09073 Summary Collection
Model-Based Design of Trustworthy Health Information Systems
— Dagstuhl Seminar —

Ruth Breu[1], John C. Mitchell[2], Janos Sztipanovits[3] and Alfred Winter[4]

[1] Universität Innsbruck, A
Ruth.Breu@uibk.ac.at
[2] Stanford University, USA
mitchell@cs.stanford.edu
[3] Vanderbilt University, USA
janos.sztipanovits@vanderbilt.edu
[4] Universität Leipzig, D
alfred.winter@imise.uni-leipzig.de

Abstract. The Dagstuhl Seminar "Model-Based Design of Trustworthy Information Systems" took place from February 11th to February 14th, 2009, at the International Conference and Research Center Schloss Dagstuhl. The goal of the seminar was to bring together experts from the domains of health care, software engineering and security in order to discuss the challenges of emerging health care scenarios. The seminar combined presentations with discussions in groups.

Keywords. Trustworthy Systems, Health Information Systems, Model-Based Design, Security Policies, Service Oriented Architecture

Summary

New technologies for Health Information Systems (HIS) offer a revolutionary new way for the interaction between medical patients and Healthcare providers. Although healthcare like other information-intensive industries has developed and deployed standards-based, secure information infrastructures it is still dependent upon paper records and fragmented, error-prone approaches to service delivery. Thus healthcare has been characterized as a "trillion dollar cottage industry". One of the main concerns is security and privacy that needs to be organically integrated into HIS architectures. Widely cited reports of the U.S. Institute of Medicine and National Research Council have documented weaknesses in information security related to healthcare, the costs and impact of medical errors (a substantial proportion of which involve a component of information mismanagement), lack of a systems approach to complex, team-oriented interdisciplinary care, and the unrealized potential of using the Internet to improve the quality and availability of healthcare services.

Dagstuhl Seminar Proceedings 09073
Model-Based Design of Trustworthy Health Information Systems
http://drops.dagstuhl.de/opus/volltexte/2009/1998

How can Health Information Systems help?

Complementing the recognition of the weaknesses are three major drivers that push the healthcare industry towards radical change: (1) the dramatic increase of genetic information and the opening opportunity to provide personalized healthcare, (2) the economic pressures to move healthcare from institutions toward homes, and (3) the rapidly increasing use of Internet and information appliances in society. This fundamental change will be enabled by advanced information technology, including ubiquitous communication and sensing, extensive use of web portals as a central point of access for communication and documentation of health care efficiency. Quality of patient specificity will be achieved via extensive use of clinical decision support systems combined with automated event monitors.

What are the key challenges?

HIS shall support patients and also doctors, nurses, paramedicals and other health care providers in diagnosing, treating and supporting patients. Health care is not only a health but also a life and death issue. In this existential situation patients have to trust on caregivers and both patients and caregivers depend on the trustworthiness of the information systems used. Not only the highly delicate relation between caregivers and patients but also the data related to this situation need particular protection from misuse. But unfortunately privacy and security requirements are frequently expressed in vague, contradictory and complex laws and regulations; it is a major concern that requires new approaches in systems design. Trustworthy HIS need to provide effective, high quality support for providing the best care for patients but without compromising their privacy, security and safety.

How to solve these challenges?

End-to-end architecture modeling integrated with privacy and security models offer new opportunities for system designers and end users. Model-based approaches to HIS are investigated extensively in Europe and in the US. While initial results show promise, many fundamental problems remained unsolved, such as modeling of privacy and security policies, and verification of their consistency, and compliance to requirements. HIS requires new architectures that are sufficiently flexible to support personalized health care without causing harm and can be adapted to changing policies.

Goals and Expected Results

The goal of this seminar was to help the computer science community understanding the unique challenges of this field and offer insight for HIS developers in the state of the art in model-based design technologies. The objective was to understand the challenges and promising approaches in HIS design as the intersection of five major areas: health information systems, model-based software and systems design, reliability, security and privacy science, enterprise information systems and legal policy. The seminar combined presentations with discussions in groups and in the plenary.

Design Aspects for Modular Surgical Assist Systems

Oliver Burgert (Universität Leipzig, DE)

The fields of Health Information Systems (HIS) and Computer Assisted Surgery (CAS) have developed rapidly in the last years. Nowadays, proprietary systems of different vendors are used for preoperative planning und intraoperative support. They partially integrate in the existing IT infrastructure (PACS, HIS), but the results of the performed actions are usually stored in proprietary formats. Furthermore, it is not possible to combine systems of different vendors to the needs of the hospital. a modular system architecture including a partial prototype implementation of a Therapy Imaging and Model Management System has been described. Afterwards, safety and security issues have been discussed.

Collaborative Security - From Risk and Compliance to the Security Architecture

Ruth Breu, Michael Hafner, Frank Innerhofer-Oberperfler
Research Group Quality Engineering, University of Innsbruck

Cooperation among stakeholders is a major issue of the new generation of applications in health care. This ranges from information exchange between attending physicians, providing patients with health-related information to agile care with automated, patient-specific alerts. The challenges to such applications are manifold. In particular, strict legal regulations and high requirements to security are in conflict with the continuous evolution of these systems. Since the borderline between IT management (IT governance), system design and system operation almost disappears, tight integration of these activities and a comprehensive approach to security is vital.

In this contribution we present a set of core principles supporting collaborative security. Collaborative security aims at the continuous management of security and compliance aspects as a cooperative task of many stakeholders. In our approach models and meta models provide the backbone for integrating different views and abstraction layers. The principles presented have been validated in the research project health@net in which a prototypic solution for a shared Electronic Health Record (EHR) has been developed as part of Austrian activities to establish a national distributed EHR.

Trustworthy health information systems - a user perspective

Thomas Bürkle, Chair Medical Informatics, University Erlangen-Nürnberg
Jochen Kaiser, IT-Security Manager University Hospital Erlangen

It starts simply: Give the clinical staff an electronic data processing system which makes some things go easier. Let's assume for example that you introduce a patient data management system (PDMS) in an intensive care unit (ICU). Let's assume, the new PDMS which reflects a department specific electronic medical record (EMR) is well accepted. With the consequence that clinical staff such as physicians, nurses, physiotherapists are going to use the system. Great !!!

But now your trouble starts: All those staff members tend to believe immediately everything the computer screen tells them. This could be data gathered from different parts of the EMR which is condensed and presented in a summary table or graphic diagram. When your PDMS is most likely user configured, how do you guarantee that you did not miss items in these summaries? Are you able to create appropriate test cases to make sure everything has been configured correctly? Think of fluid balance sheets, but also of caloric balance sheets, electrolyte balance and potential catastrophic consequences e.g. in premature babies!

If you display data in these summaries which is imported into the PDMS by interface, e.g. lab values imported via HL7: Can you assure a reliable stream of information across different computer systems? Are you absolutely sure that no values are missing, e.g. serum electrolytes included when sent from the central lab system, but not accounted for when sent from the local blood gas analyser? Again it is essential that correct information is maintained!

Let's continue questioning: Assuming that your PDMS manages data export as well. The PDMS will then collect accounting data such as respirator hours, diagnoses, procedures and score values for reimbursement purposes: Did you look inside the HL7 ORU export messages to confirm which amount of information is transmitted? ICD and ICPM codes will be transmitted fine, but additional free text diagnosis information collected in the PDMS is pretty likely to be lost in transfer.

And then, after a year of system use, the clinicians want answers from the PDMS. So you install the statistics module from the same vendor. Did you ever check its results? Would you be able to confirm the data quality of those statistics? There is a proverb small lies -big lies- statistics. But your staff will rely on these statistics.

Assume you would like to support clinical practice with medical knowledge, e.g. drug information in the PDMS as an interactive help. Fine. But are you sure: Is this the correct and appropriate medical information? If you connect medical knowledge with patient data to build knowledge based functions such as drug dosing advice: Do you ”produce“ a medical device? How do you prove that your knowledge based functions work correctly? What if not?

During installation you connected your PDMS with medical devices such as patient monitoring units and ventilators. Do you get a correct and timely readout of those parameters in the PDMS? Are you able to prove it? Every single value of 30 parameters delivered by an average ventilator? Try to catch a pulse depression in infants: It is shown on the monitor but often not in the PDMS. Or do you even find it difficult to identify from which patient the data comes? For example when medical devices are connected to the LAN and moved between patients? And there is more trouble ahead: You connected the PDMS to a Medical Device Network, e.g. to a monitoring network. Does the monitoring network work reliably? What happens if not? Do you really want monitoring data be transmitted via WLAN (offered by some vendors recently)? If so, how do you deal with data protection? How about the patch status of all the different computers which make the connection? Will you be able to patch the medical device network gateways, so that they are not affected by virus or other threats? Or do you run the risk that an internet worm may spread via unprotected medical device networks throughout your hospital?

If so, do you have your own hospital medical technology department? Or did you just outsource maintenance to the vendors? With how many vendors will you have to talk? Who is responsible? Does your IT security manager know what to do? What is your experience when talking with the vendors? How quickly did they react?

Remember: There is no such thing as a bullet proof piece of software. Any non trivial software program cannot be authenticated. Software testing results are a function of effort being used for the test.

Think about better risk management. Think about safety measures. Which measure does the least damage and provides highest efficiency? Define responsible contact persons and appropriate checklists for medical devices. Catalogue your software products in terms of risk management.

An Analysis-based Approach to Improving Medical Processes

Lori A. Clarke (Univ. of Massachusetts - Amherst, US)

As has been widely reported in the news lately, medical errors are a major cause of death and suffering. In the University of Massachusetts Medical Safety Project, we are investigating if software engineering technologies can be used to help reduce medical errors. Specifically, we are modeling medical processes, using the Little-JIL process definition language, and then analyzing these processes using a range of analysis techniques. Working with the UMASS School of Nursing and the Baystate Medical Center, we are undertaking in-depth case studies on error-prone and life-critical medical processes. In many ways, these processes are similar to complex, distributed systems in that they have many interacting, concurrent threads and exceptional conditions frequently arise that must be carefully handled.

Analysis is a cornerstone of our approach. We believe that if the system is interesting enough to warrant being modeled, then the model is probably complex enough to warrant careful scrutiny by rigorous and automated analysis techniques. Without such scrutiny one should have serious concerns about the validity of the model and any decisions made based on that model. Thus, to support rigorous analysis, the semantics of the modeling language must be formally and precisely defined. Moreover, the modeling language should provide rich semantic features, such as extensive support for concurrency, exception handling, and abstraction, in order to represent the medical processes accurately. Our experience suggests that these sorts of detailed and complex process models should be developed incrementally so that high-level, more-abstract views of the process can be validated before more-detailed models are developed. The scope and granularity of the model should be determined by the questions the model is intended to address. There is no doubt that detailed models require more effort to develop and maintain, but provide more definitive, in-depth feedback; in other words, there is no free lunch!

Since we are creating models on which to base decisions and further reasoning, it is incumbent upon us to work to make these models accurate enough to justify this trust. As the models are repeatedly validated using a range of analysis techniques, we increase our confidence in their accuracy. The analyses that we have been applying include finite-state verification to determine if all traces through a model adhere to properties that indicate the legal sequences of events, fault-tree analysis and failure mode and effects analysis to reveal vulnerabilities if steps in the process are not executed appropriately, and discrete-event simulation to determine the aggregate behavior after a large number of traces have been executed. These are by no means all the kinds of analyses that should be considered, but each of these provides distinctive kinds of feedback.

Our work has been quite satisfying in that the detailed process models and the analysis that we have applied have indeed discovered errors in actual medical processes. Indeed every step in this approach, from process modeling, to property specification, to process model verification has led to the discovery of errors of one kind or another in the actual processes. Moreover, his project is succeeding in providing benefits to both healthcare and software engineering. The medical professionals involved have reported that this project has changed the way they view, describe, teach, evaluate, and improve their processes. Moreover, several serious problems have been uncovered and the medical processes have subsequently been improved. There have also been benefits to software engineering in that it has been necessary to enhance the technologies we have used in ways that should also improve their effectiveness when applied to software systems. Moreover, we now have a new perspective on software development, particularly for human-intensive systems.

Foundations of Privacy: Contextual Integrity, The Logic of Privacy and Beyond

Anupam Datta (Carnegie Mellon University - Pittsburgh, US)

Organizations, such as businesses, non-profits, government agencies, hospitals, banks, and universities, collect and use personal information from a range of sources, shared with specific expectations about how it will be managed and used. Accordingly, they must find ways to comply with expectations, which may be complex and varied, as well as with relevant privacy laws and regulations, while they minimize operational risk and carry out core functions of the organization efficiently and effectively.

This contribution reports report on a principled approach for expressing and enforcing privacy policies in complex organizational processes. The starting point of our work is "contextual integrity", a conceptual framework for understanding privacy expectations and their implications developed in the literature on law, public policy, and political philosophy. We formalize some aspects of contextual integrity in a logical framework for expressing norms of transmission of personal information. The technical approach is based on temporal logic with semantics defined over concurrent game structures. In comparison with access control and privacy policy frameworks such as RBAC, EPAL, and P3P, these norms focus on who personal information is about, how it is transmitted, and past and future actions by both the subject and the users of the information. Our logic is expressive enough to capture naturally many notions of privacy found in legislation, including those found in HIPAA, COPPA, and GLBA. In addition to privacy, we formalize a notion of "utility" that captures the goals of the organization, e.g. since a hospital's goal is to provide health care, certain flows of personal information are necessary. We also develop automated support for policy compliance, audit, and policy analysis.

While contextual integrity and its formalization focuses on personal information about individuals, privacy policies also refer to aggregate or anonymized information about groups of individuals. I will describe some of our ongoing work on integrating database privacy concepts into formal policy models and languages. Specifically, I will report on our experiences with formalizing and lifting differential privacy (a promising recent approach to database privacy) to reactive organizational processes.

Specifying and Analyzing Workflows for Automated Identification and Data Capture

Carl Gunter (Univ. of Illinois - Urbana, US)

Humans use computers to carry out tasks that neither is able to do easily alone: humans provide eyes, hands, and judgment while computers provide computation, networking, and storage. This symbiosis is especially evident in workflows where humans identify objects using bar codes or RFID tags and capture data about them for the computer. This Automated Identification and Data Capture (AIDC) is increasingly important in areas such as inventory systems and health care. Humans involved in AIDC follow simple rules and rely on the computer to catch mistakes; in complex situations this reliance can lead to mismatches between human workflows and system programming. In this talk we explore the design, implementation and formal modeling of AIDC for vital signs measurements in hospitals.

To this end we describe the design of a wireless mobile medical mediator device that mediates between identifications, measurements, and updates of Electronic Health Records (EHRs). We implement this as a system Med2 that uses PDAs equipped with Bluetooth, WiFi, and RFID wireless capabilities. Using Communicating Sequential Processes (CSP) we jointly specify workflow and computer system operations and provide a formal analysis of the protections the system provides for user errors.

Mobile Blood Donation Registration Service: Security and Privacy Issues

Patrick Hung (University of Ontario, CA)

This seminar describes a case study of adopting Service Oriented Architecture (SOA) in the Hong Kong Red Cross to support blood transfusion services. Every time a blood donor attends a donation center, he/she must proceed through the registration process, requiring them to fill in the registration form, regardless of whether they are a first-time-donor. Based on some research studies, the registration process is a cumbersome and time consuming process which increases the drop-out rate of blood donors. The purposes of this research project are to optimize the process and minimize the drop-out rate of blood donors by using semantic rules. Using the Web 2.0 and semantic Web technologies, the system includes a novel service which supports the privacy access control and security, in protecting the donors's personal information provided throughout the blood donation process. The system is built on the mobile model of SOA and XML related security technologies. This seminar discusses the security and privacy issues of such a mobile service. The system is currently being tested and studied in the Hong Kong Red Cross blood donation center. In addition, an empirical study of technology adoption is also conducted at the site to test the usability and feasibility of such a system from the blood donors's perspectives.

Realize a trustworthy health information system

Christoph Isele (Siemens Medical Solutions - Berlin, DE)

For a dedicated application there could be a congruent context, that the manufacturer build the system according to the requirements of the heath professional and the department implements it according to the directive of the manufacturer. Because features of data privacy are part of the original requirements and the users and roles are manageable data privacy of the dedicated solution should be achieved.

For the integrating health systems like hospital information systems there are breaks in the architecture. These systems collect and interpret information from different sources with different but non published models. By explicit customizing or interfacing of different systems they are realized on the site of the hospital and every instance is different from the 'guaranteed reference instance' of the manufacturer. Guidelines of the vendor and basic models of the communication standard can help to validate and test the reliability. But in real life it is too much effort for the most hospitals.

A sufficient policy usually is too complex for the administration, too much change in roles and rights in roles occur. So in practice the predictable profiles provide a midrange access, that could be extended in an emergency mode, but that is logged and can be tracked in case of miss use.

The problems of reliability and data privacy rise when the established routine is supplemented by new technology and enhanced opportunities.

Dedicated systems integrated in an IP network with new interfaces can be attacked.

Interpret data for decision support must have a much better support for interoperability than display information / documents. A suggestion has to be explained 'possible across the interface ?'

In a network of health information systems there should be some institution that can accept responsibility for the reliability of the communication paths and the data privacy. Otherwise there can be a clear stipulation about this between the participating institutions.

Existing health information systems are mostly realized by the institution (in projects including the vendors). The 'architect' in the solution needs support to build, operate and check the reliability and the data privacy of his information system.

Data privacy means always patient data and employee data.

Electronic Identification and Process Management in Septic Patients

Jason Martin (Vanderbilt University School of Medicine, US)

Sepsis is a medical condition characterized by a systemic inflammatory response to an infection. The disease is common and occurs without gender, racial, or geographic boundaries. In the United States in 2001, approximately 750,000 cases of sepsis were reported, resulting in tens of thousands of deaths. Inpatient mortality rates remain unacceptably high despite advances in sepsis therapy and critical care. The Surviving Sepsis Campaign, a report of professional society guidelines for sepsis management, calls for various interventions. Effective implementation of these interventions is challenging in the modern intensive care unit environment; many of the actions are conditional and time-sensitive. Recent evidence suggests that these interventions are more likely to be applied correctly, and in a timely manner, when administered as formal protocols. Protocolization of sepsis management improves compliance with accepted standards of care, may shorten hospital stays (inpatient and intensive care), and may improve mortality.

In an effort to apply these accepted, evidence-based management practices to septic patients at Vanderbilt Medical Center in Nashville, Tennessee, we propose a sepsis action plan that leverages Vanderbilt's extensive technology infrastructure. Our specific aims are to 1) develop electronic tools to facilitate identification and treatment of septic patients, 2) deploy these electronic tools in Vanderbilt's Medical Intensive Care Unit and Surgical Intensive Care Unit, and 3) study the impact of these tools on compliance with standards of care and pertinent clinical outcomes. Our hypothesis is that automated identification and electronically-guided process management (protocol implementation) will facilitate greater compliance with benchmarks of quality care and improve various pertinent clinical outcomes.

We are developing two novel applications. The first application aids in identification of septic patients. The disease onset can be subtle, and delays in therapy may result in adverse patient outcomes. Our application identifies patients with certain sepsis-associated vital sign and/or laboratory abnormalities, and it prompts physicians to evaluate the patient for sepsis. Once identified, a second application aids in process management. It provides real-time, dynamic, customizable, and evidence-based decision support at the point of care.

Model-based Security Engineering and Applications to Health Information Systems

Jan Jürjens (Open University (UK) and Microsoft Research (Cambridge))

Health-care information systems are particularly security-critical. In order to make these applications secure, the security analysis has to be an integral part of the system design and IT management process for such systems.

This talk presents the experiences and results from the security analysis of the system architecture of the German Health Card, by making use of an approach to Model-based Security Engineering that is based on the UML extension UMLsec. The focus lies on the security mechanisms and security policies of the smart-card based architecture which were analyzed using the UMLsec method and tools.

Main results of the talk include a report on the employment of the UMLsec method in an industrial health information systems context as well as indications of its benefits and limitations.

In particular, two potential security weaknesses were detected and counter measures discussed.

The results indicate that it can be feasible to apply a model-based security analysis using UMLsec to an industrial health information system like the German Health Card architecture, and that doing so can have concrete benefits (such as discovering potential weaknesses, and an increased confidence that no further vulnerabilities of the kind that were considered are present).

Automatic Detection of Policies from Electronic Medical Record Access Logs

Bradley Malin (Vanderbilt University, US)

Healthcare organizations (HCOs) are increasingly adopting clinical information systems for managing patients' electronic medical records (EMRs). To support these activities, various model-based software platforms, such as Vanderbilt's Model-Integrated Clinical Information System (MICIS) have been proposed to assist in the rapid development and evaluation of formal systems based on service oriented architectures. At the same time, these systems have integrated robust privacy and security policy specification and validation languages, such as Stanford's logic based on contextual integrity. However, a significant remaining question is "what policies should be specified for data protection?" This question is difficult to address because healthcare environments are inherently dynamic, such that system have fuzzy underspecified rules, and both users and patients are constantly moving in and out of the system. This paper describes a software tool to automatically assist healthcare organizations (HCOs) in discovering and defining policies for access to their clinical information systems. The Healthcare Organizational Network Extraction Toolkit (HORNET) is an organization-nonspecific Java-based program that mines an HCOs' EMR access logs to determine the underlying workflows and relationships in the system. HORNET performs this task by extracting a social network of users from the access logs and then generating association rules to indicate probabilities and strengths of associations. The system is heavily optimized to handle large networks, such as interactions between thousands of care providers. HORNET leverages novel statistical mechanisms, based on reciprocity in networks, to discover relationships between users and rules across a hospital's departments.

We evaluated HORNET with five months of access logs from the Vanderbilt University Medical Center. The sample started in January 2006 and included 9940 unique care providers and 350,889 unique patients, resulting in over 7.5 million access events. Our findings show that the network, at an individual level is highly volatile over time 82% of relationships no longer exist after 1 week and 90% no longer exist after 5 months. At a global level, though, the network remains stable, as we see a high degree of stability in our rules. We evaluated the rules for their existence and variability over time, in order to discover meaningful rules that can form the basis of defining what is normal for more advanced auditing. This duality quantifies the difficulty with which security administrators have in defining strict access policies and shows that a data mining approach can likely generate stable rules. We have successfully generated association rules which show logical and expected relationships as having high confidence and support. Our research demonstrates the feasibility of mining HCO access logs to discover underlying relationships and workflows in a dynamic setting.

Security and Privacy in EHR and PHR systems

Lorenzo D. Martino (Purdue University, US)

This presentation analyses the impact of regulations, business and organizational factors on security and privacy in Electronic Health Record systems and Personal Health record systems. The presentation addresses both EHR systems and PHR systems. As to the former, we show the difficulty for healthcare organization to identify in an unumbigous way the healthcare information to be protected starting from the HIPAA regulation, as well as other impeding factors, such as conflicting state and federal regulations. From the technical point of view this requires to define a framework in which regulatory, business and organizational factors are dealt with at several level of abstractions and by different technical means. As to the PHR, end-user organizations advocate the right of the patients to control the security and privacy policies governing the access to their healthcare data. We propose a set of simple classifications of PHR systems which allow to understand the factors to be taken into account when evaluting some proposed technical reference architecture for PHR systems and the real possibilty of end-user control on security and privacy policies.

Run-time Provision of Organizational Security Patterns for E-health

Fabio Massacci (Università di Trento, IT)

Patterns have been proposed in mainstream Software as best practices that capture knowledge of domain experts (security and dependability in our case) intended to be used at design time and they informally discuss how system components should interact.

In many e-health scenarios (such as smart-homes for elderly patients) we need security patterns of completely different kinds:- we need to model humans in the process as they play a key role in the provision of the security and dependability of the systems - we need to provision at run-time these patterns based on the organizational context.

An example of a run-time provision is the patient is sick so we must change the traditional method of authentication for the entering the house and the emergency rescue team must get in.

In this talk I'll present some of the ideas that srtaing form our high-level requirements engineering language for socio-technical systems based on the concepts of agents, trust, dependencies arrives to runtime executable components pluggable to the application at runtine and replaceable dynamically.

If time allows I'll present some videos of how the idea is implemented in a real smart-home for impared people.

Medical information systems and privacy policy

John C. Mitchell (Stanford University, US)

Outline 1) Medical privacy problem: one part of a larger set of interesting challenges 2) Contextual integrity: a philosophical account of privacy, made precise 3) Workflow: an approach to privacy in context 4) HIPAA formalization: a sample effort to write down complex policy and use it in a prototype patient portal system

Precise Definition of Health Care Processes

Leon J. Osterweil (Univ. of Massachusetts - Amherst, US)

I am particularly interested in the use of rigorous languages to define rigorously and precisely the processes by which health care is delivered. Our group's research has indicated that these processes are often very complex. It has not been unusual to find that even the medical professionals who participate in performing these processes have been surprised at their complexity. The complexity of such processes seems to lead to errors in their performance, and we believe that such errors are at least partially responsible for considerable excess cost, unnecessary pain and suffering, and in some cases even death. Thus, a primary purpose in defining health care processes precisely is to use them as a vehicle for helping their performers to understand their roles better and to seek way to improve the processes (e.g. by removing process defects). Thus, another major goal of this research is assure that health care processes are defined sufficient well that the process definitions can be an adequate basis for effective analysis. Our group's work in applying a range of rigorous analysis approaches will be addressed in a companion statement (and presentation). But that work on analysis has emphasized the importance of using a strong, precise, well-defined language to define these processes.

The process definition language that has been used as the basis for this work is Little-JIL, a language that provides such semantic features as exception management, modularization, abstraction, concurrency control, and late binding of resources to tasks. Experience in using Little-JIL to define real health care processes, elicited from research colleagues at the Baystate Medical Center, Springfield, MA, USA, has indicated that these semantic features are essential if health care processes are to be defined sufficiently clearly, precisely, and completely to support effective detection and removal of process defects.

Thus, an overriding goal of our group's research is to use process definition and analysis to support continuous process improvement in the health care process domain. In my talk I will emphasize the contributions needed from a process definition language. Further, as our language, Little-JIL, is defined with execution semantics, the actual execution of processes is possible, and we propose to experiment with using Little-JIL process definitions to help guide the performance of health care processes.

In summary, our group's vision is to support health care with the execution of health care processes that have been rigorously defined, and iteratively improved through the removal of defects identified by the application of diverse forms of analysis.

Trust in eHealth Processes

Reinhard Posch (TU Graz, AT)

With eHealth we face basically the same assumptions as we have in other large Applications that include the general public. The most critical part is that trust and security is not the business of the actors in charge and therefore robust and mechanistic approaches are the only avenue for success. Not only the assumptions but also the building blocks like document security and robustness, identification for data protection and portal security and delivery of results, as well as archiving have the same character as with general administration. When it comes to facets like maternity payments or driving licenses the fields even overlap. Still we are far from benefiting from these synergies.

Operational Models for Security and Dependability in Electronic Health Systems

Roland Rieke , Fraunhofer SIT - Darmstadt

Security and privacy are critical aspects for the acceptance of emerging new complex technologies in the public sector, particularly the protection of personal health data is of utmost importance. In this talk four scenarios are presented where operational models for security and dependability with relevance for application in electronic health systems have been developed and analysed. These scenarios comprise,

1. a workflow and organisation based access control model for the management of medical records in hospitals,
2. an architecture with protocols for provisioning and enforcement of security policies,
3. model based test case generation for the German electronic Health Card (eHC), Health Professional Card (HPC) and Security Module Card (SMC) and their interplay,
4. a security analysis of the German Health Card infrastructure and services in particular the management services for the insurance master data.

Key priority in (1) is the inherent ambivalence between Privacy and Need to Know requirements for the processing of medical records. A compact visualisation of aspects of such a system's behaviour and examples of properties that can be verified are given.

Scenario (2) is concerned with policy provisioning and enforcement. In atypical policy controlled system, a set of policy rules, posing restrictions on the system's behaviour, is used to enforce the required security objectives. Integration of policy validation into a policy based architecture was the main goal here.

Main topic in (3) is the compliance of an implementation with the specification. The implementation on the smartcards is measured for compliance to the specification via a suite of test case sequences that are generated from the model.

In (4) the specification of the security requirements and the specification of the security mechanisms was analysed. The use case oriented specification was transfered to an asynchronous model (using APA). In order to prove that the model correctly represents the specification in such complex systems it is very useful to derive compact representations of component behaviour from global behaviour by computation of adequate property preserving abstractions.

In the finite state model of scenario (4) the modelling of timers, counters and logging mechanisms was critical for the scalability of the model and the properties that can be verified. Modelling problems approached during the course of action and open problems will be presented.

The operational finite state models of the scenarios above are based on Asynchronous Product Automata (APA), a flexible operational specification concept for cooperating systems. An APA consists of a family of so called elementary automata communicating by common components of their state (shared memory). A short coverage of the applied modelling and verification concepts and tools and of (technical) challenges is also provided.

References

[1] Peters, J., Rieke, R., Rochaeli, T., Steinemann, B., Wolf, R.: A Holistic Approach to Security Policies - Policy Distribution with XACML over COPS. In: *Proc. of the Second International Workshop on Views On Designing Complex Architectures* (VODCA 2006). Volume 168., Elsevier (2007) 143-157

[2] Ochsenschläger, P., Rieke, R.: Abstraction Based Verification of a Parameterised Policy Controlled System. In: *International Conference "Mathematical Methods, Models and Architectures for Computer Networks Security" (MMM-ACNS-7).* Volume 1 of CCIS., Springer (2007) c Springer.

[3] Ochsenschläger, P., Rieke, R., Velikova, Z.: Die elektronische Krankenakte - Eine Sicherheitsstrategie. In: *DACH Security 2008 - Bestandsaufnahme, Konzepte, Anwendungen, Perspektiven.* (2008) 90-100

[4] Rieke, R.: Abstraction-based analysis of known and unknown vulnerabilities of critical information infrastructures. *International Journal of System of Systems Engineering* (IJSSE) 1 (2008) 59-77 Copyright: c 2008, InderScience.

[5] Apel, C., Repp, J., Rieke, R., Steingruber, J.: Modellbasiertes Testen der deutschen Gesundheitskarten. In: *DACH Security 2007 - Bestandsaufnahme, Konzepte, An wendungen, Perspektiven.* (2007) 338-346

Collaborative Planning with Confidentiality

André Scedrov (University of Pennsylvania, US)

Scedrov's work with Kanovich and Rowe introduces a formal model of collaboration, which addresses confidentiality concerns. The authors draw on the notion of a plan which originates in the artifical intelligence literature. The current work uses data confidentiality policies to assess confidentiality in transition systems where actions have pre- and post- conditions of the same size. Under two natural notions of policy compliance, the work shows that it is decidable in polynomial space (PSPACE), and in principle fully automatable, to schedule a plan leading from a given initial state to a desired goal state while simultaneously deciding compliance with respect to the agents' confidentiality policies.

Data Protection Requirements for setting up EHR Systems and the "Austria ELGA Policy"

Klaus Schindelwig (Tiroler Landeskrankenanstalten, AT)

EHR systems have the potential to achieve greater quality and security in medical information than the traditional forms of medical documentation. However, from a data protection point of view the fact has to be stressed that EHR systems additionally have the potential not only to process more personal data (e.g. in new contexts, or through aggregation) but also to make a patient's data more readily available to a wider circle of recipients than before. (ARTICLE 29 Data Protection Working Party, WP 131).

There are two basic models for access to patient data at trans-regional: On one side a uniform system of a single manufacturer for all involved communication parties and on the other side no uniform system, but mandatory uniform standards and uniform interfaces.

Independently of the models the question arises: Is there a common understanding, how data protection in such models should be realized?

The presentation will first discuss several activities of other countries in the EHR domain and analyze the legal requirements for introducing a EHR. Subsequently it will outline the current situation of the ELGA (= Austrian EHR) Policy.

k-Anonymity Considered Harmful

Vitaly Shmatikov (University of Texas - Austin, US)

K-anonymity and related methods based on generalization and suppression of the so called "quasi-identifiers" such as ZIP codes and birthdates are a popular technique for protecting privacy of databases containing records of specific individuals.

k-anonymity suffers from the number of limitations: it does not hide whether a given individual is in the database, does not prevent disclosure of sensitive information about individuals, does not protect against attacks based on auxiliary knowledge, does not compose (i.e., multiple releases of the same database can lead to privacy breaches), and cannot be applied to high-dimensional data.

Nevertheless, if one assumes the extremely weakly adversarial model used in the k-anonymity literature, one may hope that k-anonymization offers some data-mining benefits over trivial sanitization, which simply separates quasi-identifiers from sensitive attributes.

Unfortunately, this is not the case. Using the same datasets from the UCI machine learning repository as were used in previous research on generalization and suppression, we show that privacy gains require almost complete destruction of the data-mining utility. In most cases, trivial sanitization provides equivalent utility and better privacy than k-anonymity, l-diversity, and similar methods based on generalization and suppression.

An Approach for Dynamic Risk Monitoring based on Key Indicators

Atle Refsdal and Ketil Stølen (SINTEF - Oslo, Norway)

Obtaining risk levels requires us to find likelihoods and consequences for the risks in question. This is often very hard. Furthermore, the values obtained may soon be outdated as the system under analysis or its environment change. We propose an approach for dynamic risk monitoring based on measurable key indicators. As the assumed correspondence between indicators and risk levels is typically based on subjective judgments, the approach also includes dynamic monitoring of the degree of inconsistency in the risk picture, with the purpose of revealing weaknesses of the analysis.

The Challenge In order for a security risk analysis to serve its purpose, we need to trust that the risk levels obtained for the identified risks are (at least roughly) correct. This requires finding good answers to the following questions: 1) How likely is the unwanted incident in question to occur? 2) What is the consequence if this incident occurs? Unfortunately, in most cases the answers obtained from a risk analysis will provide a snapshot reflecting a single point in time. Hence, the risk values may soon be outdated as the system under analysis or its environment change.

Moreover, finding correct likelihood and consequence values is often very hard. This is typically the case if we are analyzing a new system where historical data do not exist, or if the incident in question cannot easily be observed directly.

We therefore need to seek ways of obtaining good estimates of likelihood and consequence values. One way of doing this is to base the assessments on measurable indicators that are seen as relevant for the unwanted incident in question. For example, if we want to estimate the likelihood that an intruder accesses sensitive data by logging on to a computer with the username and password of a legitimate user, it may be useful to know how many passwords have not been changed during the last three months and how many of the users do not comply with the company's password strength policy. If we are able to define likelihood, consequence and risk levels as functions from sets of indicators, we also ensure that risk levels can be updated automatically as soon as the indicators are updated, rather than representing a snapshot at a given point in time. However, defining functions from sets of indicators to risk levels is clearly very challenging, and will typically have to be done based on subjective expert judgments. It is therefore important to find ways of uncovering weaknesses of the analysis.

The Proposed Approach In (1) we propose an approach for providing a dynamic risk picture and for assessing to what degree we can be confident that the risk levels obtained are correct. A basic assumption of the approach is that an infrastructure is available for defining and monitoring the measurable indicators required. Providing such an infrastructure is an important goal for the project MASTER (see http://www.master-fp7.eu/), which addresses the challenge of managing assurance, security and trust for service-oriented systems. However, the approach we present is general in the sense that we just assume the availability of a palette of monitored indicators; the infrastructure required to obtain them is not considered.

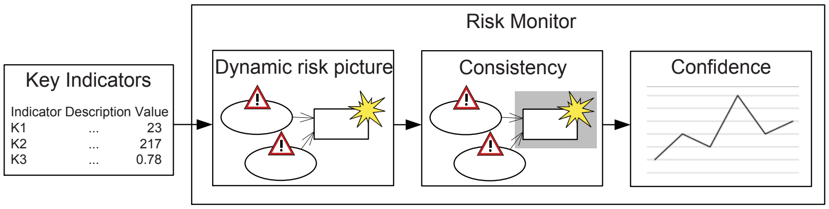

Indicator	Description	Value
K1	...	23
K2	...	217
K3	...	0.78

Fig. 1. Risk Monitor Modules.

Figure 1 outlines our vision for a dynamic risk monitor defined on the top of some monitoring infrastructure. Our envisaged dynamic Risk Monitor consists of three modules as indicated in Figure 1. The “Dynamic Risk Picture“ module allows the user to monitor the likelihood, consequence, and risk values, thereby providing a more high-level view than the ”Key Indicators” infrastructure. Values may be presented in graphical diagrams that show how threat scenarios lead up to unwanted incidents; likelihood values may be assigned to threat scenarios as well as unwanted incidents. The values are obtained from functions for calculating likelihood and consequence values from sets of key indicators, as well as for calculating risk vlaues from likelihood and consequence values. These functions are defined during the risk analysis, as the relevant risks will depend on the system in question.

The “Risk Consistency” module checks whether the risk picture is consistent at a given point in time. This can be done by comparing likelihoods for threat scenarios assumed to lead up to an unwanted incident with the likelihood of the actual incident.

Finally, the “Confidence” module offers a quantitative measure of confidence in the current risk picture, thereby providing an aggregated view from which the correctness of the analysis can be assessed. This measure is based on the degree of inconsistency detected in the risk picture.

Acknowledgements: The research leading to these results has been funded from the European Community’s Seventh Framework Programme (FP7/2007-2013) under grant agreement FP7-216917.

References

[1] Atle Refsdal and Ketil Stølen. Employing key indicators to provide a dynamic risk picture with a notion of confidence. In: *Proceedings of the 3rd IFIP International Conference on Trust Management,* 2009. To appear

Contributions of systematic information management to trustworthiness of information systems in healthcare

Alfred Winter (Universität Leipzig, DE)

The presentation will differentiate the tasks of software labs healthcare insitutuions and vendors, which are needed to achieve trustworthiness. It will shortly introduce 3LGM as a tool to support model based information management at a health care institution.

The Cancer Institute of New Jersey’s Tissue Repository: A Privacy and Security Case Study

Rebecca Wright (Rutgers Univ. - Piscataway, US)

The Cancer Institute of New Jersey is developing a biorepository of human tissue to be used as a resource for researchers. We are working with them to consider the repository as a case study of privacy, security, and trust issues.