>Authorization-ID</i>	<i>Purpose</i>	<i>Table</i>	<i>Attribute</i>	<i>Authorized-Users</i>
1	Treatment	Patient	Address	Doctor

address to be disclosed to doctors for treatment purpose, which can be represented by an authorization for table *Patient* and attribute *Address*. However, Bob may want his address disclosed for symptom *Flu* but not for *AIDS*. This second authorization cannot be represented as above because the “*Authorization-ID*” in table *Patient* is not associated with attribute “*Diagnosis*” in the fact table. My solution is to allow an authorization in the fact table to specify an attribute in the dimensional table. As shown in Table 3.6, the authorization has its “*Authorization-ID*” appearing in the fact table (that is, table *Purpose-Authorization*) but specifies an attribute “*Address*” in a dimensional table *Patient*.

From the fact table *EPR*, there exist two tuples for the same patient Bob, with different Diagnosis: *Flu* and *AIDS*. Bob doesn’t want his address information to be disclosed along with the diagnosis information. In this situation, the table *Privacy-Authorization* in the multi-dimensional HDB provides authorizations for the relationship between fact table tuples and information in dimensional tables. The authorization in table *Privacy-Authorization* (of Table 3.6) limits the access to “*Patient.Address*” for doctors, simultaneously, the tuples in

*Privacy-Authorization* limits the access at a finer-grained level, that is, only “*EPR 000001*” can be disclosed with a diagnosis *Flu*.

## Chapter 4

# HIERARCHICAL SPECIFICATION OF PRIVACY PREFERENCES

This chapter proposes a hierarchical approach to the specification of privacy preferences. First, Section 4.1 defines hierarchies that are inherent to each dimension of privacy preferences. Section 4.2 then proposes methods to simplify a user’s tasks in specifying his/her privacy preferences. Section 4.3 addresses conflict resolution for specified privacy preferences. Finally, Section 4.4 describes the implementation of this hierarchical approach.

### 4.1 Hierarchies in Privacy Preferences

In a complicated application, such as an e-Health system, we are usually dealing with data organized at different granularity levels, data recipients with different ranks and responsibilities, and purposes that may dominate each other. For example, a patient’s address naturally forms a hierarchy: “*Address*  $\rightarrow$  *City*  $\rightarrow$  *Province*  $\rightarrow$  *Country*”. Here each arrow indicates a further level of data generalization.

Two facts about hierarchies are worth noting. First, in terms of privacy, more generalized data are usually less sensitive. For example, a patient may not want his/her exact address to be published but he/she may be less concerned about the state or country. Second, the levels of generalization usually form a partially ordered set (poset)[KBR00]. As depicted in




Figure 4.1: Hierarchy on Private Data

Figure 4.1, the “*Home-Address*” and “*Work-Address*” may not be comparable in terms of sensitivity, neither do the “*State*” (“*Province*”) and “*Region*”.

Hierarchies are also inherent to the other two dimensions of privacy preferences, namely, recipients of data and purposes. First, recipients of private data in an organization typically have their privileges assigned based on their ranks or positions. For example, Figure 4.2 shows a simple hierarchy on data recipients in an e-Health system. In order to supervise nurses, a supervisor may need to access the same data accessible to the nurses.

Second, purposes can usually be refined into sub-purposes. As mentioned in [MMZ], the purpose “*Shipping*” dominates sub-purposes “*Regular-Shipping*” and “*Shipping-by-Courier*”, “*Delivery*” dominates “*Direct-Delivery*” and “*Delivery-by-Post*”, and “*Notification*” dominates “*Notification-by-Fax*” and “*Notification-by-Email*”. In an e-Health system, we consider such purposes: “*Treatment*”, “*Statistics*”, “*Research*”, “*Surgery*”, “*Prescription*”, and “*Diagnosis*”. It is not difficult to see that “*Surgery*” and “*Diagnosis*” can be regarded as the

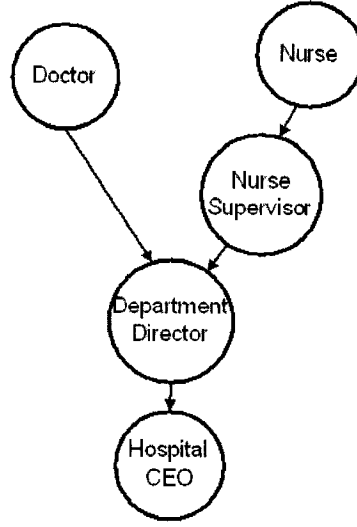


Figure 4.2: Hierarchy on Data Recipients

sub-purpose of “*Treatment*”, as shown in Figure 4.3.

We can define a single hierarchy combining all three dimensions of privacy preferences. However, additional cares need to be taken during such a combination because hierarchies on those three dimensions have different semantics. I approach this issue by taking the point of view of a user specifying his/her privacy preferences. For example, if the user agrees to disclose the address, then likely he/she may not care about the city, state, or other coarser grained data (we shall address exceptions later on). Similarly, if the user is willing to disclose the data to a nurse, then the nurse’s supervisor may also need that access; if the consent is obtained for shipping purpose, then the consent may simply extend to both sub-purposes. That is, the three dimensions should be combined in such a way that the finest-grained data will coincide with recipients with the least privileges and purposes that are not dominated by others, as formally stated in Definition 1.

*Definition 1: Let  $(D, \preceq_D)$ ,  $(R, \preceq_R)$ ,  $(P, \preceq_P)$  be posets defined on private data, recipients, and purposes, respectively; the relations  $\preceq_D$ ,  $\preceq_R$ , and  $\preceq_P$  are defined in such a way that*






Figure 4.3: Hierarchy on Purposes

$d_1 \preceq_D d_2$  if  $d_1$  is less sensitive than  $d_2$ ,  $r_1 \preceq_R r_2$  if  $r_1$  has more privileges than  $r_2$ , and  $p_1 \preceq_P p_2$  if  $p_1$  is a sub-purpose of  $p_2$ . We define the privacy preference poset as the product of these three posets  $\langle M, \preceq \rangle$  where  $M = D \times R \times P$ , and  $\preceq$  is the composition of the three relations  $\preceq_D$ ,  $\preceq_R$ , and  $\preceq_P$ .

## 4.2 Specifying Preferences Using Hierarchies

I now propose methods for specifying privacy preferences. My methods aim to simplify the specification of preferences by taking advantage of the privacy preference poset given in Definition 1. More specifically, the poset is defined in such a way that the consent on disclosures at a higher level in the poset implies the consent on disclosures at lower levels. Such disclosure inheritance may allow a user to specify a privacy preference without explicitly enumerating all the levels he/she is willing to disclose. On the other hand, exceptions to disclosure inheritance may certainly exist and users should be allowed to specify such exceptions inside a specification. The specification of exceptions can again employ the hierarchies




Figure 4.4: Hasse Diagram for Hierarchy on Private Data

if necessary.

I consider three options for specifying a preference, with increasing degree of utilization of the disclosure inheritance (notice that other options also exist but are not discussed here). For each option, I show how to generate the effective set of attributes to be disclosed according to the preference, namely, the disclosure set. I shall use the example of a poset on the data dimension as given in Figure 4.1 for simplicity. Also for simplicity, I denote the poset using an abstract version  $(H, \preceq)$ , as shown in Figure 4.4, where  $H = \{a, b, c, d, e, f, g, h\}$  with  $a$  for “Home Address”,  $b$  for “Work Address”,  $c$  for “Hospital Address”,  $d$  for “City”,  $e$  for “Province/State”,  $f$  for “Region”,  $g$  for “Country”, and  $h$  for “Continent”. The relation  $\preceq$  remains the same.

#### 4.2.1 Option 1

This option allows a user to explicitly specify either what to disclose or what to be kept private. That is, the disclosure inheritance is not employed at all. Formally, given the privacy




Figure 4.5: Disclosure Inheritance for Option 2

preference poset  $\langle M, \preceq \rangle$ , the privacy preference specification is simply a set selected by the user  $X \subseteq M$ . We can then generate the disclosure set as either  $D = X$  (if  $X$  is to be disclosed) or  $D = M - X$  (if  $X$  is to be kept private).

#### 4.2.2 Option 2

The second option allows a user to partially utilize the disclosure inheritance. More specifically, he/she can specify a range in the privacy preference poset to be disclosed but the exception, that is what to be kept private, is still specified explicitly as a set. More precisely, given the privacy preference poset  $\langle M, \preceq \rangle$ , the preference is specified with an upper bound  $X \subseteq M$ , a lower bound  $Y \subseteq M$ , and a set  $Z$  to be kept private. We can then generate the disclosure set  $D \subseteq M$  as the set  $\{d \mid \exists x \in X, \exists y \in Y, y \leq d \leq x\} - Z$ . For example, Figure 4.5 shows a specification where  $X = \{b, c\}$ ,  $Y = \{h\}$ , and  $Z = \{e, f\}$ . The disclosure set is  $D = \{b, c, d, g, h\}$ , as shown in Figure 4.6. Notice the relation  $(d, g)$  is due to transitivity.




Figure 4.6: Disclosure Set Example for Option 2

### 4.2.3 Option 3

The third option allows a user to fully utilize the disclosure inheritance by specifying ranges in the privacy preference poset both for disclosing and for keeping private. Given the privacy preference poset  $\langle M, \preceq \rangle$ , the preference is specified with a pair of upper and lower bounds  $X \subseteq M$  and  $Y \subseteq M$  for disclosing, and another pair  $W \subseteq M$  and  $Z \subseteq M$  for keeping private. We can then generate the disclosure set  $D \subseteq M$  as  $\{d | \exists x \in X, \exists y \in Y, y \leq d \leq x\} - \{p | \exists w \in W, \exists z \in Z, w \leq p \leq z\}$ . For example, Figure 4.7 shows such a specification where  $X = \{b, c\}$ ,  $Y = \{h\}$ ,  $W = \{d\}$ , and  $Z = \{g, f\}$ . The disclosure set is  $D = \{b, c, h\}$  and is shown in Figure 4.8.

## 4.3 Conflict Resolution

In my solution, each preference specified by users will be stored in the backend database in the form of its corresponding disclosure set for later reference. However, preferences specified over time may conflict with each other. Prompting the user to resolve the conflicts by choosing a preference over the other is a possible solution. However, it is usually not




Figure 4.7: Disclosure Inheritance for Option 3




Figure 4.8: Disclosure Set Example for Option 3

feasible to completely rely on users for making every decision. Instead, we can provide users general guidelines for resolving a potential conflicts, namely, meta-policies. We consider three representative meta-policies (other meta-policies are also possible), namely, “*Latest Take Precedence*”, “*Disclosure Take Precedence*”, and “*Denial Take Precedence*”.

#### 4.3.1 Latest Take Precedence

This meta-policy allows a user to update his/her preferences. That is, any potential conflict between the new preference and existing preferences will be resolved by overriding the latter with the former. Since we store previous preferences as their corresponding disclosure sets, the positive part of a new specification (that is, what to be disclosed) will not cause conflicts with existing preference. However, the negative part (that is, what to be kept private) may cause a conflict if it overlaps with existing preferences.

More specifically, given the privacy preference poset  $\langle M, \preceq \rangle$ , let  $D_0 \subseteq M$  be the disclosure set of existing preferences, then the new disclosure set  $D_1$  can be generated according to different options of specification as follows. For the first option,  $D_1 = D_0 \cup X$  (if  $X$  is to be disclosed) or  $D_1 = D_0 - X$  (if  $X$  is to be kept private). For the second option,  $D_1 = D_0 \cup \{d \mid \exists x \in X, \exists y \in Y, y \leq d \leq x\} - Z$ . For the third option,  $D_1 = D_0 \cup \{d \mid \exists x \in X, \exists y \in Y, y \leq d \leq x\} - \{p \mid \exists w \in W, \exists z \in Z, w \leq p \leq z\}$ . It is worth noting that we cannot simply compute the disclosure set of the new specification, and then union the result with the existing disclosure set. This is due to the simple fact  $A \cup (B - C) \neq A \cup B - C$ . For example, for the second option, the union between the existing and new disclosure sets would be  $D_0 \cup (\{d \mid \exists x \in X, \exists y \in Y, y \leq d \leq x\} - Z)$ , which is different from  $D_0 \cup \{d \mid \exists x \in X, \exists y \in Y, y \leq d \leq x\} - Z$ .

For example, if the existing disclosure set in Figure 4.2 is  $\{a, b, c, d\}$ , and the latest disclosure set is  $\{b, c, d, e, f\} - \{c, d\}$ , the new disclosure set would be  $\{a, b, c, d\} \cup \{b, c, d, e, f\} - \{c, d\} = \{a, b, e, f\}$  but not  $\{a, b, c, d, e, f\}$ .

### 4.3.2 Disclosure Take Precedence

This meta-policy overrides any conflicting preferences with the consent to disclosure. The meta-policy models the situation where accesses are not audited. In such a case, we must assume that any disclosure may happen immediately after it is consented by the user, so withdrawing such consent is not possible. Since we store previous preferences as their corresponding disclosure sets, any conflict caused by the negative part of a new preference will be ignored. Computing the new disclosure set by this meta-policy is straightforward, because we simply union the new disclosure set with the existing one. More specifically, for the first option,  $D_1 = D_0 \cup X$  (if  $X$  is to be disclosed) or  $D_1 = D_0$  (if  $X$  is to be kept private). For the second option,  $D_1 = D_0 \cup (\{d | \exists x \in X, \exists y \in Y, y \leq d \leq x\} - Z)$ . For the third option,  $D_1 = D_0 \cup (\{d | \exists x \in X, \exists y \in Y, y \leq d \leq x\} - \{p | \exists w \in W, \exists z \in Z, w \leq p \leq z\})$ .

For example, if the existing disclosure set in Figure 4.2 is  $\{a, b, c, d\}$ , and the latest disclosure set is  $\{b, c, d, e, f\} - \{c, d\}$ , the new disclosure set would be  $\{a, b, c, d, e, f\}$ .

### 4.3.3 Denial Take Precedence

This meta-policy is the dual to the “*Disclosure Take Precedence*” meta-policy. That is, a user can disallow a disclosure even if it has previously been consented upon. This meta-policy may be more desirable in the viewpoint of a user. Because we already store previous preferences as disclosure sets, this meta-policy behaves in exactly the same way as the “*Latest Take Precedence*” meta-policy, although they have different semantics. That is, we can generate the new disclosure set by subtracting the negative part of the new specification from the

union of the existing and new disclosure sets.

We have been assuming that preferences are stored in the form of disclosure sets. Some applications may prefer the specification itself to be stored without explicitly computing its corresponding disclosure set. For example, if most preferences involve a large portion of the privacy preference poset, then storing specifications rather than disclosure sets will lead to significant savings in storage (at the expense of more computational overhead for computing the disclosure set at run time). With such an approach, the negative part of existing preferences and that of the new preference can both cause conflicts. Moreover, the “*Denial Take Precedence*” meta-policy will need to be handled differently from the “*Latest Take Precedence*” meta-policy. Specifically, for the “*Latest Take Precedence*” meta-policy, we need to compute disclosure sets in an incremental manner. For example, for the first option, this amounts to iteratively evaluate  $D_i = D_{i-1} \cup X_i$  (if  $X_i$  is to be disclosed) or  $D_i = D_{i-1} - X_i$  (if  $X_i$  is to be kept private) starting from the first specification to the current one. For “*Denial Take Precedence*” meta-policy, on the other hand, we can simply compute the final disclosure set as  $\bigcup D_i - \bigcup X_i$ .

## 4.4 Implementation Option

This section will discuss the implementation of hierarchical authorizations, introduce how to represent those authorizations using privacy metadata of Hippocratic databases, and propose a snowflake metadata schema to record the authorizations, which is capable of substituting existing metadata when necessary.



Table 4.1: Authorization under Original HDB Design

<i>Authorization-ID</i>	<i>Purpose</i>	<i>Table</i>	<i>Attribute</i>	<i>Authorized-Users</i>
1	Treatment	Patient	City	Nurse
2	Treatment	Patient	Province	Nurse
3	Treatment	Patient	Country	Nurse
4	Treatment	Patient	City	Nurse Supervisor
5	Treatment	Patient	Province	Nurse Supervisor
6	Treatment	Patient	Country	Nurse Supervisor
7	Surgery	Patient	City	Nurse
8	Surgery	Patient	Province	Nurse
9	Surgery	Patient	Country	Nurse
10	Surgery	Patient	City	Nurse Supervisor
11	Surgery	Patient	Province	Nurse Supervisor
12	Surgery	Patient	Country	Nurse Supervisor

#### 4.4.1 Using HDB Metadata

As I have introduced in previous chapters (especially Chapter 3), we can implement fine-grained authorizations based on Hippocratic Databases. HDB privacy metadata is the basis of the hierarchical authorization as well as fine-grained authorization. In early section of this chapter, I illustrated the three dimensions in *Privacy-Authorization Table*, the disclosure set based on hierarchies and conflict resolution. We can use *Privacy-Authorization Table* to record all the authorizations and a group of tuples (authorizations) to represent the disclosure set. We extract this type of disclosure set from metadata schema to resolve conflicts resulting from authorizations executed at different time.

For example, if an authorization is given as the following:  $\{City, Province, Country\}$ ,  $\{Nurse, NurseSupervisor\}$ ,  $\{Treatment, Surgery\}$  (for the relationships of those elements, refer to Figure 4.1, 4.2 and 4.3), then totally twelve authorizations will be derived, as shown in Table 4.1.

There are twelve tuples in *Privacy-Authorization Table*. This original HDB solution performs well for fine-grained authorizations after adding an “*Authorization-ID*” column.

Nevertheless, we cannot recognize the hierarchies (especially the relationships) in those three dimensions from the given authorizations. Thus, we cannot easily extract existing disclosure sets to compare with latest disclosure set for conflict resolution. For this reason, I will discuss another solution in subsection 4.4.2.

#### 4.4.2 Using Snowflake Metadata

Hierarchies naturally exist in many aspects of an HDB, such as EPR data, authorized users, and purpose. I have proposed a hierarchical approach for simplifying the specification of privacy preferences and authorizations by leveraging such hierarchies. However, how to represent privacy authorizations in an HDB in the presence of such hierarchies is not addressed, and I have discussed the shortcoming of representing hierarchical authorizations under original HDB in section 4.4.1. I now describe a snowflake schema-based design to solve this problem.

Privacy authorizations can be derived from hierarchies on EPR data, authorized users, and purposes. For example, if “*Address*” has been authorized then less sensitive data of “*City*” or “*Country*” are implicitly authorized altogether (exceptions to such inheritance are handled in Section 4.2, and I shall not address it here). In the authorized user dimension, if “*Nurse*” is authorized to access certain EPR data, then users with higher privileges, such as “*Nurse Supervisor*” will be authorized to access the same data as well. For the purpose dimension, if a purpose “*Treatment*” has been authorized then a more specific purpose (sub-purpose) such as “*Surgery*” will also be authorized. It is worth noting that the inheritance of authorizations is reversed for the purpose dimension in contrast to the other two. For example, “*Country*” dominates “*City*” (in the sense that a country may correspond to many cities) and authorizations on “*City*” imply those on “*Country*”, whereas purpose dominates

Table 4.2: Snowflake Schema for EPR Data and Metadata

Hierarchies in EPR Schema	
<b>EPR Table</b>	<b>Attributes</b>
<i>Address</i>	Address-ID, Address, City-ID
<i>City</i>	City-ID, City, Province-ID
<i>Province</i>	Province-ID, Province, Country-ID
<i>Country</i>	Country-ID, Country

Hierarchies in Metadata	
<b>Metadata Table</b>	<b>Attributes</b>
<i>Nurse</i>	Nurse-ID, Nurse-Name, Supervisor-ID
<i>Supervisor</i>	Supervisor-ID, Supervisor-Name
...	...
<i>Sub-Purpose</i>	Sub-Purpose-ID, Sub-Purpose, Purpose-ID
<i>Purpose</i>	Purpose-ID, Purpose

sub-purpose but authorizations on purpose imply those on sub-purpose.

Neither the original HDB design nor the design in previous sections can provide efficient implementation of such hierarchical authorizations. I provide a solution based on the snowflake schema instead. Table 4.2 shows an example of my design. For simplicity, we only include the address hierarchy in EPR data. In this schema, all tables are fully normalized, and the hierarchical relationships between attributes are represented as foreign key constraints. The metadata tables include both the authorized user dimension and the purpose dimension.

In the presence of hierarchies, authorizations can be given as any combination of three sets of attributes where each set is from the EPR data, authorized user, and purpose dimension, respectively. Under the original design of HDB, this may result in a large number of authorizations. For example in Table 4.1.

With my approach, the collection of authorizations in Table 4.1 can be conveniently represented using a single record. More specifically, I redesign the *Privacy-Authorization Table* as shown in Table 4.3. In addition to the attribute “*Authorization-ID*” (required for

Table 4.3: Redesigned Schema for Hierarchical Authorizations

Table	Attributes(Columns)
<i>Privacy Authorization</i> ( <i>fact table</i> )	Authorization-ID, EPR Table, Data-ID, Purpose Table Purpose-ID, Authorized-Users Table, Authorized Users-ID

Table 4.4: Authorization under the Redesigned Schema

<i>Authorization</i> <i>-ID</i>	<i>EPR</i> <i>Table</i>	<i>Data</i> <i>-ID</i>	<i>Purpose</i> <i>Table</i>	<i>Purpose</i> <i>-ID</i>	<i>Authorized</i> <i>Users Table</i>	<i>Authorized</i> <i>User-ID</i>
1	City	All	Purpose	1	Nurse	All
2	Address	1	Sub-Purpose	2	Nurse Supervisor	2

fine-grained authorization as described in Chapter 5), each authorization now includes three pairs of attributes. Each pair indicates the exact elements in the schema is authorized in each of the three dimensions: EPR data, authorized users, and purpose.

Table 4.4 shows an example where the Data-ID is linked to Address-ID, City-ID, etc. Due to the aforementioned hierarchies, the single record ( “*Authorization 1*”) in Table 4.4 is sufficient for representing the twelve authorizations given in Table 4.1, in which “*Purpose 1*” represents “*Treatment*” in Table *Purpose* and the value “*All*” stands for all the disclosed tuples for this elements.

From viewpoint of fine-grained access control, this snowflake metadata can represent not only the hierarchical authorizations, but also the authorization in a single tuple. In Table 4.4, the values of other three columns: Data-ID, Purpose-ID and Authorized Users-ID can be used to determine authorizations on either none-or-all or exact instances. This is consistent with the fine-grained authorizations introduced in Chapter 3. The fact table *Privacy-Authorization* can record not only the elements to be disclosed (such as City, Nurse in Table 4.4) but also the exact values of elements in the disclosure sets (such as Data-ID and Authorized User-ID that refer to the fine-grained level data and recipients). As for the disclosure purpose, we must use this fine-grained level of disclosure, because the exact

sub-purpose or purpose (such as Surgery or Treatment) are what we need to complete the disclosure inheritance. For example in Table 4.4, the record “*Authorization 2*” means that “*Address 1*”(and more generalized data) is disclosed to “*Nurse Supervisor 2*” (and higher-positioned recipients) for “*Sub-Purpose 2*”.

Figure 4.9 illustrated the snowflake metadata schema for the implementation of hierarchical authorizations. For simplicity, only a few attributes in the hierarchies have been presented. Some facts are worth noting:

- *Privacy-Authorization Table* is the fact table of this schema, and it can be linked to every table of the three dimensions with constraints on three foreign keys. In order to provide better presentation of Figure 4.9, only Table *Doctor*, *Nurse*, *Address*, *City*, *Sub-Purpose* and *Purpose* have been linked to the fact table. Actually, “*Data-ID*” in the fact table can be linked to “*Address-ID*”, “*City-ID*” .etc; “*Authorized-Users-ID*” can be linked to “*Nurse-ID*”, “*Doctor*” .etc; and “*Purpose-ID*” can be linked to “*Purpose-ID*” and “*Sub-Purpose-ID*”.
- In the fact table, the attributes: “*EPR-table*”, “*Purpose-Table*”, “*Authorized-Users-Table*” have recorded the links from *Privacy-Authorization Table*. Once the names of three tables are recorded in those three columns, respectively, such as *City*, *Nurse* and *Purpose*, all the lower layer tables, such as *Nurse Supervisor*, *Province*, *Country* and *Sub-Purpose* can be authorized together with the upper layer tables.
- As mentioned before, the inheritance of authorization on Purpose dimension is reversed from the other two dimensions.
- One record could only handle one simple path of each dimension. Authorizations on complex hierarchies require more records. For example, if the patient authorizes Doctor

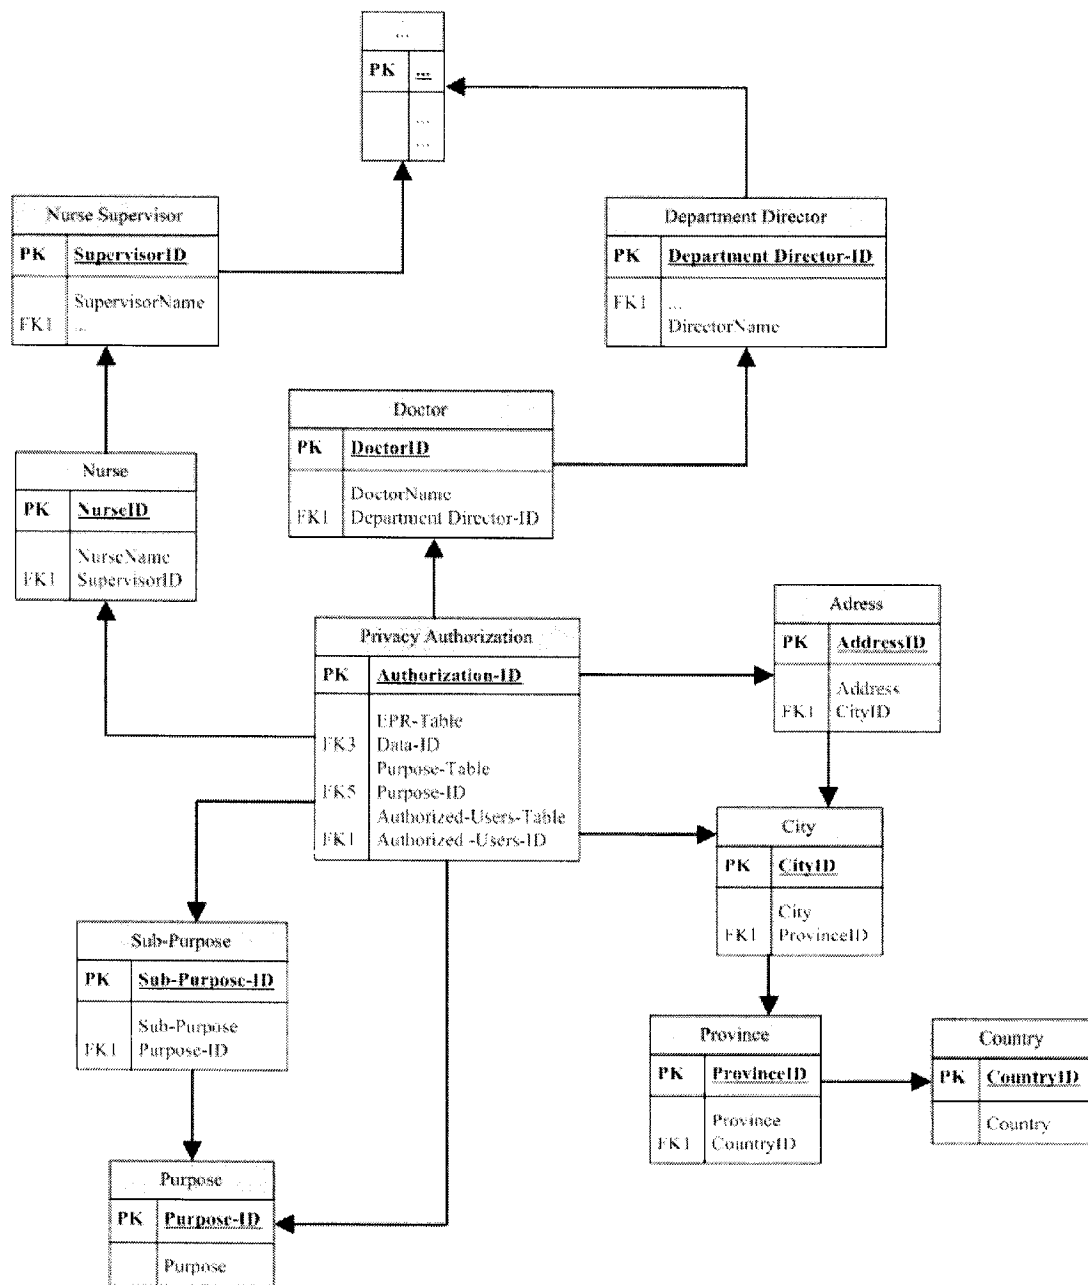


Figure 4.9: Snowflake Metadata for Hierarchical Authorizations

and Nurse to access his/her private data in one disclosure set, two tuples in the fact table's column "*Authorized User-Table*" would be *Doctor* and *Nurse*. Moreover, we can slightly modify this schema to satisfy all the disclosure options in complex hierarchies.

# Chapter 5

## IMPLEMENTATION

In this chapter, implementation issues are discussed. Firstly, I introduce the tools and environments selected to implement the e-Health system and the privacy protection components. Secondly, the schemata implementation details are illustrated with experiments on redesigned and existing schemata. Thirdly, I describe implementation details for preferences specification, which is an essential part to preserve privacy in e-Health system. Finally, a scenario for private data access is introduced in section 5.4.

### 5.1 Implementation Environment

#### 5.1.1 SOA with BEA Weblogic 8.1.2

Loosely coupled software services are more and more popular in business process handling nowadays. Web services bring us easier software integration using XML, we thus determined to use web service of Service Oriented Architecture (SOA) for our e-Health system implementation. Meanwhile, to provide a friendly interface for users (e.g. patients and doctors) and make it easier to maintain and upgrade the whole system, a web-based platform has been selected, while web pages (JSP or HTML .etc) represent the user interface for clients' operations.

There are several popular deploying platforms for web service, such as Apache, the .Net



framework, and J2EE. They can implement web service on the same platform for running the web applications. There is a test for the two leading commercial J2EE servers, BEA WebLogic and IBM's WebSphere, plus a solid also-ran, Sybase's EAServer, and the most popular open source J2EE server JBoss. The test [fJWSE] evaluates the related management capabilities as well as the support for core web service standards, SOAP, XML-RPC (Remote Procedure Call), WSDL and UDDI, flexible configuration and a set of features such as JMX, JNDI, JMS and JTA that would be expected in any enterprise-class Java platform.

We make use of a powerful and robust web platform, WebLogic Workshop 8.1.2 provided by BEA, to complete the system implementation.

Firstly, all the services are designed in the web service file (\*.jws). Database control with connection pooling has been used to access EPR data for different methods with different purposes and different recipients. The SQL statements of data usage for each method are also composed within the database control interface.

For user interfaces, we use jsp files controlled by Java Page Flow [Inca] which is a feature set built upon a Struts-based Web application programming model. Java Page Flow features include runtime support for the Web application programming model and tools that enable developers to quickly and easily build applications based upon the model and it's based on Model-View-Controller (MVC) architectural pattern. The views (jsp files) and controller (jpf file) are generated by the web service file, recomposed and extended to entire functions programming. The example of our EPR services will be illustrated in later part of this chapter.

Table 5.1: Web Service Deployment Platforms

Product	Cost	Platforms	Bottom Line
JBoss 3.2.1/ Apache Tomcat 5.0	Free	Java 1.3 (or later) JVMs	<ol style="list-style-type: none"> <li>1. Excellent support for importing and exporting existing web services</li> <li>2. Good standard support</li> <li>3. Require more expertise and initial configuration</li> <li>4. Lack many of simpler enterprise class configuration and management tools</li> </ol>
Sybase EAServer 4.2	From \$7500 to \$2800 per CPU	AIX, HP-UX, Linux, Windows 2000	<ol style="list-style-type: none"> <li>1. Good integration between app server and tools, a slick and efficient interface.</li> <li>2. it's slightly hampered by a heavy reliance on server restarts and a help system that could use a little work</li> </ol>
IBM WebSphere Application Server 5.0	From \$8000 to \$25000 per CPU	AIX, Red Hat Linux, Solaris, Windows NT/2000	<ol style="list-style-type: none"> <li>1. Allows granular control of the app server and web services and does so with almost equal skill from its GUI and CLI;</li> <li>2. Confusing installation process that can add up to more than 1GB for even a basic installation</li> </ol>
BEA WebLogic Sever 8.1	Starts at \$10000 per CPU includes WebLogic Workshop	Linux, Unix, Windows	<ol style="list-style-type: none"> <li>1. Good integration between tools and server, easy setup, solid web server standards support</li> <li>2. Excellent IDE</li> <li>3. Importing and manipulating existing web services would be easier without some additional steps</li> </ol>

Table 5.2: Main Security Features of Database Products

Product	Authentication	Access Control	Encryption	Auditing
Oracle 9i	Password -based Host-based	FGAC by VPD	TDE, OAS and SSL	Triggers Fine-grained Auditing
MS SQL Server/ 2005	Connection String Kerberos .etc	Execution Context	Digital Certificates SSL	Triggers SQL Trace
IBM DB2	Kerberos .etc Host-based	ID based, Role based, .etc	Column Level , Value Level	DB2 Audit Trace
Sybase ASE15	Password, Kerberos, LDAP	DAC Policy based	SSL, Kerberos	Audit Trail
MySQL	Password		Hashing	

### 5.1.2 Oracle 9i Database Support

Due to the stringent privacy and security requirement for data in various medical systems, especially the data processed and stored in EPR systems, the database product selection is chiefly based on the security features of various commercial database vendors.

We survey the products according to several security features: Authentication, Access Control, Encryption, Auditing, Availability .etc. The candidate commercial database products are Oracle 9i, IBM DB2, Sybase ASE15, Microsoft SQL Server 2005 and an open source database system: MySQL.

Moreover, as I have presented in Section 2.1, virtual private database (VPD) is the aggregation of server-enforced, fine-grained access control, together with a secure application context in Oracle database system. By dynamically appending SQL statements with a predicate, VPD limits access to data at record level and ties the security policy to the table or view itself.

Apparently, we need to limit the data access at the hospital level, because some hospitals share the same EPR schema in the same database. Each EPR service belongs to one

specific hospital, and the data usage for that service should be limited to the data of that hospital, except that valid authorizations have been made for special purposes, the data can be accessed across hospitals.

Assume that the table EPR is like this: EPR (EPR\_ID Number (10), Patient\_ID Number (6), Diagnosis varchar2 (30), Hospital\_ID Number (3)). Even though administrator could access all the data without any restrictions, if we want to allow applications to access only data in the hospital which it belongs to, VPD can satisfy this demand.

Consequently, when the EPR services use the data, the application connects database with the value of Hospital\_ID and execute statement: `Select * from EPR;`

Here, only the rows where  $Hospital\_ID = \{Hospital\}$  are visible for the application.

Moreover, we can also limit data access at the USER level using USER\_IDs (either Patient\_ID or Doctor\_ID) or the tuples level using Authorization\_IDs according to the proposed fine-grained authorization and access control mechanism. Since many kinds of privacy metadata have been designed in the database system, we can limit the access based on these systems and implement the access control not only at application level but also at database level. If we take advantage of these features of Oracle 9i to keep the authorization and access control implementation at database level, the privacy protection can be enhanced.

## 5.2 Hippocratic Databases Schemata Implementation

In the privacy protection subsystem of our e-Health system, the redesigned schemata based on multi-dimensional HDB, the existing schemata for multi-dimensional HDB and the proposed snowflake metadata schema for hierarchical authorizations are deployed in backend databases. First of all, some issues around schemata implementation will be described in this section. Consequently, some experiments on my redesigned schema and existing HDB

schema are illustrated.

### **5.2.1 Oracle Schemata Implementation**

Three schemata: original HDB for multi-dimensional schema, redesigned multi-dimensional HDB and snowflake privacy-authorization schema have been implemented as the EPR schema in a specific user account of the Oracle system.

All the required tables, foreign keys, triggers, sequences, and procedures in the schemata implementation constitute the schemata in my design.

To allow users to specify privacy authorizations, I implement an APPEL preference selector through Web interfaces based on BEA WebLogic 8.1.2. I generate authorizations from the selected options in the Web interfaces using a Java application. The application then stores the generated authorizations in a backend Oracle 9i database. Oracle Enterprise Edition supports fine-grained access control (FGAC) through the VPD feature, which is the aggregation of server-enforced FGAC with a secure application context. My fine-grained authorizations are implemented using this feature by defining policy functions. I implement three schemata as described in Chapter 3 for fine-grained authorizations based on a flat relational model, for multi-dimensional model without hierarchies, and for hierarchical authorizations, respectively. In practice, however, an application may choose to implement only one of those schemata based on its specific needs.

### **5.2.2 Experiments**

To evaluate the previous designs, I populate the database with the adult dataset from the UCI Machine Learning Repository [Rep] (which has been used in many related work on data privacy). There are 32561 patients in the dataset. I made following assumptions. Each

Table 5.3: Storage Requirements for Fine-Grained Authorization

#	Table	Attributes Number	Tuples Number	Storage (Bytes)
1	<i>Patient</i> in Table 3.4 or 3.5	15	32,561	5,242,880
2	<i>EPR</i> in Table 3.4	11	390,732	75,497,472
3	<i>EPR</i> in Table 3.5	9	97,683	16,777,216
4	<i>Purpose-Authorization</i> in Table 3.5	3	390,732	24,117,248
Summation of #1 and #2			80,740,352 Bytes	
Summation of #1, #3, and #4			46,137,344 Bytes	

patient owns three records in the EPR table, and each record is associated with four purposes: “*Treatment*”, “*Research*”, “*Statistics*”, and “*Prescription*”. Table 5.3 shows that my redesigned fine-grained authorization schema requires significantly less total storage than that of the original HDB design, if the record-level fine-grained authorization is to be enforced. Clearly, with more purposes or attributes that hold more storage consuming data (for example, multimedia data), the difference would be more significant.

In addition to removing the limitations of the original HDB model, our solution may also bring significant improvement to storage requirements of e-Health applications. For example, in Table 5.3, it takes over 80 Megabytes if our e-Health system adopts original HDB for privacy protection with fine-grained authorization. However, if our e-Health system take use of the multi-dimensional HDB schema for the same purpose, the storage can be reduced to nearly 46 Megabytes. The evaluation results are based upon the situation addressed in the previous paragraph. When the database is extended to a larger model, the performance improvement would be more clear.

Our contribution of hierarchical specification of privacy requirements will make patients’ job much easier in most cases. For example, if patients want to make all their private data in Figure 4.1 be disclosed to all the data recipients in Figure 4.2 for all the purposes in Figure

Table 5.4: Original HDB Preferences and Hierarchical Preferences

	Tuples Number in Metadata	Specify Preferences Times
Original HDB Preferences	$8 \times 5 \times 5 = 200$	$8 \times 5 \times 5 = 200$
Hierarchical Preferences (Snowflake Metadata)	1	1

4.3, our e-Health system is able to adopt original HDB preferences specifying approach or my solution, hierarchical preferences. The evaluation results for these two different situations are shown in Table 5.4. If a patient should specify preference using original HDB, he/she has to specify private data, data recipients and purposes one by one. Thus, when he/she specifies all the elements in those three figures, he/she must specify them by  $8 \times 5 \times 5 = 200$  times, and the 200 tuples will be generated in the metadata schema. Nevertheless, if the e-Health system takes advantage of hierarchical preferences, only one tuple in the Privacy-Authorization Table of snow-flaked schema will be generated, and the patient only need to specify the upper bounds of those three dimensions. He/she can specify it by once, which brings a much more convenient mechanism.

### 5.3 Preferences Negotiation Implementation

Patients specify their preferences in the frontend of our e-Health system. I implement the preferences negotiation using a web interface, and integrate it into the front end of our web service based portal system. There are three important components in the preference negotiation process, which are generating disclosure set, conflict resolution and APPEL preferences. The implementation issues of these three components will be introduced in this section.

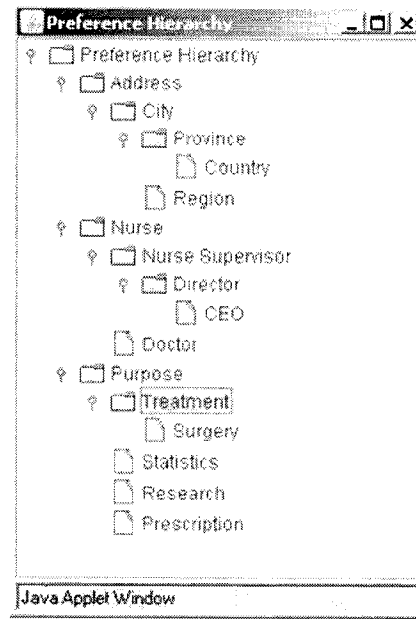


Figure 5.1: The Hierarchy Used in Implementation

### 5.3.1 Generating Disclosure Set

In the preference negotiation web interface, there is a hierarchy consisting of the nodes of three dimensions, as illustrated in Chapter 4. Patients are able to make preferences with hierarchical authorizations. A simple example for hierarchies used for disclosure inheritance is shown in Figure 5.1.

Three dimensions in the implementation are:

- Private Data: Address, City, Province, Country, Region
- Data Recipient: Nurse, Nurse Supervisor, Director, CEO, (Doctor)
- Purpose: Treatment, Surgery, Statistics, Research, Prescription

The patient specifies the preferences using a web interface shown in Figure 5.2, in which the patient selects the upper bounds of private data, data recipients and purposes, and their lower



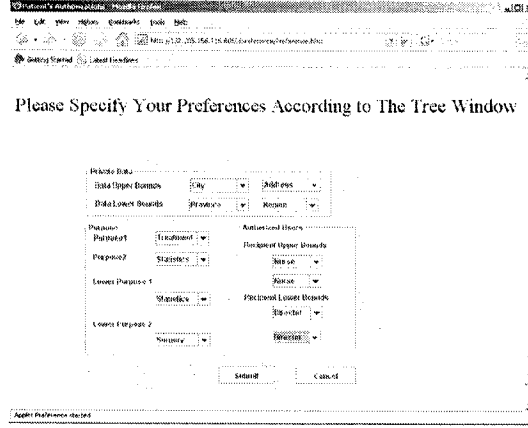


Figure 5.2: Specifying Preferences in Web Interface

bounds. Consequently, a disclosure set based on the selected bounds of those three hierarchies are generated, and the disclosure inheritance are triggered with the generated disclosure set. Finally, a group of authorizations will be inserted into the Table *Privacy-Authorization* (if we still use the original HDB metadata design), or in my redesigned snowflake metadata (the fact table *Privacy-Authorization*).

From the selected bounds shown in Figure 5.2, there are  $4 \times 3 \times 3$  authorizations generated by a single disclosure, and all of the 36 tuples have been inserted into the backend database simultaneously. The differences (before specifying preferences and after specifying preferences) for the rows of the “*Privacy-Authorization Table*” in the database schema is shown in Figure 5.3 and 5.4.

### 5.3.2 Meta-Policy for Conflict Resolution

Conflict resolution for preferences has been illustrated in Chapter 4. In this subsection, the implementation details will be discussed. There are three meta-policies defined for conflict resolution, which are “*Latest-Take-Precedence*”, “*Disclosure-Take-Precedence*” and “*Denial-Take-Precedence*”. Here, I will take the “*Latest-Take-Precedence*” as the example to describe

TOAD for Oracle - [HYPER - Schema Browser (HYPER)]

File Edit Grid SQL Editor Create Database Tools View Help DBS Debug Team Coding Window Help

Columns: Indexes Constraints Triggers Data Scope Grants Synonyms Partitions Subpartitions Stats/Size Referential Used By Policies Auditing

1 Patient Age Treatment Doctors

2 EPR Diagnosis Treatment Doctors

3 Patient Name Treatment Doctors

Row 3 of 3 total rows

SQL Editor Schema Browser

AutoConnect is OFF

Figure 5.3: Privacy-Authorization Table before Specifying Preferences

TOAD for Oracle - [HYPER - Schema Browser (HYPER)]

File Edit Grid SQL Editor Create Database Tools View Help DBS Debug Team Coding Window Help

Columns: Indexes Constraints Triggers Data Scope Grants Synonyms Partitions Subpartitions Stats/Size Referential Used By Policies Auditing

1 Patient Age Treatment Doctors

2 EPR Diagnosis Treatment Doctors

3 Patient Name Treatment Doctors

4 PATIENT Region Treatment Nurse Supervisor

5 PATIENT Region Treatment Director

6 PATIENT Region Treatment Nurse

7 PATIENT Address Treatment Nurse Supervisor

8 PATIENT Address Treatment Director

9 PATIENT Address Treatment Nurse

10 PATIENT Province Treatment Nurse Supervisor

11 PATIENT Province Treatment Director

12 PATIENT Province Treatment Nurse

13 PATIENT City Treatment Nurse Supervisor

14 PATIENT City Treatment Director

15 PATIENT City Treatment Nurse

16 PATIENT Region Surgery Nurse Supervisor

17 PATIENT Region Surgery Director

18 PATIENT Region Surgery Nurse

19 PATIENT Address Surgery Nurse Supervisor

20 PATIENT Address Surgery Director

21 PATIENT Address Surgery Nurse

Row 1 of 21 total rows

SQL Editor Schema Browser

AutoConnect is OFF

Figure 5.4: Privacy-Authorization Table after Specifying Preferences

the meta-policy implementation.

When a patient makes preferences through the web interface, he/she should select the meta-policy if there are some potential conflicts. After a hierarchical authorization has been made by the patient, the system will generate a disclosure set that represents a group of preferences. Consequently, the conflict inspection is triggered.

Conflict inspection is based on the existing preferences, either a single preference or a group of them represented in the snowflake schema, and the new preferences made by the same patient. For example, if the patient disclosed his/her data: *Address*, *City* and *Province*, and the new disclosure set includes *City*, *Country* but not *Province*. During the conflict inspection, a message is displayed: “*Conflict Found!!!*”. The system executes the conflict resolution with the pre-selected meta-policy (such as “*Latest-Take-Precedence*”). Finally, the new disclosure set will be generated. In the above example, the disclosure on *Province* for all the purposes and data recipients will be removed and *Country* will be inserted into the database with corresponding purposes and data recipients. The detailed procedures of the preference negotiation, especially for conflict resolution is shown in Figure 5.5.

### 5.3.3 APPEL Integration

In previous sections, I presented the server-centric architecture [AJSX04] [CLM<sup>+</sup>02b]. In this subsection, the tools and resources for building P3P into the service websites will be illustrated.

Referring to the P3P implementation from W3C [Con], there are some policies creating tools such as P3PEidt, P3PBuilder, .etc. P3PEdit, is a commercial P3P policy generator whose easy-to-use Web based Wizard could quickly generate P3P policies that satisfy Internet Explore’s new privacy requirements. P3PEdit generates: XML format P3P policies,

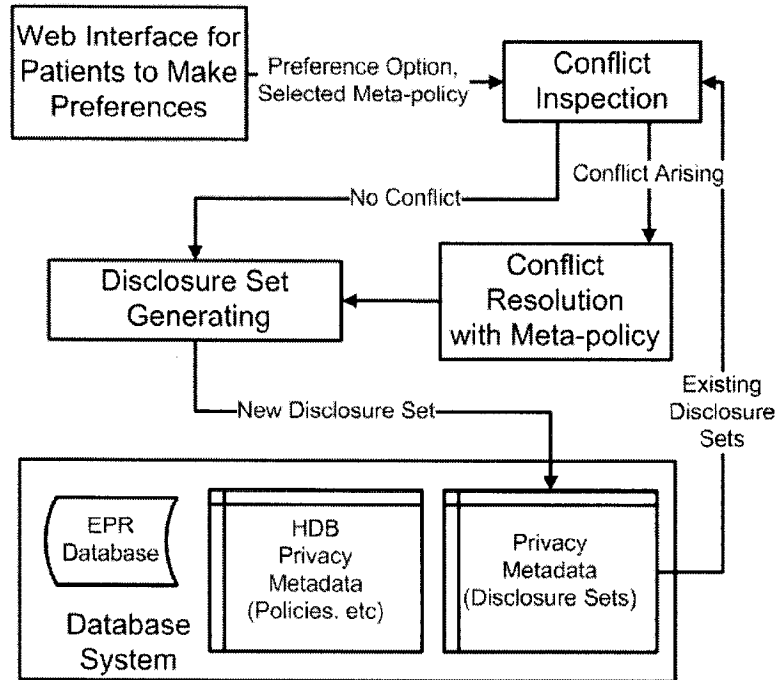


Figure 5.5: Detailed Preference Negotiation Architecture

P3P Compact Policies, HTML format Privacy Statements. It includes instructions and examples, technical support, staff review of the P3P implementation, and P3P Policy updates. P3PBuilder is a P3P policy generator that creates privacy policies to the W3C specification, and it's also web based, easy to use.

For creating APPEL preferences, a Java-based editor: JRC APPEL Preference [Cen] can assist developer to create preferences. The APPEL RULE can be created either by selecting a set of the predefined rules, or by using an advanced mode that provides options to create new rules.

JRC P3P Proxy [Cen] is an intermediary agent which handles the access control for remote web servers based upon the existing privacy preferences. It will be used for the conformance checking between policies and preferences.

All these three tools with different goals will be combined for the service development and

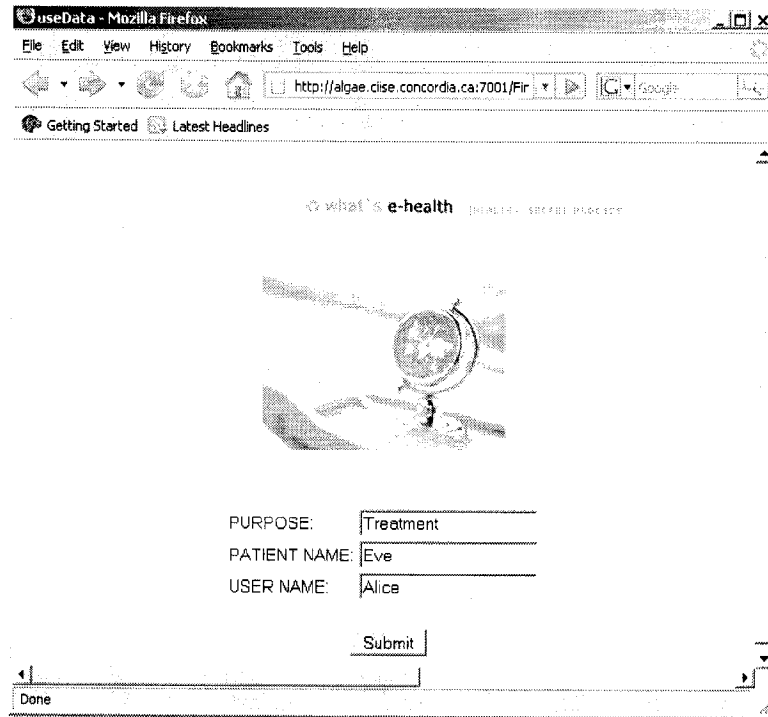


Figure 5.6: Web Interface for EPR Service to Use Private Data

data processing. They perform efficiently in those two architectures, preference negotiation and limited disclosure, which are based on the server-centric architecture of P3P.

## 5.4 A Scenario for Electronic Patient Record Data Access

This section presents a data usage scenario to introduce the privacy protection. The interface is used by various data recipients, such as doctors, nurse, etc.

From the scenario, a doctor (Alice) makes use of the web application, a specific EPR service to read the records of a patient, Eve. She will be authenticated firstly with her credential and purposes, and only the authorized data for her and the matched purposes can be accessed by Alice (a simple example for the web interface is shown in Figure 5.6).

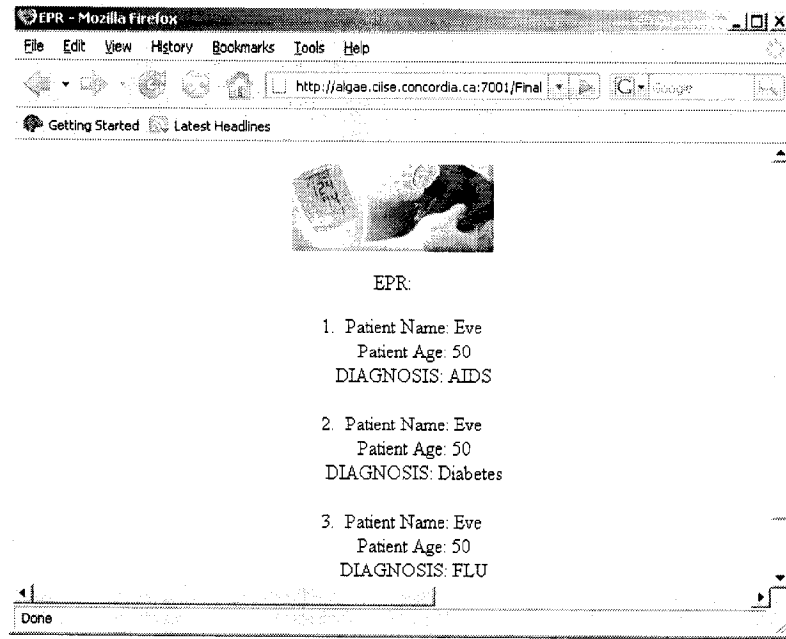


Figure 5.7: Electronic Patient Records Data Access Based on Authorization

If Eve authorizes doctors to access some of his private data (e.g. Patient Name, Age, and Diagnosis), Alice is able to access those data as a doctor. The results of the queries are presented in Figure 5.7. Note that all the records for the same patient (Eve) are created at different time and they are totally different. If Eve hasn't authorized doctors to access the requested data, Alice will be refused to view Eve's records using that query. However, if there's no matched records, but Alice is still authorized, an exception will occur that no records are found.

In addition, for the fine-grained authorization in this scenario, only the authorized tuples in the multi-dimensional EPR schema and the authorized attributes (private data) for Alice with matched purposes can be used by her.

A serial of EPR services including *"Accessing Private Data by Doctor"*, *"Access Private Data by Patient"*, *"Create EPR by Doctor"* are integrated in the e-Health architecture for

single sign-on (SSO) with different roles in the system.

## Chapter 6

# CONCLUSION AND FUTURE WORK

This thesis addressed the issue of preserving patients' privacy in e-Health systems. It proposed to allow patients to specify their privacy preferences in APPEL, which will be recorded as authorization policies in the backend HDB. Doctors are then authorized against such policies when their applications request for privacy data. It identified the lack of fine-grained authorizations as a limitation of the original HDB design and provided a solution based on a modified schema. It extended the HDB design to support the multi-dimensional model so it can be used to preserve privacy in analytical data applications. Moreover, it proposed to simplify the task of specifying privacy preferences by exploiting hierarchies inherent to various dimensions of a preference. It also described several options for users to specify a preference based on the hierarchies. It studied how to generate the effective disclosure set in each case and how to resolve potential conflicts that may arise between preferences with pre-defined meta-policies. It showed how hierarchies can be leveraged to simplify the representation of authorizations, and experimental results justify the design. It is my belief that the proposed solution can be integrated into existing e-Health systems to provide patients with better privacy protection.



## 6.1 Conclusion

Our group established a web-based e-Health portal system and my research concentrated on the privacy-preserving issues in it. There are many different roles in such a system and only two different kinds of them for privacy protection: one type to specify preferences and the other type to use data. Therefore, two kinds of architectures are required: one for preference negotiation and the other one for data use based on limited disclosure.

With my privacy-preserving model, e-Health system is capable of providing fine-grained authorization (record level) and hierarchical preferences (for more convenient preference-specifying). Furthermore, my preference negotiation architecture enables patients to specify fine-grained authorizations and preferences with hierarchies and select meta-policy to resolve conflicts. My data access architecture enables authorization checking, query processing (rewriting) to limit the data users' accesses. My metadata schemata provides support for fine-grained authorization with less storage and hierarchical preferences with better representation for hierarchies along with less storage on disclosure sets, compared with original HDB metadata model.

In short, my research shows that it is possible removes two limitations of the HDB model to make it more broadly applicable, and it proposes a novel method for more user-friendly specification of privacy requirements.

## 6.2 Future Work

With those completed contributions, I finished this thesis. What's more in the future, we can do some extra work on a few aspects that has been mentioned in this thesis.

1. Firstly, we can integrate P3P policies and APPEL preferences into my privacy-preserving model by not only building the conformance checking between policies and preferences, but also setting up matching between P3P policies to HDB policies and APPEL preferences to HDB preferences.
2. Secondly, this privacy-preserving model can be established through more mathematical computations and definitions. It is much easier to evaluate and calculate with some formal definitions.
3. Finally, I can present more comparisons and evaluations between the privacy solutions on (web-based) e-Health systems and other e-Commerce systems if possible (upon the analysis upon some other commercial systems). It contributes more to our whole e-Health system architecture design and privacy-preserving model.

# Bibliography

- [10496] United States Public Law 104-191. Health insurance portability and accountability act. 1996.
- [ABF<sup>+</sup>04] R. Agrawal, R. Bayardo, C. Faloutsos, J. Kiernan, R. Rantzaou, and R. Srikant. Auditing compliance with a hippocratic database. *Proc. of the 30th Int'l Conf. on Very Large Databases (VLDB)*, 2004.
- [AJ07] R. Agrawal and C. Johnson. Securing electronic health records without impeding the flow of information. *International Journal of Medical Informatics*, 2007.
- [AJSX04] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. Implementing p3p using database technology. *Proc. of the 19th International Conference on Data Engineering (ICDE)*, 2004.
- [AKSX02a] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. Hippocratic databases. *Proc. of the 28th Int'l Conference on Very Large Databases (VLDB)*, 2002.
- [AKSX02b] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. Server-centric p3p. *Proc. of the W3C Workshop on the Future of P3P*, 2002.
- [And96] R. J. Anderson. A security policy model for clinical information systems. *Proc. of the IEEE Symposium on Security and Privacy*, 1996.
- [AS00] R. Agrawal and R. Srikant. Privacy-preserving data mining. *Proc. of the ACM SIGMOD Conference on Management of Data*, 2000.
- [BA05] R. J. Bayardo and R. Agrawal. Data privacy through optimal k-anonymization. *Proc. of the 21st International Conference on Data Engineering (ICDE'05)*, 2005.

- [BCK03] S. Byers, L. F. Cranor, and D. Kormann. Automated analysis of p3p-enabled web sites. *Proc. of the 5th International Conference on Electronic Commerce (ICEC)*, 2003.
- [Bis05] G. Bisson. e-health portal and snomed for a more personalized integrated ehr. *Proc. UM2005 Workshop on Personalization for e-Health*, 2005.
- [BMA<sup>+</sup>01] G. Brelstaff, S. Moehrs, P. Anedda, M. Tuveri, and G. Zanetti. Internet patient records: new techniques. *Journal of Medical Internet Research*, 3(1):e8, 2001.
- [Cen] JRC P3P Resource Centre. Appel implementation. <http://p3p.jrc.it/>.
- [CLM02a] L. Cranor, M. Langheinrich, and M. Marchiori. A p3p preference exchange language 1.0 (appel 1.0). 2002.
- [CLM<sup>+</sup>02b] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler Marshall, , and J. Reagle. The platform for privacy preferences 1.0 (p3p 1.0) specification. 2002.
- [Con] World Wide Web Consortium. P3p implementations from w3c. <http://www.w3.org/P3P/implementations>.
- [Eys01] G. Eysenbach. What's e-health? *Journal of Medical Internet Research*, 3(2):e20, 2001.
- [fJWSE] Products for Java Web Service Evaluation. <http://www.infoworld.com/article/...>
- [Inca] BEA Inc. Java page flow. <http://sybooks.sybase.com/onlinebooks>.
- [Incb] IBM Inc. Ibm db2 administration guide.
- [Incc] IBM Inc. Ibm hippocratic database active enforcement user guide version 1.0 -including technical appendices.
- [Incd] Microsoft Inc. Implementing row- and cell-level security in classified databases using sql server 2005. <http://www.microsoft.com/technet/prodtechnol/sql/2005/multisec.mspx>.
- [Ince] Sybase Inc. Sybase adaptive server enterprise 12.5, system administration guide, row level access control. <http://sybooks.sybase.com/onlinebooks>.

- [IY05] K. Irwin and T. Yu. Determining user privacy preferences by asking the right questions: An automated approach. *Proc. of the 2005 ACM workshop on Privacy in the electronic society*, 2005.
- [JSSS01] S. Jajodia, P. Samarati, M.L. Sapino, and V.S. Subrahmanian. Flexible support for multiple access control policies. *ACM Transactions on Database Systems*, 26(4):1-57, 2001.
- [JZY<sup>+</sup>06] B. Jiang, D. Zhang, C. Yang, G. Yue, and F. Yu. A privacy policy of p3p based on relational database. *Proc. of the 1st International Multi-Symposiums on Computer and Computational Sciences (IMSCCS'06)*, 2006.
- [KBR00] B. Kolman, R. C. Busby, and S. C. Ross. *Discrete Mathematical Structures*. 2000.
- [KG06] G. Kaur and N. Gupta. E-health: A new perspective on global health. *Journal of Evolution and Technology*, 15(1), 2006.
- [KKST97] R. Kirkgoze, N. Katic, M. Stolba, and A. Min Tjoa. A security concept for olap. *Proc. of the 8th International Workshop on Database and Expert Systems Applications*, 1997.
- [Kyt99] T. Kyte. Fine-grained access control. *Oracle Corporation*, 1999.
- [LHL<sup>+</sup>07] S. Lu, Y. Hong, Q. Liu, L. Wang, and R. Dssouli. Access control in e-health portal systems. *Proc. 4th IEEE International Conference on Innovations in Information Technology (Innovations'07)*, 88-92, 2007.
- [LWHS06] J.-Gil Lee, K.-Young Whang, W.-Shin Han, and I.-Yeol Song. Hippocratic xml databases : A model and an access control mechanism. *Journal of Computer Systems Science and Engineering*, 21(6), 2006.
- [MMZ] F. Massacci, J. Mylopoulos, and N. Zannone. Minimal disclosure in hierarchical hippocratic databases with delegation. *Proc. of the 2005 Computer Security-ESORICS*, 3679.
- [NB03] A. Nanda and D. K. Burleson. *Oracle privacy security auditing: includes federal law compliance with HIPAA, Sarbanes-Oxley and the Gramm-Leach-Bliley Act GLB*. 2003.

- [Pen] N. Pendse. The olap report - what is olap.  
*<http://www.olapreport.com/fasmi.html>.*
- [Rep] UCI Machine Learning Repository. Adult dataset.  
*<http://archive.ics.uci.edu/beta/datasets/Adult>.*
- [San95] R. S. Sandhu. Issues in rbac. *ACM Workshop on Role-Based Access Control*, 1995.
- [SCFY96] R. S. Sandhu, E.J. Coyne, H.L. Feinstein, and C.E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38-47, 1996.
- [SdCdV04] P. Samarati and S. de Capitani di Vimercati. Access control: Policies, models, and mechanisms. *Springer Berlin / Heidelberg*, 2171/2001, 2004.
- [SLXN04] S. Sung, Y. Liu, H. Xiong, and P. Ng. Privacy preservation for data cubes. *Proc. of the 20th International Conference on Data Engineering (ICDE)*, 2004.
- [SS98] P. Samarati and L. Sweeney. Generalizing data to provide anonymity when disclosing information. *Proc. of the 7th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 1998.
- [Wik] E-health from wikipedia. *<http://en.wikipedia.org/wiki/E-health>.*

# Glossary

<b>APPEL</b>	A P3P Preference Exchange Language: A standard language to express users' preferences on their personal data based on P3P specification
<b>BCS</b>	Bio Medical Cognitive Science
<b>BPM</b>	Business Process Management: An merging field of knowledge and research at the intersection between managment and IT
<b>CIM</b>	Clinical Information Management
<b>CHI</b>	Consumer Health Informatics
<b>CORBA</b>	Common Object Request Broker Architecture: A standard enables software components written in multiple computer languages and computers to work together (Defined by Object Management Group)
<b>Closed Policy</b>	Authorizations specify permissions for an access
<b>DCOM</b>	Distributed Component Object Model: A communication technology among distributed software components across networked computers (By Microsoft Inc.)

<b>E-Health</b>	Healthcare practices supported by electronic processes and communication, such as Electronic Patient Records, Telemedicine .etc
<b>EPR</b>	Electronic Patient Record: The patients' digital records, including personal information, diagnosis, prescription .etc
<b>ETL</b>	Extract, Transform, and Load: a process in establishing data warehouses, extracting data from outside sources, transforming it to fit business requirements, and loading it to data warehouses
<b>FASMI</b>	Fast Analysis of Shared Multi-dimensional Information: Another concept to describe OLAP
<b>FGAC</b>	Fine-Grained Access Control: Implement access control with security policies or functions in a fine-grained level
<b>HIPAA</b>	Health Insurance Portability and Accountability Act: was enacted by the US Congress in 1996
<b>HDB</b>	Hippocratic Databases: a set of technologies that manages disclosure of electronic health records in compliance with data protection laws without impeding the legitimate flow of information
<b>HTML</b>	Hypertext Markup Language: a markup language for web pages
<b>JRC</b>	European Union's Joint Research Center



<b>JPF</b>	Java Page Flow: A feature set built upon a Struts-based Web application programming model
<b>JSP</b>	Java Server Pages: A Java technology that enables to dynamically generate HTML, XML or other file types for response upon the Web client request
<b>MVC</b>	Model-View-Controller: A multi-tier architecture pattern used in software engineering
<b>MLP</b>	Medical Language Processing
<b>OLAP</b>	Online Analytical Processing: An advanced for of data analysis for data warehousing environments
<b>Open Policy</b>	Authorizations specify denials for an access
<b>P3P</b>	Platform for Privacy Preferences: A specification allows users and websites to declare privacy preferences and policies in a machine-readable format (developed by the World Wide Web Consortium (W3C))
<b>REST</b>	Representational State Transfer: A style of software architecture for distributed hypermedia systems
<b>RPC</b>	Remote Procedure Call: A technology that enables a computer program to make a procedure execute in another address space without explicitly coding the details

<b>RLS</b>	Row-Level Security: Another term for Oracle's implementation of Fine-Grained Access Control, besides VPD
<b>Snowflake Schema</b>	Arranging tables in a relational databases and resembling the database schema into a snowflake shape. Since the speed of data retrieval (queries) is more important than data insertion, star schema and snowflake schema are commonly found in data warehouses
<b>Star Schema</b>	A database schema commonly found in data warehouses consisting a fact table referencing any numbers of dimension tables
<b>SQL</b>	Structured Query Language: A database computer language designed for data retrieval, database schema creation, modification or access control mechanisms
<b>SOA</b>	Service-Oriented Architecture: a software architecture that uses loosely coupled software services to support the requirements of business processes and software users
<b>SOAP</b>	Simple Object Access Protocol: A protocol for exchanging XML-based messages over computer networks, which is the basis of web services
<b>TEL</b>	Telemedicine: A developing technology for clinical medicine where medical information is transferred via telephone, Internet for remote medical diagnosis or examination

<b>VPD</b>	Virtual Private Database: Oracle's implementation of Fine-Grained Access Control
<b>W3C</b>	World Wide Web Consortium: The main international standards organization for the World Wide Web
<b>WSDL</b>	Web Services Description Language: An XML-based language which provides standard format to define web services
<b>WSRP</b>	Web Services for Remote Portlets: A network protocol standard designed for communications with remote portlets (defined by OASIS)
<b>XML</b>	Extensible Markup Language: A general purpose standardized markup language, which can be defined in many platforms and various systems