Stratified Modelling and Analysis of Confidentiality Requirements

Adeniyi Onabajo
Department of Computer Science
University of Victoria
Victoria, BC, Canada V8W 3A4
onabajo@cs.uvic.ca

Jens H. Weber-Jahnke
Department of Computer Science
University of Victoria
Victoria, BC, Canada V8W 3A4
jens@cs.uvic.ca

Abstract

In this paper we present a method for modelling and analyzing confidentiality requirements based on requirements stratification. Stakeholders with varying data usage concerns have confidentiality and privacy requirements, and these stakeholders are often in different jurisdictions, for example, national, provincial and local authorities. In addition, customers, such as patient groups and individual patients, have important confidentiality concerns which should be considered in the requirement engineering process. Our approach provides a method to model and analyze the interactions of the different requirements with their inherent stratified relationship and supports the iterative specification and analysis of the requirements. We report on a preliminary evaluation of the method with a case study in the health care domain. Our results show that our method is suitable to express most case study requirements in their natural stratification order, but it also uncovered important limitations. Nevertheless, our method was effective in detecting a potential incompleteness in the subject requirements set.

1. Introduction

Increases in the occurrence of data misuse and confidentiality breaches, which can have potential negative consequences to individuals or organizations have emphasized the need to consider security concerns at the beginning of the software development lifecycle. The negative impacts of leaking confidential data include financial loss perpetrated through identity theft, social stigmatization from disclosure of personal data such as health records, and threats to personal safety. The information age, facilitated by advancement in distributed computing and communication, is characterized by electronic data processing in various aspects of modern society such as online banking, government e-services, electronic medical systems, as well as social interactions over the internet. However, potential benefits such as prompt service delivery and ease of use, are overshadowed by data confidentiality concerns.

Government legislations, e.g., HIPAA [12] in the U.S and PIPEDA [21] in Canada, and industry regulations require organizations that use personal information to adhere to principles of information privacy and security. By extension the information systems used by the data custodians should also conform to these principles. However, government legislations merely represent a generic framework in which individual stakeholders or stakeholder groups can define more concrete security goals. Interactions of the different stakeholder security goals with each other or with legal framework can lead to inconsistencies, which need to be harmonized for secure system development.

The harmonization task can be challenging for complex applications, such as electronic patient records, due to the inherent difficulty of identifying all stakeholder confidentiality goals early on in the system lifecycle. These applications are open-ended in terms of: (1) breadth - new information is discovered or becomes relevant, (2) depth - finer-grained details are discovered or become relevant, and (3) complexity - new relationships are discovered or become relevant [22]. Hence, systems are often built with general confidentiality goals that represent "future-proof" requirements. For example, clinical information systems typically just differentiate between two different uses of patient health data: (1) clinical use (also called primary use) and (2) research use (also called secondary use). This does not give stakeholders any fine-grained control about information, e.g., "my GP can see my entire record but other clinicians cannot see the section on mental health, except for an emergency treatment".

1530-1605/08 $25.00 © 2008 IEEE

The importance of adequate requirement analysis in the software development process has been realized in the last few years and there are ongoing efforts to integrate security requirement analysis early in the development life-cycle [18]. Confidentiality requirements analysis needs to be started at the early phase of development and continued throughout the system life-cycle. This requires analysis in the presence of incompleteness - a feature of default reasoning [4].

Appropriate formalisms and methods should support a precise and concise representation of confidentiality requirements. Requirements elicitation and analysis is typically performed in an iterative process of specification, evaluation and refinement. As depicted in Fig. 1, each pass of the requirement phase provides an opportunity to consider new concerns and goals from stakeholders or to revisit existing requirements. Appropriate concerns and goals are then incorporated into subsequent phases of the development life-cycle. Methods should be able to identify inconsistencies early on and either resolve them or mark these as sources of incompleteness.

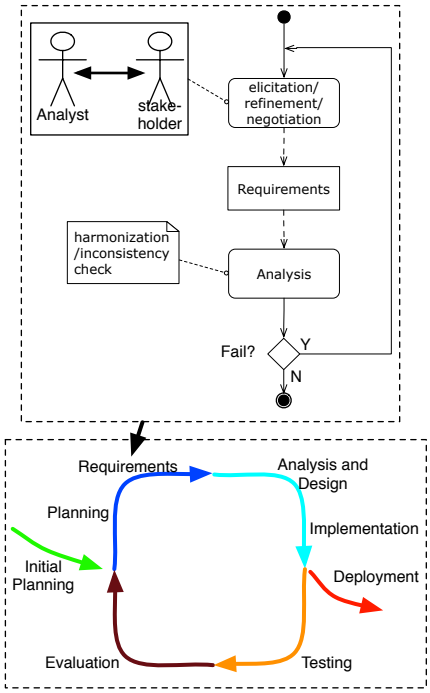

Figure 1. Confidentiality Requirement Engineering

In this paper we discuss a method, termed Confidentiality Requirement Elicitation and Engineering (*CREE*) that addresses the above-mentioned shortcoming and offers a stratified model for confidentiality requirements elicitation and analysis. Section 2 highlights methods for defining and analyzing requirements. Section 3 describes concepts and principles of a stratified and multilateral engineering of confidentiality requirements, while Section 4 is an overview of *CREE*. Formal definitions for analysis of *CREE* specifications are given in Section 5 and Section 6 describes an evaluation of the approach and its limitations. We conclude with some of the future work in Section 7.

2. Confidentiality Requirements in Practice

There have been different models to address security requirements in the software development process. Some provide extensions to the UML modelling language by adding security related features such as confidentiality and access control. For example, UMLSec [14] specifies confidentiality and integrity requirements in analysis models using UML. SecureUML [17] integrates Role-Based Access Control (RBAC) model in a model-driven process, with additional support for authorization constraints. These methods address security from a system-perspective but do not support social or organizational level modelling and analysis required for effective understanding of security issues in general and confidentiality in particular.

Goal oriented approaches to requirement engineering have also been adopted. A goal is an objective the proposed system should realize. Goals serve as sources for identifying requirements, and help to communicate these requirements to stakeholders [28]. Two frameworks which integrate goals and goal refinements in requirement models are KAOS and Tropos [27].

KAOS supports requirement analysis using concepts of objects (entity, relationship, event), operations (inputs/outputs over objects), agents (execute operations) and goals. Specification is done in a language, which consists of an outer declarative layer for conceptual modelling and an inner formal assertion layer for formal reasoning - using a theorem proving technique [29], [23].

Tropos is a framework aimed at developing an agent-oriented software engineering methodology, starting from early requirements to implementation [6]. It provides the notion of actors (e.g., agents and roles) and their associated goals and tasks for modelling and analysis. Formal Tropos provides formal analysis, and it has features of i* [16], complemented by a temporal specification. It uses model checking approach as verification technique [7]. Secure Tropos extends Tropos by providing features that enable modelling and analysis of security requirements [9]. This is done by integrating trust, security and system engineering. The

notions of ownership and offer of a service, as well as functional and trust dependencies are made explicit [10].

While these frameworks provide concepts for capturing the system, their analysis are designed with the objective of identifying inconsistencies in the classic sense. Therefore, they are "inconsistent (incompleteness) intolerant". Complete specification is not usually possible nor practical [25], hence support for evolving requirement specifications, even with inconsistency should be provided.

Some research investigating this notion has been done. [13] proposes a formal approach in the context of multi-perspective viewpoints for supporting continued reasoning in the presence of inconsistencies and tracking sources of the inconsistencies during analysis. For analysis it uses an adaptation of classical logic which allows continued reasoning in the presence of inconsistencies. Xbel [5] supports model checking using a family of multi-valued logics for reasoning over inconsistent multiple viewpoints. A framework based on default reasoning and belief revision for identifying, analyzing and managing inconsistency in evolving requirements is described in [8]. It parses requirements, specified in simple (controlled) natural language into logic formulae. Although it uses natural language in order to facilitate better communication with stakeholders, the restricted form of the language is often not expressive enough.

Our research focuses on applying the idea of tolerating incompleteness, specifically for stakeholder confidentiality concerns, in a stratified model. Analysis is provided by identifying relationships between requirements along the layers.

3. A Stratified Approach

A stratified approach for confidentiality requirements engineering addresses the inherent multilayered nature of confidentiality and privacy requirements. The requirements are layered from the general to the more specific e.g., based on political jurisdictions, and the analysis considers the degree of specificity allowed from one layer to another. We provide a precise stratified model using key concepts of confidentiality goals: (1) stakeholder who has the goal (2) actual data to be protected (3) target stakeholder who is granted or denied access to the data (4) the purpose for which access is granted or denied. Use case analysis, which describe scenarios for realizing different system functionalities, can be used as an initial step to identify the key concepts. Misuse cases can also provide insights to threats posed from an attacker's point of view [3].

3.1. Multilateral Viewpoints

The need for multiple viewpoints in requirements engineering has been highlighted by different research in the last few years [15], [26]. This is particularly important for complex and large-scale applications with different stakeholders, who have multi layered concerns. Relevant stakeholders need to be identified in modelling security requirements because it facilitates analysis of the potential vulnerabilities of the proposed system by considering threats posed from the stakeholders' interactions with the system [16].

A multilateral viewpoint allows the concerns and goals of the different stakeholders to be considered. For example, we can often distinguish between "data owners" and "data custodians", the latter being entrusted with data belonging to the former. Data owners have confidentiality goals of restricting data access by certain individuals or for some purposes to prevent misuse. Furthermore, parties not directly involved in the data collection and routine usage often have interests in the data and the usage. For example, in order to facilitate public safety, government regulations might override individual data usage directives during medical epidemic outbreaks or for law enforcement investigations.

A multilateral viewpoint requires identifying stakeholders' concerns with respect to specific data (or parts of data), individual(s) or agencies who are targets of the concerns, purposes for which the confidentiality concerns are associated and how long these concerns are for. However, complete knowledge of all the relevant stakeholders and their goals might not be available at any particular iteration of the requirement engineering phase. Hence, the confidentiality requirements are subject to modifications between iterations of requirements engineering. This is clearly the case in an evolutionary system development life-cycle, where there is continuous iteration.

Stakeholder relationships are important for confidentiality requirements because the goals for controlling data usage are not only for attackers, but for other stakeholders who do not necessarily have malicious intents. Non-permanent relationships among stakeholders, e.g., friendships, colleagues and family physician could be subjects of requirements and should be explicitly modeled for requirement specification.

An understanding of the data structures is crucial in realizing correct specification and analysis of stakeholders' confidentiality requirements. An abstract composition structure, usually hierarchical, can be used to model the data. It allows confidentiality requirements to be captured for specific data elements. However,

the constituent data elements might not be structurally distinguishable, making it difficult to have distinct requirements for individual sub elements. For example, the "Family History" section of a medical record comprises data from family members. Relationships among the data elements also have significant impact on requirements. Knowledge about a data element might be used as a basis for determining the existence of associated data due to the reduced uncertainty, e.g., a patient's medical history might be inferred from the medication history.

3.2. Stratified Viewpoints

Requirement analysis is a continuous process and iterative modelling is aimed at realizing a complete set of requirements. Iteration allows elicitation of new requirements or refinement of existing requirements. The refinements could be exceptions, which occur with respect to data elements, stakeholder(s), or purposes. For example, patient stakeholders could grant consent to clinicians, which comprise physicians and nurses, for their medical records but a refinement of this requirement could be to deny nurses access to medication history.

The existence of different stakeholders with multilayered confidentiality concerns is a motivation for stratified modelling and analysis. Privacy concerns in a national integrated health solution, such as the Canadian Health Infoway's [2] Electronic Health Record Infostructure (EHRi) involve many levels, including jurisdictions (federal, provincial and local authorities), stakeholder groups and individual patients, and this makes stratified modelling and analysis suitable.

A consequence of the stratified approach is the need for requirement representation that is non-monotonic because new refinements can contradict the existing set of requirements [4]. Stratified requirement modelling enables interpretation of requirements and their refinements (exceptions) at different layers. Analysis should identify possible conflicts among the levels. In addition, the approach enables specification of incomplete set of requirements which occur during the early iterations. Different stakeholders' requirements are included at different iterations and analysis of their interactions with existing requirements is performed.

4. CREE

CREE - Confidentiality Requirement Elicitation and Engineering, addresses issues of precise specification of confidentiality requirements from a multilateral perspective with a stratified method for model and analysis of the requirements. The method presented in this paper is a further development of the method described in [11]. This paper adds concepts and theory for stratified modelling and analysis of requirements.

4.1. Multilateral Requirement Definition

Fig. 2 shows the meta model containing the key concepts for modelling confidentiality goals. This was developed using various sources which provide details of confidentiality and privacy concerns in different domains, including healthcare and finance [20]. The sources included privacy legislations, publications on information security as well as interviews with practitioners.

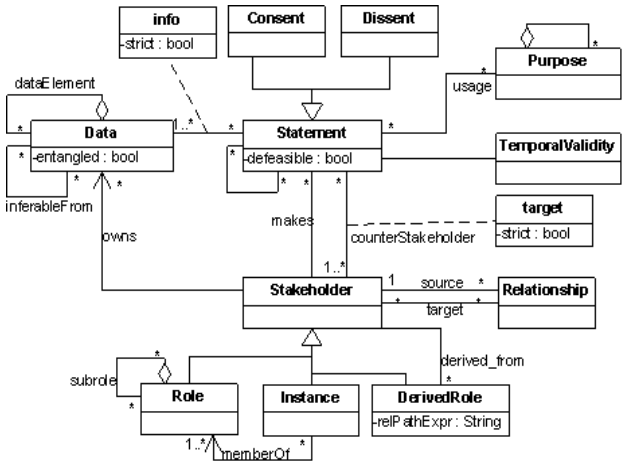

Figure 2. Confidentiality Properties

The following is a description of some of the concepts and how they address the issues of stratified multilateral confidentiality requirement engineering described in Section 3.

Stakeholder - relevant stakeholders are identified as individuals, which are represented as *Instance* or groups termed *Role*. The groups comprise instances or sub groups (sub roles). For example, the individual "J. Doe" can be a member of "Patients" role, while "Physicians" and "Nurses" are sub roles of the "Clinicians" role. A *Relationship* is a connection from a stakeholder (*source*) to other stakeholders (*targets*) and it represents social or organizational associations among stakeholders. *DerivedRole* is a stakeholder group based on relationships, e.g., derived role "Care Team" for "J. Doe" (the source of the relationships) is based on the relationships "Family GP" and "Home Nurse". This concept supports dynamic groups which might change as requirements evolve.

Data - this represents the object of concern. It is represented by a hierarchical structure, in which data elements can be further decomposed into sub elements. Data elements composed of indistinguishable data elements are indicated in the model with the *entangled* property. Requirements specified for entangled (semi-structured) data elements are applied to all the sub elements.

If a data element X can be inferred from Y, then knowledge about Y can be used to deduce the existence of X. This association can be a rationale for additional requirements to minimize undesired information leakage. The concept of inference between data elements is represented by the *inferableFrom* relationship between data elements.

Statement - this expresses the actual confidentiality requirement, and it can represent a permission (*Consent*) or a denial (*Dissent*) of access to data. A statement is attributed to a stakeholder with the requirement and is directed to one or more stakeholders, called *counterstakeholder*.

Purpose - represents functional usage of data. It is the objective for which the data is needed. Purposes might be associated to operational tasks used to realize the purposes. Purposes such as clinical encounter, emergency and research can be identified for the clinical domain.

Confidentiality goals (statements) are also associated with temporal validity, which indicates the length of time the statement is valid.

Visual Modelling

The conceptual model shown in Fig. 2 is used as a basis for diagrams to specify stakeholders, data compositions and stakeholders' requirements. The diagram types are:

i) Stakeholders diagram - this is used to model stakeholders and their relationships. A stakeholders diagram for the example above is shown in Fig. 3.

ii) Stakeholder information model diagram - this represents a stakeholder's data. It is a hierarchical model and each data element can be composed of other data elements. An example of a data model for the role "Patients" is shown in Fig. 4. Any of the data elements in a stakeholder's information model diagram can be the object of a confidentiality statement.

iii) Stakeholder statement diagram - this is used to depict a stakeholder's confidentiality requirements, and it uses elements from stakeholders and information model diagrams. Fig. 5 shows two requirements for the "Patients" stakeholder:

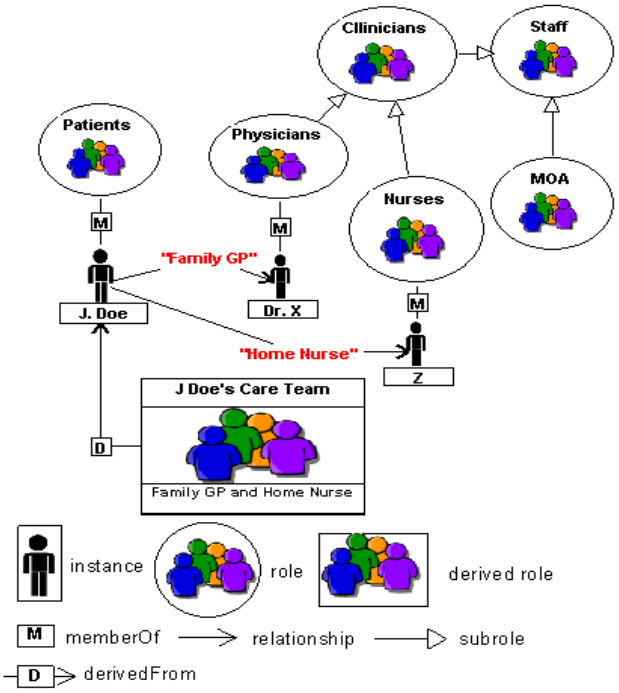

Figure 3. Stakeholders Diagram

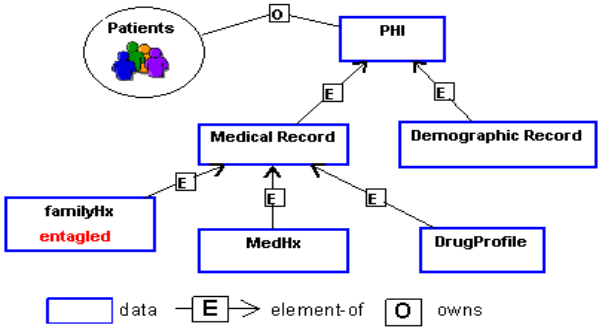

Figure 4. Information Model Diagram

r2.1 is a denial of access to personal health information (PHI), while r2.1.1 gives consent for demographic data. Both are directed at the "MOA" (Medical Office Assistant) role.

The diagrams support multilateral requirement specification by separately modelling each stakeholder's confidentiality viewpoints, which can be analyzed with other requirements for inconsistencies [27]. The different perspectives from the diagrams also allow easier communication with stakeholders, a key aspect of requirements engineering [19], by avoiding a presentation of overwhelming information.

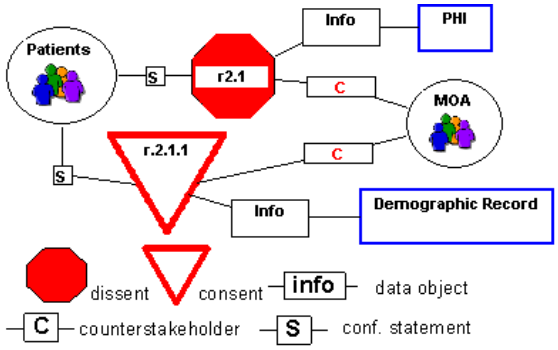

Figure 5. Statement Diagram

Table 1. Stratification with Data property

label	*stakeholder*	*data*	*statement type*	*counter stakeholder*
r1.1	Patients	PHI	consent	Nurses
r1.2	Patients	Family Hx	dissent	Nurses

Table 2. Stratification with Stakeholder property

label	*stakeholder*	*data*	*statement type*	*counter stakeholder*
r2.1	Patients	PHI	dissent	MOA
r2.2	J. Doe	PHI	consent	MOA
r2.3	J. Doe	Family Hx	dissent	Clinicians
r2.4	J. Doe	Family Hx	consent	Physicians

4.2. Stratified Requirement Definition

Requirement stratifications are a result of specifying more specific requirements based on the hierarchical data composition, stakeholder relationship structure and/or purpose hierarchy. A more specific requirement, at a lower level, is based on data sub elements, sub roles (instance-role membership) or sub purposes. This approach of specifying requirements is in line with system evolution. Each evolutionary phase provides increased knowledge that reveals new and/or relevant data elements, finer details or relationships [22]. While identifying and addressing all stakeholders' concerns is desired, this is often not practical The development process has to adapt to continuous discovery and refinements of stakeholders' concerns. For each iteration, additional requirements, can be unrelated, enforcements or exceptions to existing requirements.

Exceptions and enforcements are more specific to the particular context, hence, at a lower level in relation to general requirements. However, it may be desired that a confidentiality statement is not subject to exception refinements and should always be satisfied. This notion is represented by the *defeasible* property of a statement. A non-defeasible statement does not support exceptions.

For the rest of the paper, we omit the temporal validity property, but the discussions are generally applicable. *CREE* supports stratified requirement specification by allowing more specific requirements for data, stakeholders and purposes. This is illustrated with the following examples.

Data stratification - Table 1 shows two requirements: r1.1 specifies that patients generally give consent for nurses to access their personal health information (PHI). This is a default requirement for all patients. However, r1.2, which specifies a denial of access to the family history (Family Hx) section of the health record, is an exception to r1.1 - the data property of r1.2 is a sub element of the data for r1.1.

Stakeholder stratification - this occurs with stakeholders who make confidentiality statements or the counterstakeholders of the statements. The stakeholder for the refined requirement is either a member or a sub role of the role in a higher level requirement. In Table 2, r2.2 is a stratified redefinition of r2.1 because "J. Doe" is a member of "Patients" of requirement r2.1. "J. Doe" gives consent to the health record even though the general requirement for patients (r2.1) does not for the "MOA" role. Similarly, r2.4 is a stratification of r2.3 because 'Physicians" is a sub role of "Clinicians".

Purpose stratification - this is based on purpose refinement, in this case the purpose property of the refined requirement is a sub purpose of a higher level requirement. In Table 3, r3.2 is a stratified redefinition of r3.1 because "Emergency" is sub purpose of "Clinical Encounter".

Table 3. Stratification with Purpose property

label	*data*	*statement type*	*counter stakeholder*	purpose
r3.1	Family Hx	dissent	Nurses	Clinical Encounter
r3.2	Family Hx	consent	Nurses	Emergency

5. Formal Definition

In this section we present formal definitions of CREE requirements specifications.

Definition 1: A confidentiality requirement specification is a tuple $(D, S, P, R, \leq_D, \leq_S, \leq_P, \leq_R)$ with:

- D - a partially ordered set of data elements, where for d_1,$d_2 \in D$, $d_2 \leq d_1$ if d_2 is element-of d_1
- S - a partially ordered set of stakeholders (roles, instances, derived roles), where for s_1,$s_2 \in S$,

$s_2 \leq s_1$ if s_2 is specialization or instance member of s_1

- P - a partially ordered set of purposes, where for p_1,$p_2 \in P$, $p_2 \leq p_1$ if p_2 is a sub purpose of p_1
- R - a partially ordered set of statements of the form $r(s, type, def, o, d, p, c)$:
  - $s \in S$: stakeholder with confidentiality concern
  - $type$ = {consent(+), dissent(−)}: statement type
  - def = {T, F}: indicates if a statement is defeasibe
  - $o \in S$: the data owner
  - $d \in D$: data object of the statement
  - $p \in P$: purpose of the statement
  - $c \in S$: counterstakeholder of the statement

  and for r_1,$r_2 \in R$, $r_2 \leq r_1$ if $s_2 \leq s_1$, $o_2 \leq o_1$, $d_2 \leq d_1$, $p_1 \leq p_1$ and $c_2 \leq c_1$

Definition 2: A confidentiality requirement r_2 is a conceptual refinement of r_1 (denoted as $r_2 \leq r_1$) if $s_2 \leq s_1$, $o_2 \leq o_1$, $d_2 \leq d_1$, $p_2 \leq p_1$ and $c_2 \leq c_1$

Analysis of stratified confidentiality requirements is aimed at identifying conflicts and ambiguities. The requirement specification can be depicted as a tree structure where a node is a conceptual refinement (Definition 2) of its parent(s). The following definitions are used in the analysis of the stratified confidentiality requirements model.

Definition 3: A requirement r_2 is an *Exception* of r_1 if r_2 is a conceptual refinement of r_1, def_1 = T and $type_1 \neq type_2$.

Stratified modelling supports an open world assumption of evolving information which could be considered in the requirement specification. While exceptions need to be tolerated, identifying these can help in detecting specification errors (errors by analyst/stakeholder). Exceptions are permitted only for defeasible requirements.

Definition 4: A requirement r_2 is an *Enforcement* of r_1 if r_2 is a conceptual refinement of r_1, def_1 = T, def_2 = F and $type_1 = type_2$.

Stratified specification allows refinements which enforce defaults for a specific stakeholder, data or purpose. Although, this might be considered redundant, an enforcement supports the notion of specificity. The enforcing requirement is more specific and non-defeasible, thus would not allow any exceptions. Incompleteness is still tolerated because the enforced requirement can still be refined with exceptions.

Definition 5: A requirement r_2 is a *Conflict* with r_1 if r_2 is a conceptual refinement of r_1, def_1 = F and $type_1 \neq type_2$

A conflict in the stratified confidentiality model is similar to concept of conflict in traditional reasoning. This occurs with non-defeasible requirements, which do not allow exceptions.

Definition 6: r_1 and r_2 are *Ambiguous* if at least one of $s_1 \leq s_2$, $o_1 \leq o_2$, $d_1 \leq d_2$, $p_1 \leq p_2$, $c_1 \leq c_2$ holds, and least one of $s_2 \leq s_1$, $o_2 \leq o_1$, $d_2 \leq d_1$, $p_2 \leq p_1$, $c_2 \leq c_1$ holds, and $type_1 \neq type_2$

Stratified modelling might present a situation where two requirements refine some of each others properties, i.e., each of the requirements satisfies some of the conditions for conceptual refinements of the other. If these requirements have different types, then an ambiguity occurs.

6. Evaluation

We have evaluated our approach to stratified modelling and analysis of confidentiality requirements with a complex real-world case study in the area of electronic health information management. While a single case study does not allow us to draw general conclusions about the suitability of our approach, it has generated valuable insight and indicates the importance of stratified confidentiality requirements analysis. The case study investigates privacy requirements for establishing an electronic health record (EHR) infrastructure in the Canadian context. Canada is a federation of ten provinces and three territories. Funding and administration of health care services falls under the jurisdiction of these individual provinces and territories, which are further divided into different health authorities affiliated to a set of health care organizations (hospitals, clinics, etc.). Privacy legislation and policies exist on all jurisdictional levels and requirements are defined by various stakeholder groups, which may cross-cut these levels. The Assembly of First Nations is a stakeholder group with requirements that cross jurisdictional boundaries [24] .

In 2001 Canada Health Infoway (Infoway) [2] was formed with a mandate to facilitate the development and adoption of interoperable EHR solutions. Infoway has generated a Security & Privacy requirements specification and an architecture blueprint for the pan-Canadian EHR defining 28 privacy and 87 security requirements in natural language [1]. We have investigated the Infoway privacy requirements specification as well as several other multi-lateral view points, including federal, provincial, and clinic-wide privacy requirements.

6.1. Modelling

The translation of the multilateral viewpoints into CREE was possible for most of the requirements but we were not able to represent all of them (details will follow at the end of this section). We found that confidentiality requirements make frequent use of the term "should" to represent a recommendation and "must" to represent a strict obligation. Consider the following Infoway requirement (page 30 of the specification) as an example:

"Except where inappropriate (e.g. specifically exempted by law or professional code of practice), organisations connecting to the EHRi, and organisations hosting components of the EHRi should obtain the knowledge and consent of each patient/person for the collection, use or disclosure of his or her PHI - and where required by law, must- obtain the knowledge and consent of each patient/person for the collection, use or disclosure of his or her PHI."

We modeled such requirements as defeasible consent statements for collecting, using or disclosing personal health information (PHI). Individual provinces, health authorities or even patient groups would be able to refine this requirement. This is represented in our formal notation as:

$$R1 = (CanadaInfoway, consent, T, Patients, PHI, Legitimate, User)$$

Some jurisdictions have already decided to use "deemed" consent for any collection, use or disclosure of PHI for the purpose of caring for a patient's health. Other jurisdictions, such as Ontario's Personal Health Information Protection Act (PHIPA) require seeking explicit consent, which is covered by the second part of the above requirement. We have expressed requirements of the form "must not collect/use/disclose without consent" as a defeasible dissent against the collection/use/disclosure of data. The *CREE* formal notation is:

$$R2 = (Ontario, dissent, T, Patients, PHI, Legitimate, User)$$

Patients or groups of patients can again refine such a requirement by adding explicit consent. As an example, we took the NSMobile call group of physicians who share a subset of their individual GP patient records to provide 24/7 on call services. Patients wanting to join this group and receive this on-call service would have to give explicit consent for electronically collecting, using and disclosing their PHI among the call-group members and for the purpose of the service. We model such requirements as non-defeasible consent. The formal representation for this is as follows:

$$R3 = (OnCallPatients, consent, F, Patients, PHI, Treatment, User)$$

6.2. Analysis

After sorting all multi-lateral confidentiality requirement statements according to the partial order defined on data, stakeholders and purposes, we detected several instances of confidentiality requirements stratification. As an example, Fig. 6 shows the refinement relationship between the three confidentiality requirements discussed above. It shows NSMobile's patient strict consent as an exception to Ontario's defeasible dissent, which again is an exception to Infoway's defeasible consent statement.

The analysis also revealed a potential incompleteness of the Infoway requiremens specification: The federal Personal Information Protection and Electronic Documents Act (PIPEDA) requires the individual's consent for collection of personal information except for specific cases such as if the "collection is clearly in the interests of the individual and consent cannot be obtained in a timely way". We translated this requirement as a (defeasible) dissent for the collection of personal information, which can be redefined by individual stakeholders.

$$R4 = (CanadaPIPEDA, dissent, T, J.Doe, PII, Legitimate, User)$$

PII = Personally Identifiable Information.

Page 32 of Infoway's specification states: "Where POS systems connected to the EHRi record a patient/person's consent directives, including the withholding, withdrawal or revocation of consent, such POS systems must transmit these consent directives to the EHRi, in a consistent form, whenever they transmit the associated PHI to the EHRi." We model this statement as a (non-defeasible) consent for the collection of consent information.

$$R5 = (CanadaInfoway, consent, F, J.Doe, ConsentDirective, Legitimate, User)$$

Infoway's specification clearly specifies that the consent information includes personal identifiable information. Thus, it falls under PIPEDA. A person withholding consent for collecting his consent information would, however, be in conflict with the "non-defeasible" Infoway requirement.

Note that protecting consent information as confidential personal information is not merely an academic thought but has great practical importance. Disclosing the mere fact that some patients have withheld or

withdrawn their consent about collecting certain types of information may be a valuable and potentially harmful piece of information.

6.3. Limitations

The case study also revealed current limitations of the CREE method. Some of the requirements did not readily map to defeasible or non-defeasible consents or dissents or data collection/use/disclosure. Several requirements demanded explicit disclosure, use, or collection of data. For example, the Infoway specification demands that EHR access logs for patients with elevated risk of incurring invasions to their privacy (e.g., celebrities) are to be checked by auditors on a mandatory basis. This type of "must disclose/use"statement is more than a strict consent and cannot be adequately modeled with our current method.

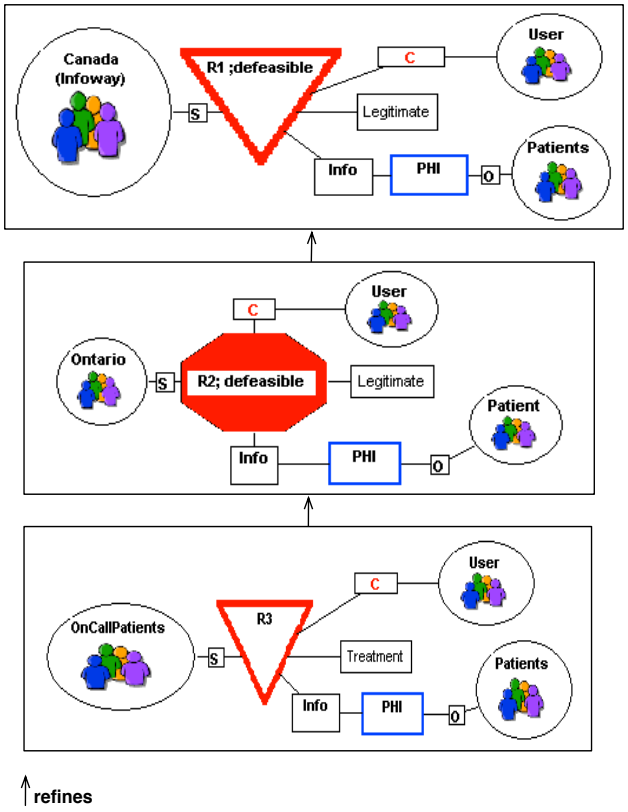

Figure 6. Requirement Refinements

7. Conclusion

Privacy and confidentiality requirements are usually stratified due to different levels of jurisdictions, legislation and policy. However, current methods do not adequately support this inherent relationship between requirements. An approach, based on stratified relationships, enables specification of the requirements as they naturally occur. Specification of concerns from stakeholders e.g., customers, need to be addressed as they are identified over the system life-cycle. *CREE* supports iterative requirement specification with the defeasibility concept, which allows default requirements that can be refined.

Our method has been tested with the Infoway example (Section 6) with partial success and the existing limitations are being considered in our ongoing research. Furthermore, some concepts in our confidentiality property meta model (Fig. 2) are not currently used in our analysis but are planned for future work. For example, the *strictness* attribute of the data and counterstakeholder properties can be used to improve the conciseness of the requirement specification. It denotes the implication of the requirement statement for the data/stakeholders not explicitly indicated. A dissent with *strictness* on data means the counterstakeholder may obtain data not explicitly specified, while a consent with strictness on data requires the unspecified data to be kept confidential. In addition, analysis with the *inferableFrom* property requires techniques to determine conditional uncertainty for disclosure.

References

[1] "An Overview of the Electronic Health Record Privacy and Security Conceptual Architecture," http://knowledge.infoway-inforoute.ca/EHRSRA/doc/EHR-Privacy-Security-Overview.pdf, accessed 10 June, 2007.

[2] "Canada Health Infoway," http://www.infoway-inforoute.ca/en/home/home.aspx, accessed 10 June, 2007.

[3] I. Alexander, "Initial Industrial Experience of Misuse Cases in Trade-Off Analysis," in *Proc. IEEE Joint Int'l Conf. Requirements Eng*, Essen, Germany, 9-13 September 2002, pp. 61–68.

[4] G. Antoniou and G. Aditya, "What is Default Reasoning Good For? Applications Revisited," in *Proc. 32nd Hawaii Int'l Conf. on System Sciences*, Maui, Hawaii, USA, 5-8 Jan 1999.

[5] S. Easterbrook and M. Chechik, "A Framework for Multi-Valued Reasoning over Inconsistent Viewpoints," in *Proc. 23rd Int'l Conf. Software Eng (ICSE '01)*. Toronto, Ontario, Canada: IEEE Computer Society, 2001, pp. 411–420.

[6] A. Fuxman, L. Liu, J. Mylopoulos, M. Pistore, M. Roveri, and P. Traverso, "Specifying and Analyzing Early Requirements in Tropos," *Journal Requirement Eng*, vol. 9, no. 2, pp. 132–150, May 2004.

[7] A. Fuxman, M. Pistore, J. Mylopoulos, and P. Traverso, "Model Checking Early Requirements Specifications in Tropos," in *Proc. 9th IEEE Int'l Requirements Eng Conf.*, Toronto, Canada, 27-31 August 2001, pp. 174–181.

[8] V. Gervasi and D. Zowghi, "Reasoning About Inconsistencies in Natural Language Requirements," *ACM Transactions Software Eng and Methodology*, vol. 14, no. 3, pp. 277–330, 2005.

[9] P. Giorgini, F. Massacci, J. Mylopoulos, and N. Zannone, "Requirements Engineering Meets Trust Management: Model, Methodology, and Reasoning," in *Proc. 2nd Int'l Conf. Trust Management (iTrust 2004)*, Oxford, UK, 2004, pp. 176–190.

[10] P. Giorgini, F. Massacci, J. Mylopoulos, and N. Zannone, "Modeling Security Requirements Through Ownership, Permission and Delegation," in *Proc. 13th IEEE Int'l Requirements Eng Conf.*, Paris, France, 2005, pp. 167–176.

[11] S. Gürses, J. H. Jahnke, C. Obry, A. Onabajo, T. Santen, and M. Price, "Eliciting Confidentiality Requirements in Practice," in *Proc. 15th Centers for Advanced Studies Conf. (CASCON 2005)*, Richmond Hill, ON, Canada, October 2005, pp. 207–222.

[12] HIPAA., "Health Insurance Portability and Accountability Act of 1996," http://aspe.hhs.gov/admnsimp/pl104191.htm, accessed 16 March, 2007.

[13] A. Hunter and B. Nuseibeh, "Managing Inconsistent Specifications: Reasoning, Analysis and Action," *ACM Transactions Software Eng and Methodology*, vol. 7, no. 4, pp. 335–367, October 1998.

[14] J. Jürjens, "UMLsec: Extending UML for Secure Systems Development," in *Proc. 5th Int'l Conf. Unified Modeling Language (UML 2002)*, Dresden, Germany, 2002, pp. 412–425.

[15] G. Kotonya and I. Sommerville, "Viewpoints for Requirements Definition," *Software Eng Journal*, vol. 7, no. 6, pp. 375–387, November 1992.

[16] L. Liu, E. Yu, and J. Mylopoulos, "Security and Privacy Requirements Analysis within a Social Setting," in *Proc. 11th IEEE Requirements Eng Conf.* Monterey Bay, CA, USA: IEEE Press, 8-12 September 2003, pp. 151–161.

[17] T. Lodderstedt, D. A. Basin, and J. Doser, "SecureUML: A UML-Based Modeling Language for Model-Driven Security," in *Proc. 5th Int'l Conf. Unified Modeling Language (UML 2002)*, Dresden, Germany, 2002, pp. 426–441.

[18] H. Mouratidis, P. Giorgini, and G. A. Manson, "Integrating Security and Systems Engineering: Towards the Modelling of Secure Information Systems," in *Proc. 15th Int'l Conf. on Advanced Information Systems Eng (CAiSE 2003)*, Klagenfurt, Austria, 16-20 June 2003, pp. 63–78.

[19] B. Nuseibeh and S. Easterbrook, "Requirements Engineering: A RoadMap," in *The Future of Software Eng, Companion vol Proc. 22nd Int'l Conf. Software Eng (ICSE 2000)*, A. C. W. Finkelstein, Ed. Limerick, Ireland: ACM Press, NY, USA, 4-11 June 2000, pp. 35–46.

[20] A. Onabajo and J. H. Jahnke, "Properties of Confidentiality Requirements," in *Proc. 19th Compuer-Based Medical Systems (CBMS 2006)*. Salt Lake City, Utah, USA: IEEE Computer Society, 22-23 June 2006, pp. 841–846.

[21] PIPEDA., "Office of the Privacy Commissioner of Canada: Personal Information Protection and Electronic Documents Act," http://www.privcom.gc.ca/legislation/02_06_01_e.asp, accessed 11 June, 2007.

[22] A. L. Rector, "Clinical Terminology: Why is it so Hard?" *Methods of Information in Medicine*, vol. 38, no. 4-5, pp. 239–252, Dec 1999.

[23] W. N. Robinson, S. D. Pawlowski, and V. Volkov, "Requirements Interaction Management," *ACM Computing Surveys*, vol. 35, no. 2, pp. 132–190, 2003.

[24] B. Schnarch, "Ownership, Control, Access and Possesion (OCAP) or Self-Determination Applied to Research: A Crtical Analysis of Contemporary First Nations Research and some Options for First Nations Communities," *Journal of Aboriginal Health*, vol. 1, no. 1, pp. 80–95, January 2004.

[25] R. W. Schwanke and G. E. Kaiser, "Living With Inconsistency in Large Systems," in *Proc. Int'l Workshop Software Version and Configuration Control*, Grassau, Germany, January 27-29 1988, pp. 98–118.

[26] I. Sommerville and P. Sawyer, "Viewpoints: Principles, Problems and a Practical Approach to Requirements Engineering," *Annals of Software Eng*, vol. 3, pp. 101–130, January 1997.

[27] A. van Lamsweerde, "Requirements Engineering in the year 00: A Research Perspective," in *Proc. 22nd Int'l Conf. Software Eng (ICSE 00)*. Limerick, Ireland: ACM Press, NY USA, 2000, pp. 5–19.

[28] A. van Lamsweerde, "Goal-Oriented Requirements Engineering: A Guided Tour," in *Proc. 5th IEEE Int'l Symposium Requirements Eng*, Toronto, Canada, 27-31 August 2001, pp. 249–262.

[29] A. van Lamsweerde, R. Darimont, and E. Letier, "Managing Conflicts in Goal-driven Requirements Engineering," *IEEE Transactions Software Eng*, vol. 24, no. 11, pp. 908–926, 1998.