Reflections on Security of Database and Datatransfer Systems in Health Care

Joachim Biskup, Gerrit Bleumer

Institut für Informatik, Universität Hildesheim, Samelsonplatz 1, D-31141 Hildesheim, Germany, Fax: ++49 5121 883-732, e-mail: {biskup, bleumer}@informatik.uni-hildesheim.de

Abstract

Health care is a complex and transnational task of societies. The various institutions and individuals involved are not and should not be governed by one central authority, but have legitimate and sometimes diverging interests. Balancing these interests has been subject to organizational and legal procedures traditionally. Today, institutions like hospitals and practices are no longer employers of isolated IT-systems, but they increasingly collaborate by means of a transnational network of various highly efficient database and datatransfer components, etc. Balancing all legitimate interests demands too much of organizational measures and, thus, requires direct support by the IT-systems themselves. The most important implication for the specification of such systems is that it must explicitly allow different parties to distrust different components of the system. We start from a simple specification of a "secure" system and unfold it with respect to typical requirements of health care.

Keyword Codes: F.3.1; H.4.3; J.3
Keywords: Specifying, Verifying and Reasoning about Programs; Communications Applications; Life and Medical Sciences

1. A PERSPECTIVE OF HEALTH CARE INFORMATICS

Health and *life* are a basic value of a society. E.g., the potential of personal development strongly depends upon a person's state of health and sanity, upon one's own and other people's knowledge about it. The elementary relationship of health care is the one between patient and physician. Naturally, it is characterized by the patient's trust into the physician's personal skills and the physician's obligation to use his skills, his knowledge and the technology available for him to cure the patient, at reasonable cost. In these respects the patient-physician relationship differs from most other relationships, e.g., in business or military, which are potentially less individual. Essential prerequisites for effective patient care are *professional secrecy* of physicians (oath of Hippocrates) and *informational self-determination* of patients, physicians and medical staff. The informational self-determination of physicians and medical staff, however, must be carefully balanced against the society's interest in preventing misuse. Generally, informational self-determination is another basic value of a society. A lot of institutions like hospitals, health insurances, regional, national, and international authorities are supporting the patient-physician relationships. Accordingly, financial clearing between all patients, physicians, and institutions involved is a tremendous task which gives rise to consider a third basic value of a society: *property*.

Informatics has made more and more information technology available to health care. The IT-systems' capability of automatically and efficiently processing mass data and the upgradability of IT-systems have legitimately attracted people to employ this technology for many tasks of health care. Today most information technology is still employed separately on a local basis. In near future, however, formerly separated components, even if located in different regions and countries, will be integrated into global health care IT-systems. This will be achieved by heterogeneous and federated database and datatransfer systems. At least in the long run health care will reach an essentially new stage of development where the roles of individual communication and direct interaction among patients and medical staff will be newly defined.

Traditionally, societies have kept their basic values by organizational regulations and legal proceedings. However, integrated IT-systems will demand too much of organizational regulations. A fundamental challenge to health care informatics is to replace those organizational measures by technical means such that the IT-systems support to keep the basic values directly rather than by external organizational measures.

2 . SECURITY

During the past decade security has increasingly been recognized as an important property of information technology. Despite of a lot of scientific, industrial, and governmental activities (e.g., [1,2,3,4], [5], and [6]), there is still no commonly accepted understanding of the notion of security. We propose a framework for secure system specification and verification which reflects the different interests of the parties involved.

We argue for expanding the notion for hierarchical levels of abstractions, which of course, as is usual with reasonings of this kind, are intrinsically interrelated. The intuitive top level deals with a society, a community, and an IT-system. The *society* is entitled to keep *basic values* which might conflict with each other. Important examples are presented in chapter 1, others are participation, co-determination, protection of environment, and secrets. Those *parties* of the society who employ or use the IT-system to achieve their specific *(sub)purposes* form the *community*. It might be one institution, or a federation of institutions, or a federation of institutions and individuals. We think of an *IT-system* as a distributed set of components (like database and datatransfer systems, networks, personal devices, etc.) controlled by potentially different parties of the community and possibly affecting other parties. Parties controlling are, e.g., database employers, companies maintaining them, physicians using them, patients holding their chipcards, etc. Parties affected are, e.g., patients whose data is processed whether they have access to the IT-system or not.

A typical purpose is to require a (directed) information flow from some sender to some receiver. At least in general, also communal purposes serve socially respected basic values, and these values may in their turn conflict with other basic values. On this level of abstraction, therefore, our notion of security is based on the fundamental *value assumption* that, by some decision procedure, the community pursues legitimate and consistent purposes. Then we tentatively define:

> A system is *secure* iff it satisfies the intended purposes without violating other relevant basic values.

At the operational bottom level all concepts should be fully formalized. Firstly, we assume that the IT-system is precisely given as an executable formalism F, say as a set of CONCURRENT

PASCAL-programs (syntactic part) together with a (virtual) CONCURRENT PASCAL-machine (semantic part). Secondly, all the communal purposes, i.e., the collection of all subpurposes of all parties, should be refined to a family of technical services *Req*. As far as possible, and in practical situations we can achieve this goal only approximately, these technical services should be specified completely and closely. Supposing that a statement of the form "an executable formalism satisfies a specified service" is well-defined and, at least in principle, provable, we can now precisely relate the executable formalism F to its specification. Thirdly, in order to do so, each party of the community has to identify the threats of its own concern. Since the parties do not necessarily trust each other (completely) they may consider different threats as relevant. Hence, we allow a *decentralized threat model*. Each party p may specify which subsystem F_p of F it trusts and against which threats it needs security. Ideally, a threat to party p is thought as an adversary environment E, which comprises the respective distrusted part of F, and a set of forbidden services $Forb(E)$. Each adversary environment E is again an executable formalism which can be combined with F_p, and the forbidden services $Forb(E)$ are negative specifications. In general, each party specifies a whole family of threats.

Finally, we can elaborate the tentative definition of security as an "even if (the F_p's are combined with adversary environments) – nothing else (than the respective required services happens) – property". Roughly speaking, we say that an executable formalism securely implements the specified required services under the stated threats, if the formalism, combined with whatever the stated adversary environments are, satisfies all required services, but none of the forbidden ones. A little bit more precisely, we define:

Let P be a set of (not necessarily disjunct) parties of a community,
$Req = (Req_p)_{p \in P}$ a family of subspecifications for services,
F an IT-system, given as an executable formalism, and for each party $p \in P$:
let Req_p be the services required by p,
F_p the subsystem of F trusted by p,
$Threat_p = (Env_p, (Forb(E))_{E \in Env_p})$ the threats concerned by p, consisting of a set Env_p of adversary environments and a corresponding family of forbidden services.

Then we say that F *securely implements* the required services *Req* for the community P under the threat model given by F_p and $Threat_p$ for all $p \in P$, iff

reliable correctness: for each $p \in P$ and each environment $E \in Env_p$ the combined system of F_p and E satisfies all services $S \in Req_p$, and

confinement: for each $p \in P$ and each environment $E \in Env_p$ the combined system of F_p and E does not satisfy any forbidden service $S \in Forb(E)$.

This framework is intended to include, e.g., error probabilities in the formulation of services or restrictions of the computational complexity of adversary environments in the formulation of threats. Furthermore, one has to take care that the set of all required and forbidden services is free of contradictions.

It should be clear that the state of art is far away from being able to apply this notion of security for many realistic examples directly. Nevertheless, the definition properly highlights the ultimate goal. Furthermore, classical aspects of security can be interpreted within our framework: for instance, availability, integrity as enforcement of semantic constraints, and non-repudiation are aspects of reliable correctness; and integrity as prevention from unauthorized modification, confidentiality, authentication, and anonymity are aspects of confinement.

3. GENERIC SERVICES OF DATABASE AND DATATRANSFER SYSTEMS

A community employs both database and datatransfer systems as technical tools for communication, i.e., to support information flow among its many and various parties (including the machines working on behalf of them). Thus, the generic purposes of database and datatransfer systems can often be described as intended information flow from some sender(s) to some receiver(s). However, there are considerable differences between the two kinds of systems:

In contrast to datatransfer systems, database systems generally do not allow “senders”, i.e., users who insert data into a database, to explicitly address some intended “receiver(s)”, i.e., users who can later on read that data as a result of a query. A database system potentially delays information flow to keep data available across time. A datatransfer system usually delivers data just in time. A database system stores its data in some virtually central place. For a distributed database system, this place is actually splitted into the repositories of its components. A datatransfer system explicitly deals with separated places, indeed it is designed to transmit data across space. Database and datatransfer systems together can overcome the two fundamental physical obstacles to human communication: time and space. Therefore, current technology is directed to integrate both kinds of systems.

We summarize the generic services which are most important to meet the generic purposes outlined above. A database system should provide: data modeling, views, querable database scheme, insertion and deletion of data, physical distribution, enforcement of semantic constraints, persistent and shared storage, data aggregation, deductive query evaluation, duplication of assembled data, transactions. A datatransfer system should provide: application specific data formats and protocols, fault tolerant and confirmed end-to-end data transport, logical addressing, synchronization.

Security matters and particularly the adoption of a decentralized threat model require databases to provide additional services like access control and encryption of shared data. Likewise datatransfer systems have to provide access control, authentication including digital signatures, as well as encrypted and anonymous datatransfer.

4. SECURITY OF IT-SYSTEMS IN HEALTH CARE

As before, many considerations can only be sketched. More instructive details and valuable insights about secure IT-systems in health care are provided by [7,8,9,10,11,12].

4.1. Purposes in health care

In health care numerous parties try to achieve various legitimate purposes. We outline some characteristics.

- Patients and physicians who communicate or cooperate by an IT-system need flexible data representations. Otherwise, the systems would be found to be obstacles to interaction and curing rather than a support. E.g., a (multinational) IT-system would by no means be satisfactory if it forced patients and physicians to express everything in only one (standard) medical language.

Patient data serves for multiple purposes, i.e. numerous authorized parties need it for several legitimate ends:

- Physicians plan treatment and medication.

- Administration and physicians within a hospital need support in accurate and timely documentation for every patient to properly place accounts and to assure continuity of treatment.
- Surgeons and organ database managers need real time available datatransfer for database queries and database updates, respectively.
- For statistical analysis and research regional, national, and international authorities, disease registers, and sometimes physicians need aggregated or sporadically identifiable data.
- For education and training cross-sections of anonymized case studies are needed.
- Hospitals' managements require aggregated data concerning their current situation with respect to finances, efficiency, quality of service, etc.
- Administration and physicians are interested to settle the accounts of financing institutions like national health services or health insurances without having to release unnecessary patient details. Settlement of accounts must be supported across countries' borders.
- In principle, patients should keep full control over their personal medical data, i.e., employing IT-systems should not weaken their autonomy in this respect.

Health care is characterized by a long tradition of self-determination and personal responsibility and accountability. Thus, many single member parties like each patient, physician, financing institution, etc. pose legitimate security purposes upon IT-systems. However, posing such security purposes is not restricted to single members, but is also reasonable for multi-member parties. For example, two physicians, one being on vacation the other acting for him have common security purposes with respect to confidentiality of patient details.

- Whoever deals with personal medical data has an interest or the obligation to keep this data confidential, i.e., to *prevent* its disclosure to unauthorized parties. If the data is transferred it should not be released to other parties than those addressed. If it is stored or archived only authorized parties should be able to process it. Since personal medical data tends to be highly sensitive (e.g., genetic information) in many cases a disclosure to unauthorized parties cannot be repaired satisfactorily (e.g., by smart-money).
- Likewise if personal medical data is received or retrieved one usually wants to be sure about its origin (authenticity), i.e. if its originator is the one he claims to be. In many situations the receiver will require a legally binding proof of the originator's identity. Checking the origin of data implies to check its integrity, i.e., if it has not been modified. Particularly in the health care world the need for data authentication is essential since fraudulent, inaccurate, or faulty data might decide over health, permanent harm, or even death.

4.2. Required services and appropriate threat models in health care

We will consider several parties, their required services and their respective threat models. Generally, service availability is reasonably required by a multi-member party comprising the service provider(s) and the service consumer(s). Confidentiality and integrity (including authentication) of data are usually required by single members *independently* from others. In the following we discuss some characteristic features of health care rather than to provide a complete list.

- Multi-member parties require specific integration of database and datatransfer services. This situation includes deductive query evaluation together with anonymous or identified remote consultation by patients. Particularly, authentication and encryption of medical data must compatibly span transfer, shared storing, deductive query evaluation, and duplication.
- The entire health care community requires a data modeling which reflects its culture of self-determination and personal responsibility. In particular, the various individual members

(patients, physicians, …) and their roles (seeking for care, treating, researching, receiving payment,…) specified by rights, obligations, etc. should be directly simulated. Not only the mere data, but, particularly, interaction of members as the source of most medical data must be reflected.

- Patients require that views, anonymity, and access control ensure that operations on medical data, although it is often multi-purpose, are suitably restricted according to the specified roles.
- Sometimes different sites transfer patient details to a central register (e.g., a disease register) which must be able to link corresponding details, but must not be able to find out the respective patient behind it. Hence, the patient and the register require aggregation of pseudonymized data.
- In case of an emergency patients suffering acutely and the respective duty physicians require that certain security measures must be overridable. In that case at least accurate and reliable logs are required.
- Usually small parties, say a patient and the health authority, require authentic data archiving (and retrieving) over long periods of time. For medical documentation archiving periods of 30 years are the rule rather than an exception.
- Legally valid receipts for answers or results produced by an IT-system are primarily required by single members who are personally accountable for actions they take or omit (e.g., physicians). Furthermore, there are parties who require legally valid proof about another party's decisions (e.g., patients who suffer from wrong treatment, etc.).

There are several *sources of threats* to such services. Firstly, identifiable medical data has to be transmitted between many and various parties and it is subject to a chain of many steps of aggregations and conclusions. It is, thus, not always obvious how sensitive aggregated or derived data still is. Secondly, in contrast to business or the military, parties of the health care world cannot be divided into “good” ones and “bad” ones in which case participation of the bad ones could simply be prohibited. Rather there is a characteristic openness and fluctuation of personnel in health care which makes it rather impossible to simply control access by multilevel clearances or closed groups of members. In fact, *any* other party can become a threat to a party's required services. Thirdly, the IT-systems installed increasingly become open and highly heterogeneous.

To have a closer look at the threats against the parties and their required services we consider an IT-system as given by some executable formalism F. The culture of self-determination and personal responsibility in health care fundamentally implies that the security of the system must be proved against a *decentralized* threat model.

For example, consider a two-member party p, consisting of a patient and a physician. An IT-system securely implements a datatransfer service available for p only if its availability does not rely upon some component numbered among p's adversary environments Env_p. An IT-system could achieve this goal by supporting datatransfer directly from a patient's device to a physician's device and vice versa.

Or consider any multi-member party p, consisting of a hospital patient and the staff of a ward. An IT-system securely implements storage of medical data confidential to p only if no-one except members of p can see that data. In particular, the company in charge of maintenance of the system must not be able to get the data. This sounds natural, but hardly any operational IT-system today is securely implemented in that sense. A more ambitious requirement could demand “confidential query evaluation” upon confidential data.

Another interesting case is data authentication. Consider any three-member party $p = (m_1, m_2, c)$, where say m_1 is a physician, m_2 is a pharmacist, and c is a court like quality assuring agent. An IT-system securely implements authenticated datatransfer from m_1 to m_2 with respect to c only if no-one (within or without p) except m_1 can produce valid authenticators (i.e., authenticators that are afterwards accepted by c) and if all authenticators produced by m_1 were accepted by c. Again, a more ambitious requirement could demand "authenticated query evaluation" upon authenticated data. E.g., even if the database manager or the maintaining company illegitimately modified data this fact could eventually be detected by end users.

5. SOME ACHIEVEMENTS AND OUTLOOK

Obviously, these reflections need further elaboration. Theoretically oriented researchers will criticize that we do not consider their specific formalization of security whereas practically oriented engineers will criticize that our proposals are hardly practical. Nevertheless, the continuous introduction of new IT-systems into health care can only be a blessing to our societies if they also fulfill all the legitimate security requirements, and if they *credibly* do so. This is the driving force why entire IT-systems must be provably correct against a specification agreed by all parties involved. Hence, the above reflections should encourage system specification activities to focus upon practically relevant security needs and encourage system development activities to ground upon formal specifications rather than to rely on ad hoc solutions. This is one of our own main topics of research.

The need for a personal model of data gave rise to the project DORIS (Data Security ORiented Information System) [13,14]. Any data is understood to be personal knowledge of some human agent, and every human agent is directly represented by an encapsulated software object. Additionally, roles and acquaintances, rights and restrictions of human agents are represented by their respective objects. The model has remained rudimentary, particularly, the concept of anonymity is missing. Nevertheless, it demonstrates that the personal model of data is feasible.

A comprehensive approach to the development of secure IT-systems in health care and their operation regarding medical, legal, and financial requirements is taken by the project SEISMED (SEcure Information Systems in MEDicine) within the EC-framework AIM (Advanced Informatics in Medicine). It provides a High Level Security Policy which states the basic values of the health care world and derives guidelines for the design and the operation of IT-systems in a clinical environment. Several technical guidelines cover security of database and datatransfer systems. Secure implementations of IT-systems with respect to an underlying decentralized threat model require cryptographic mechanisms like encipherment and digital signatures which are investigated and recommended. A prototypical software implementation demonstrates that the performance of such systems on standard hardware is satisfactory for many services [15].

These activities have to be accompanied by *technology assessment* since not only the specification of an IT-system, but also its actual integration and operation within a social environment could conflict with basic values. An inherent problem is, e.g., that new diagnostic technologies tend to force physicians to diagnose whatever possible because otherwise patients could sue them. Looking from the back end, much of the diagnoses will have been unnecessary which in turn will cause the patients to withdraw trust from their physicians. Another example are advanced medical cards which appear to be a favorite solution to many problems. In order to give patients full control over their personal medical data they could be equipped with personal medical databases by means of advanced medical cards. However, patients not only become more independent, but also become more vulnerable since they may loose their cards or may feel to be compelled to show it to third parties like employers.

Acknowledgement: We sincerely thank all colleagues of the research group "Information systems and security" at the University of Hildesheim and our SEISMED partners. Many of the ideas presented here result from valuable discussions and joint work with them.

REFERENCES

1. R. Dierstein: The concept of secure information processing systems and their basic functions; IFIP/Sec'90, Helsinki, May 1990, North-Holland, Amsterdam 1991, 133-149.
2. D.B. Parker: Restating the foundation of information security; 8th IFIP International Conference on Computer Security, IFIP/Sec '92, Singapur, September 1992, North-Holland, Amsterdam 1992, 139-151.
3. J. Biskup: Sicherheit von IT-Systemen als "sogar wenn – sonst nichts – Eigenschaft"; G. Weck, P. Horster (ed.): Verläßliche Informationssysteme, Proceedings der GI-Fachtagung VIS '93; DuD Fachbeiträge 16, Vieweg, Wiesbaden 1993, 239-254.
4. B. Pfitzmann, M. Waidner: A general framework for formal notions of "secure" systems; March 1994, submitted for publication.
5. European Computer Manufacturers Association: Standard ECMA-138; Security in Open Systems – Data Elements and Service Definitions; December 1989.
6. Commission of the EC: ITSEC: Information Technology Security Evaluation Criteria (Provisional Harmonised Criteria, Version 1.2, 28 June 1991); Office for Official Publications of the European Communities, Luxembourg 1991.
7. C.D. Stromberg: Access to hospital information: Problems and Strategies; Frontiers of Health Services Management 4 (1987), 3-33.
8. J. Biskup: Protection of Privacy and confidentiality in medical information systems: problems and guidelines; D.L. Spooner, C. Landwehr (eds.): Database Security III: Status and Prospects, Elsevier Science, Amsterdam, 1990, 13-23.
9. Commission of the EC DG XIII/F AIM (ed.): Data Protection and Confidentiality in Health Informatics – Handling Health Data in Europe in the Future, Proc. of the AIM Working Conference, Brussels, 19-21 March 1990, IOS Press, Amsterdam 1991.
10. E.H. Kluge: Advanced patient records: Some ethical and legal considerations touching medical Information Space; Methods of Information in Medicine 32 (1993), 95-103.
11. B. Barber, A.R. Bakker, S. Bengtsson (ed.): Caring for Health Information: Safety, Security and Secrecy, Elsevier Science, Amsterdam 1994.
12. N. Gaunt, F. R. France: The need for security - a clinical view; in [9], 189-194.
13. J. Biskup, H.H. Brüggemann: The personal model of data; Computers & Security 7/6 (1988), 575-597.
14. J. Biskup, H.H. Brüggemann: Das datenschutzorientierte Informationssystem DORIS: Stand der Entwicklung und Ausblick; Proc. Verläßliche Informationssysteme (VIS'91), IFB 271, Springer, Heidelberg 1991, 146-158.
15. G. Bleumer: Security for decentralised health information systems; in [11], 139-146.