

available at www.sciencedirect.comjournal homepage: www.elsevier.com/locate/cose
**Computers
&
Security**


Patient-centric authorization framework for electronic healthcare services[☆]

Jing Jin^a, Gail-Joon Ahn^{b,*}, Hongxin Hu^b, Michael J. Covington^c, Xinwen Zhang^d

^aDeutsche Bank Global Technologies, Cary, NC, USA

^bArizona State University, 699 S. Mill Ave, Tempe, AZ, USA

^cIntel Corporation, Hillsboro, OR, USA

^dSamsung Information Systems America, San Jose, CA, USA

ARTICLE INFO

Article history:

Received 13 April 2010

Received in revised form

12 August 2010

Accepted 5 September 2010

Keywords:

Electronic Health Records(EHRs)

Patient-centric authorization

Selective sharing

Policy composition

Policy anomaly analysis

ABSTRACT

In modern healthcare environments, a fundamental requirement for achieving continuity of care is the seamless access to distributed patient health records in an integrated and unified manner, directly at the point of care. However, Electronic Health Records (EHRs) contain a significant amount of sensitive information, and allowing data to be accessible at many different sources increases concerns related to patient privacy and data theft. Access control solutions must guarantee that only authorized users have access to such critical records for legitimate purposes, and access control policies from distributed EHR sources must be accurately reflected and enforced accordingly in the integrated EHRs. In this paper, we propose a unified access control scheme that supports patient-centric selective sharing of virtual composite EHRs using different levels of granularity, accommodating data aggregation and privacy protection requirements. We also articulate and address issues and mechanisms on policy anomalies that occur in the composition of discrete access control policies from different data sources.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

In much of the developed world, healthcare has evolved to a point where patients can have many different providers – including primary care physicians, specialists, therapists, and even alternative medicine practitioners – to address their diverse medical needs. It is not uncommon for patients to visit providers who are physically separated from one another; some are located across town, while others are across the country or on another continent. As a result, medical records can be found scattered throughout the entire healthcare sector. From the clinical perspective, delivering proper patient care requires access to integrated and unified patient

information that is often collected in real-time to ensure the freshness of time-sensitive data. Yet the data dispersion in current healthcare settings typically results in painstaking, time-consuming efforts to obtain a patient's complete medical history, or unnecessary duplication of tests and other investigations. There is a strong need to create an infrastructure that uniformly integrates this heterogeneous collection of medical data and delivers it to the healthcare professionals who need it at the point of care (IEEE-USA's Medical Technology Policy Committee Interoperability Working Group, 2006). The adoption of standardized Electronic Health Records (EHRs) (Gates and Slonim, 2003; Ciena, 2008) has become an extremely important prerequisite for

[☆] A preliminary version of this paper appeared under the title "Patient-centric Authorization Framework for Sharing Electronic Health Records," in Proc. of the 14th ACM Symposium on Access Control Models and Technologies, Stresa, Italy, June 2009.

* Corresponding author.

E-mail address: gahn@asu.edu (G.-J. Ahn).

0167-4048/\$ – see front matter © 2010 Elsevier Ltd. All rights reserved.

doi:10.1016/j.cose.2010.09.001

bringing interoperability and effective data integration to the healthcare industry.

Effective management of EHRs is a very complex and sensitive issue. Patient privacy concerns, along with threats that could expose medical information, highlight the need for security and privacy technologies that are well-integrated into healthcare systems and enforceable across a variety of heterogeneous systems and networks. In particular, a shared EHR involves a complex composition of sensitive information, including patient demographic details, medical histories, laboratory test results, radiology images (X-rays, CTs), and so on. There is also a strong need for protection models that comply with legal and regulatory policies, while simultaneously ensuring that access to sensitive information is limited only to those entities who have a legitimate need-to-know privilege allowed by patients. For instance, a patient may opt to explicitly hide his medical information pertaining to an HIV/AIDS diagnosis from general medical information sharing session unless a specific treatment option is indicated. It is, therefore, essential to seek a secure, usable, and straightforward mechanism that allows patients to quickly and easily authorize a variety of medical affiliates to access their sensitive records in whole or partially. In addition, as a patient's medical records are distributed at different sites and virtually aggregated at the point of care, such an access control mechanism must be uniformly applied not only to the EHRs residing at each local site, but also to the aggregated EHR that is generated and shared on-the-fly. In this paper, we propose an access control model for selectively sharing such EHRs. We refer to the aggregated EHR with complex data and policy composition features as the virtual composite EHR and the local data sources that contribute to the virtual composite EHR as EHR instances. A key characteristic of our model is that we formulate the semantics and structural composition of EHRs in a hierarchical structure, where internal data objects are distinguished and associated with properties to address important criteria for medical data sharing such as data types, intended purposes and information sensitivities. Such hierarchical structure is further explored with a filtration mechanism, called authorization zone, that provides a flexible and efficient means to select and authorize a portion of EHRs to be shared with specific property criteria.

Considering the complexity of data aggregation and policy composition with multiple distributed EHR instances, we need a more sophisticated information model, as well as a unified policy scheme, to uniformly regulate selective sharing of EHRs at different levels of granularity. Such an access control model is the first contribution of this paper. In addition, the dynamic aggregation of distributed EHR instances requires seamless integration of access control policies from multiple data sources. Our second contribution is to articulate and propose mechanisms that identify and resolve potential policy anomalies for composed access control policies at the virtual composite EHR level. Finally, a virtual composite EHR sharing system is designed and implemented for integrated and federated healthcare networks, where a patient consent mechanism is incorporated to control and provide only authorized "views" of patient medical information to requesters.

The rest of the paper is organized as follows. In Section 2, we provide a brief overview of some related work. In Section 3, we present our unified patient-centric authorization framework. Section 4 discusses our policy anomaly analysis approach and

a policy evaluation mechanism is further elaborated in Section 5. Our prototype EHR sharing system is described in Section 6. Section 7 concludes this paper with future research directions.

2. Related work

2.1. EHR standards

There are several standards currently under development to structure and specify the clinical content of an EHR for the purpose of exchange, such as openEHR and HL7 Clinical Document Architecture (CDA) (openEHR Community, 2010; Dolin et al., 2004). openEHR uses a two-level methodology to model the EHR structure. In the first level, a generic reference model is designed to express the generic data content needed in clinical contexts and it also provides an explicit representation of the semantic and vocabulary of all EHR instances. In the second level, the notion of archetype is introduced to model specific healthcare concepts such as blood pressure and lab results. These archetypes are the fundamental building blocks to form the contents in various clinical EHR instances. Similarly, HL7 V3 Standards define an underlying Reference Information Model (RIM) that forms the generic information domain used across all HL7 messages, while CDA defines detailed structure and semantics of medical documents in terms of a set of coded components (called vocabulary) to model basic medical concepts.

2.2. Access control for EHR systems and e-Consent

A number of solutions have been proposed to address security and access control concerns associated with EHR systems (Eyers et al., 2006; Becker and Sewell, 2004; Bhatti et al., 2006). All of these approaches, to some extent, utilize role-based access control (RBAC) to address organizational security management requirements and authorize access to various healthcare parties. However, selective sharing of composite EHRs requires clear understanding of the internal clinical information and their structural relationships. None of these approaches took into account the structure and semantics of composite EHRs, and thus cannot support a more fine-grained access control to selectively share EHRs in whole or only partially. In this paper, we focus on the "selective" feature of an EHR system where a logical structure of a composite EHR is captured with its internal data elements being clearly distinguished and organized, so that our access control policies can be specified to authorize any portion of an EHR for data sharing.

Achieving privacy preservation in medical information sharing is also a critical concern for an EHR system. Several purpose-based access control models have been proposed recently to protect sensitive data (Byun et al., 2005; Yang et al., 2007). These models associate the intended purpose information with a given data element, and access is granted when the access purpose is consistent with the data element's intended purpose. However, as healthcare is such a complex domain involving various parties with different duties and objectives, the purpose-based access control alone cannot meet all the patient's privacy protection requirements. In this paper, we incorporate more applicable factors beyond purpose to control the selective sharing of EHRs.

To enable the patient control of medical information sharing, “e-Consent” mechanisms have been proposed to allow patients to issue or withhold authorization policies as electronic consents to those who wish to access their electronic health information (Coiera and Clarke, 2004; Ruan and Varadharajan, 2003; O’Keefe et al., 2005; Pritts and Connor, 2007). Several consent models with associated consent templates have been identified (Coiera and Clarke, 2004; Ruan and Varadharajan, 2003), and a few e-Consent based systems have been built upon these guidelines (Pritts and Connor, 2007; O’Keefe et al., 2005). However, it is still essential to develop a systematic approach to determine how a patient’s consent is expressed and at what granularity the consent is applied to the EHRs. Meanwhile, with dispersed EHR instances across many caregivers, it is also required for a patient to manage his consents in a unified and consistent manner within a shared EHR environment.

2.3. Policy anomaly discovery and resolution

A number of policy analysis tools have been introduced with the goal of detecting policy conflicts. A tool called Firewall Policy Advisor (Al-Shaer et al., 2003) was proposed to detect pairwise conflicts in firewall policies. Yuan et al. (2006) presented FIREMAN, a tool to check for policy misconfigurations through static analysis. These approaches for firewall policies cannot be directly applied to our policy analysis at both the EHR-instance level and the aggregation level. The resolution of policy conflicts also remains as an important issue. Some work presented general conflict resolution methods for access control in various areas (Fundulaki and Marx, 2004; Jajodia et al., 1997; Moses, 2005). In this paper, we propose a strategy chain to achieve more complete and effective conflict resolution while accommodating the features from these approaches.

3. Patient-centric authorization framework

3.1. Unified logical EHR model

A patient’s EHRs are typically dispersed over a wide range of distributed clinical systems and data structures. As suggested in dbMotion (2008), a Unified Data Schema (UDS) can be specified for all EHR instances so that a unified medical record can be maintained without the need to be adapted for these different environments. Similar to the generic reference models in openEHR and HL7, UDS defines generic semantics and logical relationships between data elements drawn from medical domains such as patient demographics, labs, medications, encounters, imaging and pathology reports, and a variety of other medical domains from primary, specialty and acute care settings. Based on these predefined categories, EHR instances are aggregated and integrated into a unified patient record as a virtual composite EHR.¹ In Fig. 1, a virtual

¹ Since data integration is not the focus of this paper, we do not consider heterogeneity in schema integration and assume all EHR instances and the corresponding aggregated virtual composite EHR uniformly conform (or are converted to conform) to a predefined UDS.

composite EHR aggregates two EHR instances from hospitals h_1 and h_2 based on a simple UDS defining three categories, Demographics, History and Labs.

In our model, both EHR instances and the aggregated virtual composite EHR are uniformly modelled as a labelled hierarchical structure. The nodes represent the clinical data elements that need to be protected for sharing. Their relations are captured as the association links between the nodes within the hierarchy. Each node is associated with specific properties to address essential features in term of the sources of data and their sensitivity levels. The properties can be categorized into three dimensions: origin, sensitivity, and object type. The origin property is specified to indicate the source(s) of data within the composition. As patient information may be duplicated in multiple EHR instances, such data elements should be merged as one element within a virtual composite EHR. We also use multiple origins to indicate such data merge, while a node with a single origin indicates that the data is unique from the respective origin. Using the category of Illness history in Fig. 1 as an example, the information of asthma comes from both h_1 and h_2 , while HIV information is uniquely from h_2 . The sensitivity property is designed to label a node based on the sensitivity of the content contained in it, which can be eventually used to prevent the patient’s sensitive medical information from being disclosed unintentionally. In the practice of Iowa HISPC (Iowa Foundation for Medical Care, 2007), the sensitivity classifications of medical data include general medical data, drug and alcohol treatment, substance abuse treatment, mental health, communicable disease (HIV, STDs, etc.), decedent, immunizations, and so on. Based on these classifications, data elements representing the patient’s HIV history and CD4 lab test should be marked with a property of “communicable disease” (“HIV” for simplicity). Finally, the object type property gives another dimension on data selection and protection. The nodes can be primitive types such as plain texts, dates and images. They can also be a composite type in the hierarchical structure including other types of data nodes. Formally, an EHR can be uniformly modelled and defined as follows:

Definition 1. (Logical EHR Model) An EHR composition is a tuple $C = (v_c, V_o, E_o, \tau_{V_o})$, where

- v_c is the root representing the whole EHR object;
- V_o is a set of nodes within the composite structure;
- $E_o \subseteq V_o \times V_o$ is a set of links between nodes; and
- $\tau_{V_o} : V_o \rightarrow P$ is a node labelling function to specify the property of a node. P is a set of properties defined in Definition 2.

Definition 2. (Property) Let O , S , and T be the sets of data origins, sensitivity classifications, and object types, respectively. And let $n = |V_o|$ be the number of nodes in an EHR composition C .

- $P_o = \{po_1, \dots, po_n\}$ is a collection of origin sets, where $po_i \subseteq O$ is a set of origins associated with a node, $i \in [1, n]$;
- $P_s = \{ps_1, \dots, ps_n\}$ is a collection of sensitivity classification sets, where $ps_i \subseteq S$ is a set of sensitivity classifications associated with a node, $i \in [1, n]$; and
- $P = P_o \times P_s \times T$ is a set of three dimensional properties of origin, sensitivity, and data type.

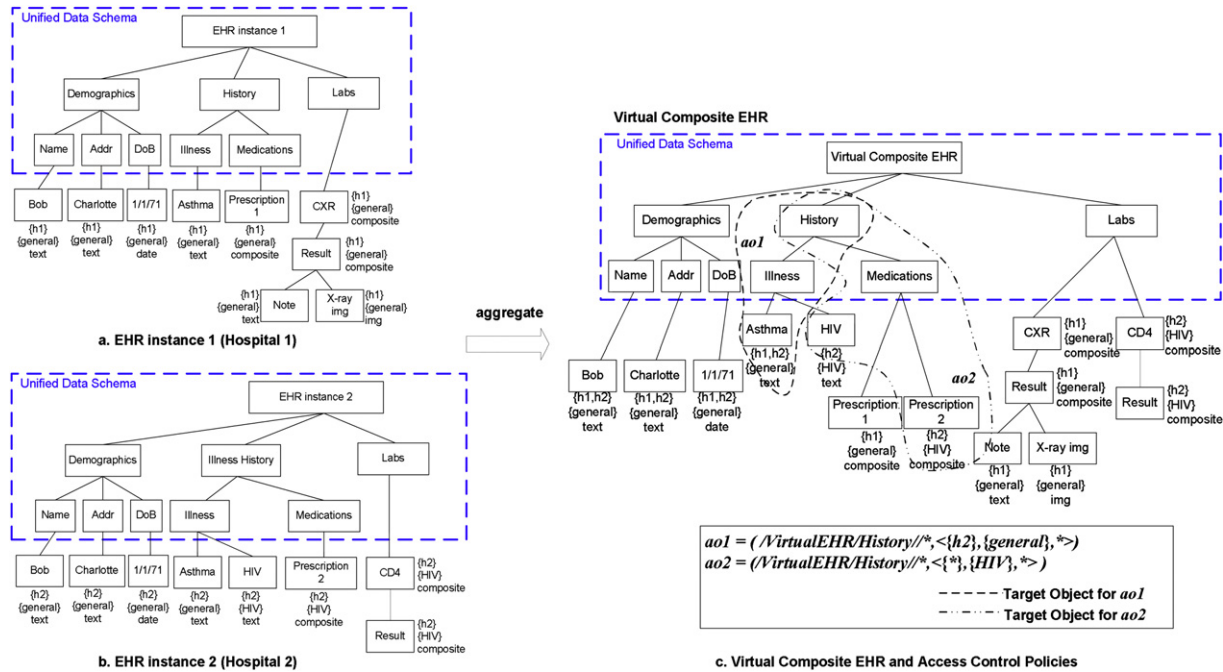


Fig. 1 – Virtual composite EHR in a hierarchical structure.

Given a node $v_i \in V_o$ inside an EHR composition C , the function $\tau(v_i) = p$ is defined to retrieve the property label p for the node. And we use the dot notation to refer to a specific property dimension. For instance, $p.p_o$ refers to the data origin property; $p.p_s$ refers to the sensitivity property; and $p.t$ refers to the object type. Within a logical EHR structure, nodes can be explicitly denoted by their identifiers, or can be implicitly addressed by means of Path Expressions. We apply an XPath-like expression for the path representation. Table 1 describes the notions and examples we use to select nodes inside a virtual composite EHR illustrated in Fig. 1(c).

3.2. Policy specification

To enable an authorized and selective sharing of patients' EHRs, it is essential for an authorization policy to determine a subject's access privileges for specific portion(s) of an EHR composition. Our policy specification scheme is built upon the identified logical EHR model so that access policies can be effectively defined at different granularity levels within the structure. In this paper, we assume that EHR instances are virtually

aggregated at the point of care for a practitioner to review, therefore we mainly focus on read-only access permission.

3.2.1. Subjects

In the context of healthcare, some authorizations may be given by patients in relation to identified individuals. For instance, a patient may want to indicate the following intent: "Dr. Smith is allowed to access my medical data". In other circumstances, the authorization is for a role such as "general physician", "cardiologist", "nurse", and so on. As healthcare practitioners are always associated with certain organizations, such a unique property may also be articulated to further constrain the subject. Formally, the subject specification is defined as follows:

Definition 3. (Subject Specification) Let E , R and O be the sets of user IDs, roles, and origins, respectively. A subject sub is defined as a tuple $sub = (e,so)$ or $sub = (r,so)$, where $e \in E$, $r \in R$, and optional subject origin set, $so \subseteq O$. Overall, the subject set Sub is defined as $Sub = (E \times 2^o) \cup (R \times 2^o)$.

3.2.2. Objects and filtration property

In order to support a flexible selection of data objects in an EHR composition, we utilize XPath-like path expressions to specify the scope of data elements to which an authorization policy applies. Meanwhile, the filtration properties are defined to be compared with the property label of each node within the EHR, and only matched nodes are selected as the Target Objects of the authorization. We formally define these concepts as follows:

Definition 4. (Filtration Property) Let O , S , and T be the sets of data origins, sensitivity classifications, and object types,

Table 1 – Path expression for node selection.

Expression	Description	Example
<i>nodename</i>	Select the named nodes	CXR
/	Select the node through absolute path from root node	/EHR/Labs/CXR
//	Select the node through relative path	//Labs/CXR
*	Select all immediate children nodes	//Labs/CXR/*
//*	Select all descendant nodes	//Labs/CXR/**

respectively as defined in [Definition 2](#). A filtration property is specified as a tuple $prop = \langle po, ps, pt \rangle$, where $po \subseteq O$ is the filtration property for origins; $ps \subseteq S$ is the filtration property for sensitivity classifications; and $pt \subseteq T$ is the filtration property for object types.

Definition 5. (*Property Match*) Suppose $prop = \langle po, ps, pt \rangle$ is a filtration property specification, and $p' = \langle po', ps', t' \rangle$ is the property label of a node, the node matches the filtration property if the following conditions are satisfied:

1. $p'.po' \subseteq prop.po$;
2. $p'.ps' \subseteq prop.ps$; and
3. $p'.t' \in prop.pt$.

Definition 6. (*Object Specification*) Let scp_expr be a scope expression to denote a set of nodes within the composition, and $prop$ be a filtration property specification, the object selection specification is defined as a tuple $ao = \langle scp_expr, prop \rangle$. Given an EHR composition $C = \langle v_c, V_o, E_o, \tau_{V_o} \rangle$ and an object selection specification ao , we define a function: $select(C, ao) \rightarrow V_a$, where $V_a \subseteq V_o$, to select the matched nodes within the specified scope as the Target Objects.

In specifying filtration properties, we also allow patterns to be used. Pattern “*” is to indicate any value within a property dimension, and pattern “{” is to specify any set within a property dimension.

Example 1. The following are two examples of object selection specifications against the EHR structure in [Fig. 1\(c\)](#).

ao1: $ao1 = \langle \text{VirtualEHR/History//}^*, \langle *, \{general\}, \text{text} \rangle \rangle$; and
ao2: $ao2 = \langle \text{VirtualEHR/History//}^*, \langle *, \{HIV\}, * \rangle \rangle$.

Both object specifications select the same scope, History category in the virtual composite EHR. *ao1* selects the nodes that come from any origins with *general* level of sensitivity and *text* as an object type. *ao2* selects the nodes from any origins with *HIV* level of sensitivity and any object types. [Fig. 1\(c\)](#) also illustrates the target objects being selected as two dashed zones. In particular, *ao1* results in the Asthma node under Illness being selected, and *ao2* results in the nodes of HIV and Prescription2 being selected under Illness and Medications, respectively.

Purposes: To further address a patient’s privacy concerns in sharing his medical information, an attribute of “purpose” is necessary to be specified in the authorization policy to confine the intended purposes/reasons for data access in healthcare practice. According to [Dimitropoulos \(2007\)](#), business practices for health information exchange can be organized by eleven purposes including payment, treatment, research, and so on. Formally, the intended purpose is specified as follows.

Definition 7. (*Intended Purpose*) Let P be a set of purposes for business practices in a healthcare domain. And let m be the total number of authorizations in the system. The intended purpose set $P_p = \{pp_1, \dots, pp_m\}$ is a collection of possible intended purpose sets, where $pp_i \subseteq P$ specifies the intended purposes for a particular authorization, $i \in [1, m]$.

Access Control Policy: To summarize the above-mentioned policy elements, we introduce the definition of an access control policy as follows.

Definition 8. (*Access Control Policy*) An access control policy is a tuple $acp = \langle sub, ao, pp, effect \rangle$, where

- $sub \in Sub$ is a subject;
- ao is an object selection specification resulting in a set of nodes $V_a \subseteq V_o$ being selected as target objects;
- $pp \in P_p$ is the intended purposes; and
- $effect \in \{permit, deny\}$ is the authorization effect of the policy.

Example 2. Given *ao1* and *ao2* in [Example 1](#) and suppose a patient articulates following access control policies in *h1*:

P1: $\langle \langle GP, * \rangle, ao1, \{research\}, permit \rangle$;
P2: $\langle \langle SP, \{h1\} \rangle, ao2, \{treatment, research\}, permit \rangle$; and
P3: $\langle \langle \text{Dr. Butcher}, \{h1\} \rangle, ao2, \{treatment, research\}, deny \rangle$.

In **P1**, a patient allows all general practitioners (GP) to view his *general* medical history for a *research* purpose. In **P2**, the patient allows all specialists (SP) in *h1* to view his *HIV* history for *treatment* and *research* purposes. Suppose Specialist Dr. Butcher in *h1* is a relative of the patient, the patient defines **P3** to deny his access to the *HIV* data.

In healthcare practices, HIPAA regulations are widely adopted by healthcare practitioners in the United States, by default allowing healthcare providers to share clinical information without the individual’s explicit permission for treatment, payment and healthcare operations ([Pritts and Connor, 2007](#)). In addition, in order to accommodate the emergency situations, a “break-of-glass” policy (“BG” policy for simplicity) should be specified to allow staffs in emergency rooms to access the patient’s medical information without the patient’s explicit authorizations. Both the default HIPAA policy and “BG” policy can be specified conforming to our unified policy schema.

Example 3. The default policy and BG policy can be specified as follows:

P_D: $\langle \langle \text{HP}, \{*\} \rangle, \langle \{*\}, \{*\}, * \rangle, \{treatment, payment, HCO\}, permit \rangle$;
P_{BG}: $\langle \langle \text{ERStaff}, \{*\} \rangle, \langle \{*\}, \{*\}, * \rangle, \{treatment\}, permit \rangle$.

4. Policy anomaly analysis

The access control policies in our model can be uniformly specified both at the EHR-instance level within different origin domains and at the aggregation level within a virtual composite EHR. Yet bringing in both positive and negative policies raises potential conflicts, and the flexibility in data selection may result in possible policy overlaps. We generalize such critical issues as *policy anomalies*.

In practice, there are two stages for policy anomaly analysis. The first stage is during the policy specification phase when a patient inserts, modifies or removes a policy from a local policy set. Policy anomaly analysis should be conducted to highlight any potential anomalies that may be introduced in adding and updating policies. Such anomalies can be resolved by consulting the patient to make changes to certain policy specifications

against the detected anomalies. The second stage of policy anomaly detection is during the policy composition and enforcement phase when the EHR instances are aggregated and policies are evaluated against the virtual composite EHR. Redundancies in composite policies could be identified and removed immediately in this stage. However, considering the large number of policies associated with the virtual composite EHR, it may be very difficult, or even impossible, to manually resolve all conflicts. And without a priori knowledge on the patient's authorization requirements, we cannot automatically correct the conflicts either. In this case, a practical method of resolving policy conflicts is to identify which policy involved in a conflict situation should take precedence in policy evaluation. The details of conflict detection and resolution in the second stage are discussed in Section 5. In this section, we focus on the policy anomaly analysis.

4.1. Anomaly classification

Example 4. We additionally define two object selection specifications as

ao3: $ao3 = (/VirtualEHR/History/*,\{h2,\{*,*\}\});$ and
ao4: $ao4 = (/VirtualEHR/History/*,\{h2,\{HIV\},text\}).$

Assume the patient defines four policies in $h2$ as follows:

P4: $((SP,*, ao3, \{treatment, research\}, deny);$
P5: $((Dr. Jones,\{h2\}, ao2, \{research\}, permit);$
P6: $((SP,*, ao3, \{treatment, research\}, permit);$ and
P7: $((Dr. Jones,\{h2\}, ao4, \{treatment, research\}, deny).$

Given the above example, we elaborate the following policy anomalies:

1. **Contradictory:** Two policies are contradictory to each other if they have different effects (permit or deny) over the same subjects (sub), target objects (ao) and intended purposes (pp). This is often considered as a typical policy conflict. For example, **P4** is contradictory to **P6** as all specialists (SP) are permitted to access the data matching aos for $treatment$ and $research$ purposes in **P4**, but are explicitly denied in **P6**.
2. **Exception:** A policy is an exception of another policy if they have different effects, but one policy is the subset of the other. Suppose *Dr. Jones* is a specialist (SP). **P7** is then an exception of **P6**, since all specialists are allowed to view the data matching $ao3$ for $treatment$ and $research$ purposes (**P6**), except for *Dr. Jones* (**P7**). The exception is not necessarily a policy conflict as it is commonly used to exclude a specific access request from a general access permission.
3. **Correlation:** Two policies are correlated if they have different effects, but one policy intersects with the other. In this case, the intersection of the two policies is permitted by one policy, but denied by the other. This is considered as a partial policy conflict. For example, **P5** and **P7** fall in this category. The intersection of **P5** and **P7** is a data element of HIV history. **P5** permits *Dr. Jones* in $h2$ to access this data element for $research$ purpose, but **P7** denies it.

4. **Redundancy:** A policy is redundant if there is another same or more general policy available with the same effect. For example, **P7** is redundant since all specialists (**P4**), including *Dr. Jones* (**P7**), are already denied to access the data matching $ao2$, which is a subset of $ao3$.

4.2. Anomaly detection

According to Definition 8, an access control policy essentially determines a relation of $sub \times V_a \times pp$, where V_a is derived from ao . We further define such a relation as an Authorization Zone (AZ) and a utility notation $P_\alpha[\beta]$ is used to indicate a particular field of a policy. $P_\alpha[\beta]$ implies that a field (or a set of fields) β from a policy P_α . There are four possible relationships between the authorization zones of two access control policies.

- **Exactly Match (\vee_{EM}):** An authorization zone AZ_x determined by a policy P_x exactly matches another authorization zone AZ_y derived from a policy P_y , if and only if the fields of sub , ao and pp in P_x are equal to the corresponding fields in P_y . Formally, $\forall i : P_x[i] = P_y[i] \Rightarrow AZ_x \vee_{EM} AZ_y$, where $i \in F = \{sub, ao, pp\}$.
- **Inclusively Match (\vee_{IM}):** An authorization zone AZ_x determined by a policy P_x inclusively matches another authorization zone AZ_y derived from a policy P_y , if and only if the fields of sub , ao and pp in P_x do not exactly match but are a subset of the corresponding fields in P_y . Formally, $\forall i : P_x[i] \subseteq P_y[i]$ and $\exists j : P_x[j] \subset P_y[j] \Rightarrow AZ_x \vee_{IM} AZ_y$, where $i, j \in F$ and $i \neq j$.
- **Partially Match (\vee_{PM}):** An authorization zone AZ_x determined by a policy P_x partially matches another authorization zone AZ_y derived from a policy P_y , if and only if the fields of sub , ao and pp in P_x do not exactly and inclusively match, but have intersections with the corresponding fields in P_y . Formally, $\forall i : P_x[i] \cap P_y[i] \neq \emptyset$ and $\exists j : P_x[j] \not\subseteq P_y[j] \wedge P_y[j] \not\subseteq P_x[j] \Rightarrow AZ_x \vee_{PM} AZ_y$, where $i, j \in F$ and $i \neq j$.
- **Disjoint (\vee_{DJ}):** An authorization zone AZ_x determined by a policy P_x is disjoint with another authorization zone AZ_y derived from a policy P_y , if and only if the fields sub , ao and pp in P_x have no intersection with the corresponding fields in P_y . Formally, $\forall i : P_x[i] \cap P_y[i] = \emptyset \Rightarrow AZ_x \vee_{DJ} AZ_y$, where $i \in F$.

By formalizing the relationships between the authorization zones and their effects, we could detect the policy anomalies between two policies P_x and P_y , as follows:

1. $AZ_x \vee_{EM} AZ_y$ or $AZ_x \vee_{IM} AZ_y$, and $P_x[effect] = P_y[effect] \Rightarrow$ **Redundancy:** If authorization zones determined by P_x and P_y exactly match or inclusively match and P_x and P_y define the same effect, then P_x is a redundant policy.
2. $AZ_x \vee_{EM} AZ_y$ and $P_x[effect] \neq P_y[effect] \Rightarrow$ **Contradictory:** If authorization zones determined by P_x and P_y exactly match and P_x and P_y define different effects, then P_x and P_y are contradictory to each other.
3. $AZ_x \vee_{IM} AZ_y$ and $P_x[effect] \neq P_y[effect] \Rightarrow$ **Exception:** If authorization zones determined by P_x and P_y inclusively match and P_x and P_y define different effects, P_x is then regarded as an exception of P_y .
4. $AZ_x \vee_{PM} AZ_y$ and $P_x[effect] \neq P_y[effect] \Rightarrow$ **Correlation:** If authorization zones determined by P_x and P_y partially match and P_x and P_y define different effects, then P_x and P_y are correlated.

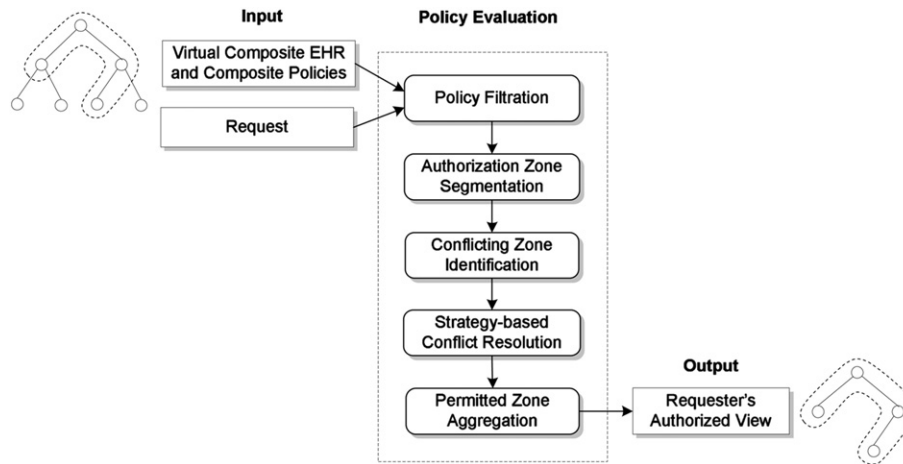


Fig. 2 – Policy evaluation process.

5. $AZ_x \vee_{PM} AZ_y$ and $P_x[effect] \neq P_y[effect]$ or $AZ_x \vee_{D_j} AZ_y \Rightarrow Normal$: If authorization zones determined by P_x and P_y partially match and P_x and P_y define the same effect, or authorization zones determined by P_x and P_y are disjoint, there is no anomaly between P_x and P_y .

5. Policy evaluation

Once authorization policies are specified, an authorization view of a virtual composite EHR can be computed through the policy evaluation. Meanwhile, conflicts in composite policies should be identified and resolved as well when enforcing policies to generate authorization views. As illustrated in Fig. 2, our policy evaluation mechanism computes a requester's authorization view with five steps: policy filtration, authorization zone segmentation, conflicting zone identification, strategy-based conflict resolution, and permitted zone aggregation. In the step one, all applicable authorization policies are selected by a policy filter from the policy pool based on an access request. These policies serve as the basis for deriving the requester's authorization view. In the step two, a policy-based segmentation technique is adopted to divide the entire authorization zone into disjoint segments. By identifying conflicting zones in these disjoint segments, conflicting policies can be determined in the step three. A strategy-based conflict resolution approach is then introduced in the step four to resolve all identified conflicts and generate the evaluation result. Finally in the step five all permitted zones are aggregated to yield the requester's authorization view. The details of our policy evaluation mechanism are illustrated in Algorithm 1.

5.1. Authorization zone segmentation and conflicting zone identification

Policy-based segmentation technique converts a list of policies into a set of disjoint authorization zones. As shown in lines 26–44 in Algorithm 1, a function called *Partition()* accomplishes this procedure. The function works by adding authorization zones z_p derived from each policy p to an

authorization zone set Z . A pair of authorization zones must satisfy one of the following relations: subset (line 31), superset (line 36), partial match (line 39), or disjoint (line 43). Therefore, one can utilize set operations to separate the overlapped zones from disjoint zones.

Definition 9. (*Conflicting Authorization Zone*) A conflicting authorization zone cz for a set of policies P is a collection of all nodes matching at least two policies that have different actions: *Permit* and *Deny*.

Conflicting zones are identified as shown in lines 6–9 in Algorithm 1. To illustrate our approach using the policy pool in Example 4, assume that a patient has removed P_4 from the policies of h_2 to resolve the contradictory conflict between P_4 and P_6 . Then, the policies from h_1 and h_2 defined by the patient are aggregated together along with the virtual composite EHR. In addition, suppose that Dr. Jones in h_2 sends a request to access this patient's EHR for a research purpose. The matched policies, P_1 , P_5 , P_6 and P_7 are then selected to generate the authorization view. Fig. 3 gives a representation of the zones derived from these four policies. We can notice that five unique disjoint zones are generated.² dz_1 is a denied zone defined by P_1 . pz_1 and pz_2 are two permitted zones derived from P_5 and P_6 , respectively. Moreover, two conflicting zones cz_1 and cz_2 are identified. They represent two policy conflicts, where a conflicting zone cz_1 is associated with two conflicting policies P_1 and P_6 , and a conflicting zone cz_2 is related to three conflicting policies P_5 , P_6 and P_7 .

5.2. Strategy-based conflict resolution

Once the conflicting zones are identified, the policy conflicts can be resolved by checking which policy involved in the conflict situation should take precedence in the policy

² For the purposes of brevity and understandability, we employ a two dimensional geometric representation for each zone. Note that an object selection specification in an access control policy of our model typically utilizes four fields to define the scope of object selection, thus a complete representation of authorization zone should be multi-dimensional.

Algorithm 1. (Policy Evaluation Algorithm)

```

Input: Virtual composite EHR,  $EHR$ ; Composite policies,  $CP$ ; Requester's ID,  $ID$ ; Requester's
purpose,  $PR$ .
Output: Requester's authorization view.
1 /* Policy filter */
2  $P \leftarrow Filter(CP, ID, PR)$ ;
3 /* Partition the entire authorization zone */
4  $S \leftarrow Partition(P)$ ;
5 /* Identify the conflicting zone */
6 foreach  $z \in Z$  do
7    $P' \leftarrow GetPolicy(z)$ ;
8   if  $\exists p_i \in P', p_j \in P', p_i \neq p_j$  and  $Effect(p_i) \neq Effect(p_j)$  then
9      $CZ.Append(z)$ ;
10 /* Identify the permitted zone */
11 foreach  $z \in Z$  do
12    $P' \leftarrow GetPolicy(z)$ ;
13   if  $\exists p_i \in P', p_j \in P', p_i \neq p_j$  and  $Effect(r_i) = Effect(r_j)$  and  $Effect(r_i) = "Permit"$ 
then
14      $PZ.Append(z)$ ;
15 /* Strategy-based conflicting resolution */
16 foreach  $cz \in CZ$  do
17    $P'' \leftarrow GetPolicy(cz)$ ;
18    $cz.Effect \leftarrow ApplyStrategy(P'')$ ;
19 /* Permitted zone aggregation */
20 foreach  $cz \in CZ$  do
21   if  $cz.Effect = "Permit"$  then
22      $AV \leftarrow AV \cup cz$ ;
23 foreach  $pz \in PZ$  do
24    $AV \leftarrow AV \cup pz$ ;
25 return  $GenerateView(EHR, AV)$ ;

26 Partition( $P$ )
27 foreach  $p \in P$  do
28    $z_p \leftarrow AuthorizationZone(p)$ ;
29   foreach  $z \in Z$  do
30     /*  $z_p$  is a subset of  $z$  */
31     if  $z_p \subset z$  then
32        $Z.Append(z \setminus z_p)$ ;
33        $z \leftarrow z_p$ ;
34       Break;
35     /*  $z_p$  is a superset of  $z$  */
36     else if  $z_p \supset z$  then
37        $z_p \leftarrow z_p \setminus z$ ;
38     /*  $z_p$  partially matches  $z$  */
39     else if  $z_p \cap z \neq \emptyset$  then
40        $Z.Append(z \setminus z_p)$ ;
41        $z \leftarrow z_p \cap z$ ;
42        $z_p \leftarrow z_p \setminus z$ ;
43    $Z.Append(z_p)$ ;
44 return  $Z$ ;

```

enforcement. We introduce the following strategies to resolve policy conflicts in our patient-centric EHR sharing system.

1. *Recency-overrides*: This strategy states that newer authorizations prevails older ones. As the authorization requirements may change over a period of time, a patient may specify certain policies in one origin, and he might define new policies in other origins along with his changed

authorization requirements. Obviously, when aggregating these policies together, newer policies should take precedence to respond an access request that matches policies from different origins. However, this strategy may not always resolve the policy conflicts since no authorization wins when two conflicting policies have the same timestamps. We call a strategy that cannot always derive a solution as a *nondeterministic* strategy.




Fig. 3 – Authorization zone segmentation.

2. *Specificity-overrides*: This strategy states that a more specific authorization overrides a more general authorization. As we discussed earlier, *exception* in composite policies is not necessarily a conflict, but is often used to exclude specific parts of a larger set for a certain effect. This strategy is also nondeterministic if there is no exception in composite policies.
3. *Deny-overrides*: This strategy indicates that the decision for an access request is “Deny” if any matched policy’s evaluation results in “Deny”. As the fundamental reason of policy conflict is that two policies specify different effects, it implies that at least one policy defines an effect as “Deny”. Therefore this strategy is deterministic and always returns “Deny”. In order to protect the patient’s privacy and confidentiality in healthcare systems, it is plausible to apply such a strict strategy in some special cases.

Since *recency-overrides* and *specificity-overrides* strategies are nondeterministic, and *deny-overrides* strategy is too restricted in general for conflict resolution, it is possible to combine these strategies together to achieve a more effective conflict resolution. Fig. 4 demonstrates a composite strategy for policy conflict resolution that sequentially applies *recency-overrides*, *specificity-overrides*, and *deny-overrides* strategies.

In the example depicted in Fig. 3, if we assume that the policies defined in $h2$ is newer than the policies defined in $h1$, the conflicting authorization zone cz_1 can be resolved by applying *recency-overrides* strategy, where $P6$ takes precedence over $P1$ by granting SP an access to the data covered by the conflicting authorization zone cz_1 .

Regarding the conflict represented in cz_2 , the conflicting policies $P5$, $P6$ and $P7$ are defined in $h2$ with the same timestamps. Therefore, *recency-overrides* strategy cannot be applied but *specificity-overrides* strategy can be applied to resolve this conflict by denying an access to the data covered by cz_2 since $P7$ is an exception of both $P5$ and $P6$.

5.3. Permitted zone aggregation

To generate the requester’s authorization view (AV), all permitted zones should be aggregated together. The following




Fig. 4 – Composite strategy for policy conflict resolution.

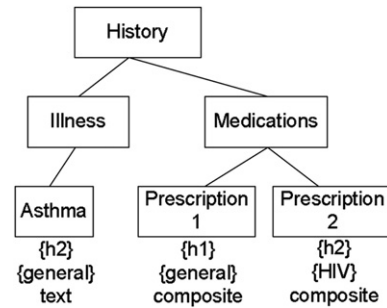


Fig. 5 – Dr. Jones’s authorization view.

relation is used to aggregate all permitted zones for constructing the authorization view:

$$AV = CZ^P \cup PZ \quad (1)$$

CZ^P indicates the conflicting authorization zones with “permit” evaluation result when applying conflict resolution strategies to solve associated conflicts. PZ denotes all permitted zones originally derived from the selected policies.

Using the same example shown in Fig. 3, the conflicting zone cz_1 belongs to CZ^P . Therefore, the requester’s authorization view can be derived as $AV = cz_1 \cup pz_1 \cup pz_2$. Fig. 5 shows an authorization view of the virtual composite EHR returned to the requester Dr. Jones in $h2$. The view contains a subset of nodes in the virtual composite EHR, including three nodes, Asthma node under Illness, and Prescription1 and Prescription2 nodes under Medications.

6. EHR sharing system

As part of our ongoing research efforts, we have designed and implemented a proof-of-concept system, called InfoShare, which is a simplified clinical information sharing system that utilizes our proposed model for a patient to control accesses of his medical information. In particular, InfoShare collects a patient’s access control policies as patient consents, and uses these consents to selectively share the patient’s medical information as contained in the virtual composite EHR with different requesting parties through a point of service (POS) web application.

Architecturally, health information as EHR instances generally is maintained and managed at each geographically distributed healthcare provider’s local site with the notation of federation. The data elements in the EHR instances should be associated with special properties for data filtration and authorization purposes. We therefore establish an EHR data labelling system in the federation to facilitate the unified property labelling for the healthcare providers’ EHR instances. Fig. 6(a) illustrates the overall architecture of an integrated InfoShare information system.

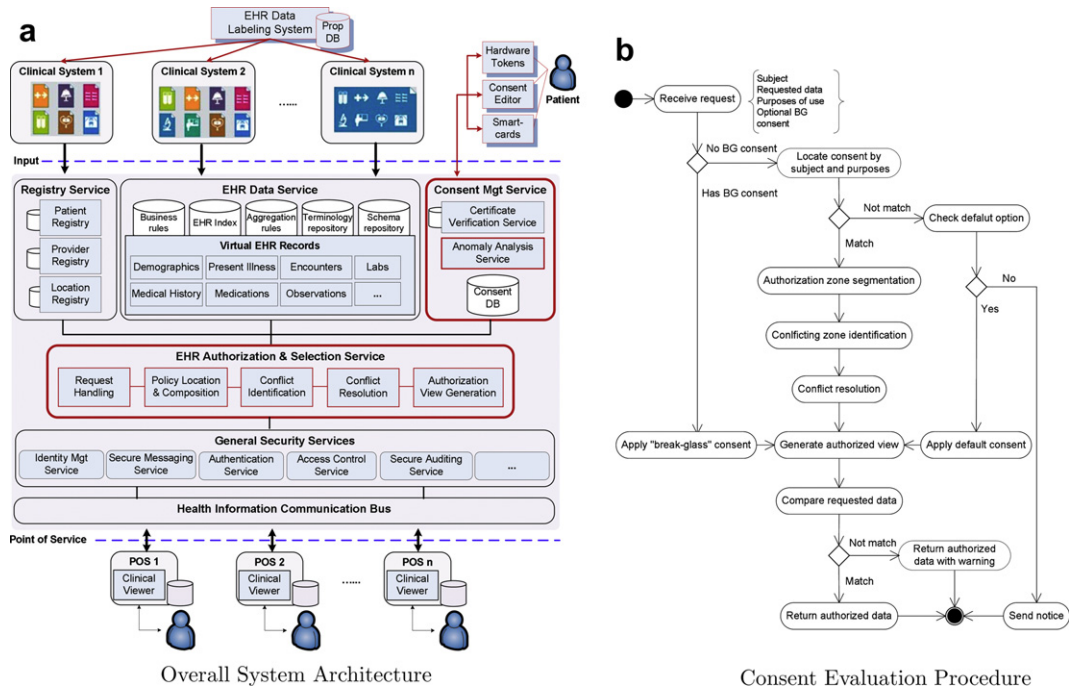


Fig. 6 – InfoShare system.

As a general clinical information sharing system, the Registry Service, EHR Data Service, General Security Service and Health Information Communication Bus are common system components to achieve the required functionalities of secure data retrieval, virtual composite EHR creation and communication with requesting POS applications. Especially, we inject the Consent Management Service and EHR Authorization & Selection Service as the major system modules to convey the core features of our proposed approach. In particular, the Consent Management Service enables the patient control by collecting and analyzing the patient’s access control policies encapsulated in consents. A web-based consent editor tool is implemented to facilitate a patient to edit his policy consents, and interact with the Consent Management Service for the patient to store and update his consents. Also, anomalies in policy consents can be pointed out in this consent editor tool. The EHR Authorization & Selection Service is responsible to handle the data access requests, and composite and evaluate relative access control policies to derive the authorization view for the requester. Conflicts can be identified and resolved in the policy evaluation.

There are three types of consents in the system: the patient’s specific consents, the default consent, and the “BG” consent, where the default consent and “BG” consent specify the default policy P_D and “BG” policy P_{BG} , respectively as shown in Example 3. The precedence order of evaluating these consents is defined as $BG_consent \geq patient_consent \geq default_consent$. We have articulated our policy evaluation approach for patient’s specific consents in Section 5. Fig. 6(b) illustrates the procedures for the EHR Authorization & Selection Service to handle access requests and derive the authorized data to be shared, considering all three types of consents. An access request includes information of the requester subject, the requested data, the intended purposes of use, and an optional “BG”

consent in emergency situations. The “BG” consent has the highest priority in execution, therefore such consents are directly evaluated to get the authorized data. In other situations, the authorization service interacts with the Consent Management Service to locate the related patient consents based on the specified subject and intended purposes. If certain matched consents are located, the policies are aggregated and evaluated to derive the authorized portion of data within a virtual composite EHR. If there are no patient consents being located, the default consent is evaluated to derive the authorized data. After the authorized data portion is determined, it is compared with the requester’s requested data and only matched data portion is returned to the requester. If the requester is not authorized to access all the data he requested, the requester is notified with a warning, so that the requester may further ask for new patient consents to access the data for the need of his practice. Such an effective mechanism is utilized to balance the data integrity concern of the practitioners and the privacy concern of the patients for shared EHRs.

In terms of implementation details, the XML-based HL7 Clinical Document Architecture (CDA) (Dolin et al., 2004) is utilized in our *InfoShare* system for the formal representation of EHR instances as well as the virtual composite EHR. We use *Jaxe XML editor* as the EHR data labelling service to associate properties with data elements in CDA EHRs. We implement the patient consents as X.509 attribute certificates (Housley et al., 2002), where the access control policies are encapsulated as attributes within the certificate. We utilize Ordered Binary Decision Diagrams (BDDs)³ to represent policies and

³ BDD has been demonstrated as an efficient data structure to deal with a variety of policy analysis, such as policy conflict detection in firewall (Yuan et al., 2006) and XACML policy verification (Fisler et al., 2005).

Table 2 – Experimental result for policy evaluation process.

# of Matched Policies	Preprocessing Time (ms)	# of Segments	# of BDD Nodes	# of Conflicting Segments	Processing Time (ms)
3	11	3	7	0	46
7	14	13	21	2	67
12	19	26	39	5	82
18	23	34	47	6	105
25	32	45	62	10	123
31	38	78	92	13	159

perform various set operations required by the zone segmentation algorithm, such as unions (\cup), intersections (\cap), and set differences (\setminus). A Java-based BDD library, called *JavaBDD* (2007), is employed in our implementation. The *InfoShare* system employs a Java Servlet based web portal as the POS application for a healthcare practitioner to query and view the authorized medical information of a patient.

As discussed earlier, policy evaluation is the core functionality of the EHR Authorization Service in our *InfoShare* System. Thus, the efficiency of our policy evaluation approach should be evaluated. In our experiments, we built a policy pool with 200 composite policies along with a virtual EHR dataset constructed based on our unified logical EHR model. The experiments were carried out on a desktop PC running Windows XP SP2 with 3.25 GB RAM and 3.00 GHz Intel Core 2 Duo CPU. By randomly triggering the policy evaluation process with different access requests, we measured the response time for each request. Since time required by the policy evaluation process highly depends upon the number of matched policies for a request, we selected six representative samples, which are shown in *Table 2*, with respect to the different number of matched policies. The preprocessing time of each request in this table indicates the time for locating and aggregating matched policies. The processing time for a request includes the time for segmenting authorization zone, the time for identifying conflicting authorization zones, the time for resolving conflicts, and the time for aggregating permitted zones to generate the requester's authorization view. *Table 2* also summarizes the information for the policy segmentation including the numbers of all generated segments, the numbers of constructed BDD nodes, and the numbers of conflicting segments. From *Table 2*, we observe that our policy evaluation approach performs fast enough to handle an access request matched a large number of policies. For example, evaluating a request with 31 matched policies only takes totally around 200 ms to generate corresponding authorization view in our experiments.

7. Conclusion and future work

In this paper, we proposed an innovative approach to supporting selective sharing of virtual composite EHRs. The access control policies are specified around the unified logical EHR model, taking into consideration of critical issues such as distributed data integration and privacy protection concerns. We also proposed a mechanism to identify and resolve policy anomalies in the process of policy composition. Our approach has been demonstrated in a proof-of-concept prototype

InfoShare system that applies e-Consent mechanism to enable the patient-centric medical information sharing with different parties in the healthcare environments.

For the future work, rigorous experiments need to be conducted to evaluate the performance and storage efficiency of our *InfoShare* system. Meanwhile, a variety of analytical and empirical methods from the area of usability study could also be adopted to investigate usability of our system.

Acknowledgments

This work was partially supported by the grants from National Science Foundation (NSF-IIS-0900970 and NSF-CNS-0831360) and Department of Energy (DE-SC0004308 and DE-FG02-03ER25565).

REFERENCES

- Al-Shaer E, Hamed H. Firewall policy advisor for anomaly discovery and rule editing. In: *Integrated network management, 2003*. In: *IFIP/IEEE eighth international symposium*; 2003. p. 17–30.
- Becker MY, Sewell P. Cassandra: flexible trust management, applied to electronic health records. In: *Proc. of IEEE 17th computer security foundations workshop*; 2004. p. 139–54.
- Bhatti R, Moidu K, Ghafoor A. Policy-based security management for federated healthcare databases (or RHIOs). In: *Proc. of the international workshop on healthcare information and knowledge management*; 2006. p. 41–8.
- Byun JW, Bertino E, Li N. Purpose based access control of complex data for privacy protection. In: *Proc. of 10th ACM symposium on access control models and technologies (SACMAT)*; 2005. p. 102–10.
- Ciena. The national health information network creating a new vision. In: *White Paper, healthcare information and management systems society (HIMSS) conference 2008*; 2008.
- Coiera E, Clarke R. e-Consent: the design and implementation of consumer consent mechanisms in an electronic environment. *Journal of the American Medical Informatics Association* 2004; 11(2):129–40.
- dbMotion. White paper: the critical role of integrated patient information in the delivery of high quality healthcare; January 2008.
- Dimitropoulos LL. Privacy and security solutions for interoperable health information exchange: interim assessment of variation executive summary, http://www.rti.org/pubs/avas_execsumm.pdf; July 2007.
- Dolin RH, Alschuler L, Boyer S, Beebe C, Behlen FM, Biron PV. HL7 clinical document architecture, release 2.0. ANSI Standard; 2004.

- Eyers DM, Bacon J, Moody K. OASIS role-based access control for electronic health records. In IEEE proceedings – software; 2006. p. 16–23.
- Fisler K, Krishnamurthi S, Meyerovich LA, Tschantz MC. Verification and change-impact analysis of access-control policies. In ICSE '05: Proceedings of the 27th international conference on software engineering; 2005. p. 196–05.
- Fundulaki I, Marx M. Specifying access control policies for XML documents with XPath. In: Proceedings of the ninth ACM symposium on access control models and technologies; 2004. p. 61–9.
- Gates C, Slonim J. Owner-controlled information. In: Proc. of the 2003 workshop on new security paradigms; 2003. p. 103–11.
- HL7. HL7 reference information model, http://www.hl7.org/Library/data-model/RIM/modelpage_mem.htm.
- Housley R, Polk W, Ford W, Solo D. Internet x.509 public key infrastructure certificate and certificate revocation list (crl) profile. RFC3280, <http://rfc.net/rfc3280.html>; 2002.
- IEEE-USA's Medical Technology Policy Committee Interoperability Working Group, editor. Interoperability for the national health information network (NHIN). IEEE-USA EBOOKS; 2006.
- Iowa Foundation for Medical Care. HISPC state implementation project summary and impact analysis report for the state of Iowa, http://www.ifmc.org/news/StateImpactReport_11-27-07.doc; 2007.
- Jajodia S, Samarati P, Subrahmanian VS. A logical language for expressing authorizations. In IEEE symposium on security and privacy, Oakland, CA; May 1997. p. 31–42.
- JavaBDD, <http://javabdd.sourceforge.net/>; 2007.
- Jaxe XML editor, <http://jaxe.sourceforge.net/>.
- Moses T. eXtensible access control Markup Language (XACML), version 2.0, Oasis Standard. Internet, <http://docs.oasis-open.org/xacml/2.0/accesscontrol-xacml-2.0-core-spec-os.pdf>; 2005.
- O'Keefe CM, Greenfield P, Goodchild A. A decentralised approach to electronic consent and health information access control. *Journal of Research and Practice in Information Technology* 2005;37(2):161–78.
- openEHR Community. openEHR, <http://www.openehr.org>.
- Prittis J, Connor K. The implementation of e-Consent mechanisms in three countries: Canada, England, and The Netherlands. SAMHSA report, <http://ihcrp.georgetown.edu/pdfs/prittise-consent.pdf>; 2007.
- Ruan C, Varadharajan V. An authorization model for e-Consent requirement in a health care application. In: *Applied cryptography and network security*, LNCS, vol. 2846; 2003. p. 191–205.
- Yang N, Barringer H, Zhang N. A purpose-based access control model. In: Proc. of 3rd international symposium on information assurance and security (IAS); 2007. p. 143–8.
- Yuan L, Chen H, Mai J, Chuah C, Su Z, Mohapatra P, Davis C. Fireman: a toolkit for firewall modeling and analysis. In: 2006 IEEE symposium on security and privacy; 2006. p. 15.
- Jing Jin** received the Ph.D. degree at the College of Computing and Informatics, University of North Carolina at Charlotte, Charlotte. She was a member of the Laboratory of Information Integration, Security, and Privacy (LIISP), University of North Carolina at Charlotte. Her current research interests include access control and trust management, identity and privacy management, network and distributed system security, and security in health informatics.
- Gail-Joon Ahn** received the Ph.D. degree in information technology from George Mason University, Fairfax, Virginia, 2000. He is currently an Associate Professor in the School of Computing, Informatics, and Decision Systems Engineering and the Director of Security Engineering for Future Computing (SEFCOM) Laboratory at Arizona State University (ASU), Tempe. His current research interests include information and systems security, vulnerability and risk management, access control, and security architecture for distributed systems. His research has been supported by the U.S. National Science Foundation, National Security Agency (NSA), U.S. Department of Defense (DoD), U.S. Department of Energy (DoE), Bank of America, Hewlett Packard, Microsoft, and Robert Wood Johnson Foundation. Dr. Ahn is a recipient of the U.S. Department of Energy CAREER Award and the Educator of the Year Award from the Federal Information Systems Security Educators Association (FISSEA). He was an Associate Professor in the College of Computing and Informatics, and the Founding Director of the Center for Digital Identity and Cyber Defense Research, and Laboratory of Information Integration, Security, and Privacy (LIISP), University of North Carolina at Charlotte, Charlotte.
- Hongxin Hu** is currently working toward the Ph.D. degree at the School of Computing, Informatics, and Decision Systems Engineering, Arizona State University, Tempe. He is a member of the Security Engineering for Future Computing (SEFCOM) Laboratory, Arizona State University. His current research interests include access control models and mechanisms, network and distributed system security, secure software engineering, and security in social network and cloud computing.
- Michael J. Covington** received his Ph.D. and MScS degrees from the Georgia Institute of Technology's College of Computing in Atlanta, Georgia. He also holds a B.S. degree from Mount Saint Mary's College in Emmitsburg, Maryland.
- Xinwen Zhang** is a research scientist at Samsung Information Systems America at San Jose, CA. His research interests include security policies, models, architectures, and mechanism in general computing and networking systems. His recent research focuses on secure and trusted mobile platforms, applications, and services. He has a PhD in information technology from George Mason University, Fairfax, VA.