Privacy Contracts incorporated in a Privacy Protection Framework

Hendrik JG Oberholzer
Tshwane University of Technology
oberholzerhjg@tut.ac.za

Martin S Olivier
University of Pretoria
http://mo.co.za

Abstract

Privacy policies do not really seem to guarantee the protection of personal information. As individuals are getting more and more concerned about the protection of their personal data, the authors propose privacy contracts as a solution. A privacy contract constitutes a set of legal agreements between an individual and an organisation and as such should be enforceable by law.

This paper extends the notion of privacy contracts. The paper discusses privacy contracts in the context of a privacy protection framework that also includes the management of privacy policies and the enforcement of policies and contracts.

The authors conclude with a comprehensive privacy protection framework built on the fundamental tenets of privacy contracts. Privacy contracts enable the individual to customise his/her personal information on a fine-grained level of detail. Subsequently the individual has more control over his/her personal information.

1. Introduction

Every individual has the right to privacy. This right to privacy is recognised in the privacy legislation of many countries [1, 2, 3]. Privacy rights comprise both the rights each person is entitled to expect and protect, and the obligations of organisations and others to respect these rights [4].

As an ever-increasing number of privacy violations surface in literature, individuals are becoming more and more concerned about the privacy of their personal information [5]. These concerns might lead to a situation where the clients do not trust the organisation any more [6, 7] and take their business elsewhere [8].

In an attempt to ease the privacy concerns of their clients and to adhere to new legislative measures, organisations publish privacy policies, stating what they would or would not do with the personal information of their clients. However, simply publishing privacy policies is not sufficient to convince potential clients to disclose their personal information to the organisation. In the context of this article, a privacy policy is one or more statements that a data controller (for example an organisation, enterprise, or any institution or public body that handles data) makes to a data subject (the individual or person whose data will be handled) stating how the data controller will handle the personal information of the data subject.

The primary question that has to be answered is therefore: What could organisations do to ensure that their clients trust the organisation to protect their personal information against potential misuse? More specifically, how can the personal information of an individual be protected against privacy violations in a relational database, while at the same time allowing business processes to be performed on the same personal information? Agrawal et al. [9] suggest that databases must take responsibility for the privacy of data as a fundamental tenet and that such databases should be called Hippocratic databases. Within the context of this article the question that is raised is: how could the Hippocratic database be enhanced to protect personal information stored in a relational database?

The second question is: How can clients exercise privacy preferences and as such have control over their personal information? More specifically, how could the perception that individuals view privacy differently be incorporated into the solution? How can guarantees (and subsequently trust) be instilled between individuals and those that need to access their personal information?

Although some work has already been done on Hippocratic databases [9, 10, 11], the main aim of this paper is to develop a Hippocratic privacy protection (from now on referred to as HPP) framework that could be used to protect personal information stored in relational databases. Our framework will allow the individual or data subject through personalisation the freedom (although not totally) to decide what could be done with his/her personal information stored in the relational database. At the same time the data controller must be allowed (see Section 3.4) to perform business transactions as required to run the business.

As an initial solution to the problems relating to privacy policies, individual privacy contracts were proposed [12] and eventually incorporated into our HPP framework. A privacy contract is a concrete and legal manner through which to contest privacy breaches, should any occur.

The personalisation of privacy contracts is realised through the concluding of privacy agreements on different levels. People have different perceptions about privacy as stated by Ishitani et al. [13]. Based on these perceptions, one can argue that data subjects require their personal information to be treated in different ways. These perceptions induced the 'concept of agreements to be consented to at different levels of privacy as preferred by each data subject'. Although contracts might lead to more personal interaction on behalf of the data subject, some data subjects might view the conclusion of a contract as too much of an effort. As a compromise data subjects have to choose between four different levels of privacy ranging from level 0 (mandatory level) to level 3 (advanced level) (see Section 3.3.6).

Although the data subject can choose the level of privacy preferred, the choice is limited. Distinction is to be made between mandatory and optional transactions, and mandatory and optional purposes. Mandatory transactions are transactions without which the data controller cannot run his business. The data subject must conclude a privacy agreement for each mandatory transaction otherwise no optional transactions could be performed against the database. An optional transaction is a transaction that is not mandatory. Every transaction has one or more purposes for which the transaction may be used, as defined in the privacy policy of the data controller. When consenting to any transaction, the data subject must consent to all mandatory purposes (if any) relating to the specific transaction otherwise the data subject is not allowed to perform the specific transaction. In contrast to mandatory purposes, an optional purpo**se** is a purpose that is not absolutely necessary for the day-to-day running of the business, but it might be to the benefit of the data controller or the data subject if the data related to a transaction could be used for the specific optional purpose. Therefore a data subject may or may not consent to an optional purpose.

Every data subject has to enter into a privacy contract (consisting of privacy agreements) with the data controller, otherwise no transactions can be performed by the data controller or between the data controller and the data subject. A data subject must consent to a privacy agreement before the data controller can use the data of the transaction associated with the agreement.

Two interacting parties could only respect and trust each other if mutual trust, based on a clear set of principles, has been established between them. Various sources in literature were found that list data protection principles. These sources were studied and an extended set of principles proposed [14]. The concept of a revised Hippocratic database was proposed into which this extended set of Hippocratic principles was incorporated and eventually modelled through an entity relationship diagram (ERD).

In order to have an overall view of the privacy protection framework (that includes privacy contracts), it is necessary to review the ERD that was developed.

The remainder of the paper is structured as follows. Section 2 discusses work related to Hippocratic databases and the principles they are modelled on as well as some earlier work based on privacy contracts [12]. Section 3 distinguishes between the three main components (identified as packages) of the framework and presents a brief description of each package followed by a more detailed discussion of the packages. The paper ends with a summary (see Section 4) and some concluding remarks.

2. Related work

As already stated, Agrawal et al. [9] suggest that databases must take responsibility for the privacy of data and and that such databases should be called Hippocratic databases. Originally, Hippocratic databases were built upon ten principles (from now on referred to as Hippocratic principles) that might become the founding principles to protect and manage personal information that reside in databases. These principles define what a data subject could expect from a database that claims to be Hippocratic. But again, promises and expectations are not good enough to guarantee privacy protection and *"Technology can't create trust where society sows doubt and disbelief, it can only hope to minimise the risk"* [15]. Thus, in order to minimise risk further, privacy contracts incorporate these Hippocratic principles. Further discussions in the remainder of this paper will indicate how these principles support privacy contracts.

The most significant difference between privacy contracts and Hippocratic databases is that privacy contracts are primarily based on transactions as opposed to Hippocratic databases that are purpose-driven [9]. In the context of this paper a transaction is defined as any kind of interaction or action (for example view, display, print, add, modify, delete, audit) that is performed on the personal information (stored in a relational database) of the data subject either by the data subject or the data user (a data user is the person who is authorised to handle the personal information of the data subject). The reason why privacy contracts are associated with transactions is, because whenever a transaction is performed, the transaction is performed on a certain subset of the personal data of a data subject. Once a data subject has consented that the data controller may perform the specific transaction, only then might the data subject consent to one or more specific reasons why the transaction may be performed on his/her personal information.

Consent associated with a transaction (performed for a specific reason) is supported by five of the Hippocratic principles namely: Purpose Specification (purpose is associated with the stored information by consenting to the transaction that uses the stored data), Consent (data subject has agreed to the purposes associated with the transaction), Limited Collection (only data needed by the transaction, for the purpose consented to, will be stored), Limited Use (stored information can only be used by the transactions consented to) and Limited Disclosure (stored information can only be disclosed through disclosing transactions consented to).

The Hippocratic Principle of Safety (i.e. personal information should be protected by security safeguards against theft and other misappropriations) should be addressed as a security issue and is not within the scope of this paper. The rest of the Hippocratic principles not listed here are addressed at later stages when more appropriate.

Massacci, Mylopoulos & Zannone [11] argue that the Hippocratic database as proposed by Agrawal et al. lacks three major principles namely hierarchies of purposes, delegations of tasks and authorisations, and the minimal disclosure of private information. The authors agree that the principles of 'Hierarchies of Purposes' and 'Delegations of Tasks and Authorisations' are covered at the level of detail that they should be, but argue that the principle of 'Minimal Disclosure of Private Information is adequately covered.

The architecture of the Hippocratic database (HD) as published by the IBM Almaden Research Center [16] involves three main components namely: the specification of the company's privacy policy in a 'privacy language'; the ability of users to specify their privacy preferences; and secure querying capabilities to enforce corporate privacy policies against users' preferences. Ashley, Powers & Schunter [17] proposed a similar framework for enterprise-wide privacy management. Although their first three components match those listed as published by the IBM Almaden Research Center, two additional components were listed that are incorporated in our proposed framework. These additional two components are the creation of an audit trail of access to privacy sensitive information and the generation of enterprise wide and individualised reports that show accesses to privacy sensitive information and conformance to the governing privacy policy. We integrated these last two components into our proposed HPP framework and do not describe them as separate components.

Our proposed HPP framework is similar to the architecture of the IBM Almaden Research Center in its general structure. Both the HPP framework and the IBM architecture are based on the fundamental concept that databases should take responsibility for the protection of personal information. However, our proposed HPP framework involves three main components: the management of privacy policies; the management of privacy contracts; and the enforcement of privacy policies and contracts. The main difference between our proposed framework and the original Hippocratic database is that the Hippocratic database is based on purposes, while our proposed framework is based on privacy contracts and privacy agreements, which are directly associated with database transactions.

The next section describes the proposed Hippocratic Privacy Protection Framework. Section 3.1 gives an overview of the revised Hippocratic database, while the subsequent sections describe the three main components of our HPP framework in more detail.

3. Privacy protection framework

The aim of this section is to introduce our proposed Hippocratic Privacy Protection (HPP) framework in order to protect the personal information stored in relational databases. This framework must enable the individual the freedom (although not totally) to decide what is allowed to be done with his/her personal information stored in the relational database. In addition, the data controller must be able to perform transactions as required to run the business. The primary question that has to be answered is therefore: How can the personal information of an individual be protected against privacy violations in a relational database without restricting the normal execution of database transactions or queries?

In answer to the stated question, a privacy protection framework was developed. A data model that supports the creation of privacy contracts based on the privacy rules is embedded in the ERD (see figure 1). Algorithms were developed in pseudocode to demonstrate how the privacy contracts are generated from the privacy policies and how these privacy contracts and privacy policies are enforced when transactions are performed on personal information stored in the relational database.

3.1. Revised Hippocratic Database

This section extends the concept of a Hippocratic database to a comprehensively revised Hippocratic database (our HPP framework) based on the extended principles (see Table 1) as proposed by Oberholzer & Olivier [14]. As most personal information or transactional data relating to an identifiable individual are generally stored in relational databases, an entity relationship diagram (ERD) has been developed in order to model these principles as metadata. The ERD (showing the entities, their attributes, and relationships) on which the proposed principles were mapped, can be seen in figure 1.

Transactions and their related purposes, based on the revised Hippocratic principles, were used to propose privacy contracts as an extension to privacy policies to protect the personal information of data subjects stored in relational databases [12]. As a privacy contract is based on privacy agreements between a data controller and a data subject, the privacy contract should be legally binding and enforceable. The data subject has control over what information may be used for what purpose when performing a specific transaction. Although the enforcement of the privacy contract resides with the data controller, data subjects exercise control over their personal information as they consent to specific transactions performed on their personal data for specific purposes. In order to perform a transaction on the data of a data subject, the data subject must consent to the data controller's use of the data associated with the specific transaction. The data controller still has to define its internal privacy policies, as the agreements of a privacy contract are based on and extracted from the installed privacy policy of the data controller.

Following is a description of the entities used to model the data (see figure 1) and the concepts on which the revised Hippocratic database was developed including a brief description of the concepts that privacy contracts are based on. The entity names are printed in italics. As these entities were implemented through algorithms that were grouped into the three main components (referred to as packages from now on), the entities will be referred to as tables from section 3 onwards.

Although Hippocratic databases are purpose-driven, our revised Hippocratic database is based on contracts, agreements and transactions. A privacy *contract* is based on a minimum set of mandatory privacy *agreement*s (directly associated with transactions) without which any privacy contract could not exist. In addition to the mandatory privacy agreements, optional privacy agreements could be added when required by the data subject. Without a privacy contract no transactions can be performed by the data controller or between the data controller and the data subject.

Data subjects and data users interact with the database through transactions. We distinguish between mandatory and optional *transaction*s where a mandatory transaction is defined as a transaction without which the data controller cannot perform any business functions. An example of a mandatory transaction might be to 'Capture Credit References'. When the data subject needs to open an account, the subject has no other choice than to consent to this mandatory transaction, otherwise no account can be opened. An example of an optional transaction might be to 'Capture Home Address' because the transaction is not really necessary in order to conduct business. Whenever a transaction is performed, the transaction is performed on a certain subset of the personal data of a data subject. Once a data subject has consented that the data controller may perform the specific transaction, only then might the data subject consent to one or more specific reasons or *purpose*s why the transaction may be performed on his/her personal information. Transactions and purposes are associated through *transactional_purpose*s. A purpose must be identified as either a mandatory or an optional purpose. The transaction 'Capture Credit References' might have a mandatory purpose 'To Verify Credibility' and an optional purpose 'To Share With Third Parties'. The subject may or may not consent to the optional purpose, but must consent to the mandatory purpose.

A transaction is limited to what data (*columns* related to *tables* and associated to a transaction through *transactional_column*) might be used. For example, for the transaction 'Capture Credit References' the columns company, clientnr, and credit_status might be used.

People have different perceptions about privacy as stated by Ishitani et al. [13]. Based on these perceptions, one can argue that data subjects require their personal information to be treated in different ways. These perceptions induced the 'concept of agreements to be consented to at different levels of privacy as preferred by each data subject'. Privacy agreement levels are defined at levels 0,1,2, or 3. The motivation for the four levels is based on two arguments, and because an agreement directly relates to one transaction, the arguments will be stated in terms of transactions and not agreements. The first argument is that central to the privacy protection framework a distinction has to be made between mandatory and optional transactions - thus an initial need for two levels or categories. Regarding the level of optional transactions, data subjects have the opportunity to exercise some choices. Based on the different perceptions that data subjects have about privacy and the protection of personal data, Westin [18] divided data subjects into three groups or categories namely the privacy fundamentalists (usually unwilling to share information freely), marginally concerned (willing to share information), and pragmatic majority (people fitting somewhere between the other two groups). Each of these groups is then assigned a level. Therefore four levels are needed – one level (0) for mandatory transactions and three levels (1,2,3) for optional transactions.

The minimum set of mandatory transactions (referred to earlier), are defined at privacy level 0. Marginally concerned data subjects may prefer privacy agreements (based on fair information practices) defined at level 1. For every optional transaction, the data subject can decide what level of privacy is preferred for the specific transaction to be performed on his/her data. These perceptions are represented by entities ATP2 and ATP3. The entity *ATP2* defines privacy agreements consented to at privacy level 2. Level 2 agreements give the pragmatic data subject the choice to consent to optional purposes that a transaction may be used for at level 2. At privacy level 3 the privacy fundamentalist can decide which attributes as defined in entity *column*s can be used for the specific transaction that was consented to at level 2 for a specific purpose. These privacy agreements are represented by entity *ATP3*.

Each time a data user accesses an agreement reached between the data subject and data controller, the access is logged in *agreement_audit_log* for auditing purposes.

Data *users* are authorised for specific *access_kind*s through the association *transactional_user* on a transaction.

The revised Hippocratic database also provides for the specification of the *privacy_law*s that a specific transaction must adhere and relate to (see entity *transactional_law*). P*rivacy_breach_resolution*s are also incorporated to inform the data subject what steps could be taken if a privacy breach has occurred and how the data controller will handle such privacy breaches.

Now that we have introduced the data model by describing the entities that belong to the ERD (see figure 1) we can focus our attention to the main components of our proposed Hippocratic privacy protection framework. The HPP framework is partitioned into three packages, each of which addresses a particular component of the framework. The first package named MANAGE_PRIVACY_POLICIES (see Section 3.2) is mainly used to implement the privacy policies defined by the data controller. The second package named MANAGE_PRIVACY_CONTRACTS (see Section 3.3) enables the creation and manipulation of a single privacy contract for every data subject that needs to interact with the data controller. It also allows for the individual customisation of the privacy contract through privacy agreements concluded for a transaction by the data subject. The third package named ENFORCE_POLICIES_AND_CONTRACTS (see Section 3.4) protects the personal information of the data subject through enforcement of the customised privacy contract that belongs to the specific data subject, while at the same time enforcing the privacy policies. This package includes algorithms to be used when personal or business transactions are processed and the obligations to be performed when the data controller updates a privacy policy.

The following subsections describe and discuss these packages in more detail.

3.2. MANAGE_PRIVACY_POLICIES

From an organisational point of view, a data controller normally publishes a general privacy policy stating how personal or private information of the data subjects will be handled. In addition to the privacy statements proclaimed by the data controller, data subjects sometimes have the opportunity to opt-in or opt-out from certain uses of their personal information, but these opportunities are limited and do not allow for fine-grained levels of choices or consent. In addition, privacy policies might be misleading [19], they are too lengthy and written in a legalistic way that is both difficult to comprehend (Kobsa, as cited by [20]) and boring to read [21]. Many Internet users consent to these policies without reading them by just clicking on the ‘I agree’ or ‘OK’ button simply in order to be able to continue with the session or transaction.

This package mainly consists of two components namely INSTALL_PRIVACY_POLICIES and MAINTAIN_PRIVACY_POLICIES.

3.2.1. Install privacy policies

Privacy policies are installed after they have been defined by the data controller. Installation consists of an automated process that creates the relational tables needed for the policies, contracts, and audit logs, followed by addition of privacy policy elements as well as associations between transactions and some of these policy elements. The addition of privacy policy elements as well as associations between transactions and some of these policy elements is realised through SQL ‘Insert’ statements.

3.2.2. Maintain privacy policies

Maintaining privacy policies consists of adding, modifying, and deleting privacy policy elements and associations between policy elements. The adding, modifying, and deleting relates respectively to the SQL ‘Insert’, ‘Update’, and ‘Delete from’ statements while enforcing database integrity constraints simultaneously. Before an element or association can be added a check is performed to verify that such an element or association does not already exist. The elements that can be added are kinds of access (read, modify add, delete), columns that belong to a specific table, privacy laws, purposes, privacy resolutions (in case of a privacy breach), tables, transactions, and data users. The associations of transactions with purposes, columns, privacy laws, and data users are also handled by adding the associations as privacy policies and are also stored and maintained in relational database table structures.

When adding a new transaction, it is critical to distinguish between mandatory and optional transactions. If the transaction is mandatory, then all existing privacy contracts must be frozen (see attribute *frozen* in entity *contract*) . This means that data subjects with existing privacy contracts must first consent to the new transaction (policy) that has been added before the subjects would be allowed to commence with any other transactions against the database. If the data subjects do not consent to the new mandatory transaction, their privacy contracts will stay frozen, otherwise their contracts will be unfrozen and they would be able to commence with transactions against the database as before. The data controller will only be able to handle data according to privacy agreements of contracts that are not frozen. For those contracts that are frozen, the data controller will only be able to handle data according to privacy agreements that were cocluded before the date when the new mandatory transaction was added or modified.

When adding a new association between an optional transaction and a mandatory purpose all existing contracts that include this optional transaction must be frozen. The data related to this transaction may not be used for this purpose, because the data subjects whose contracts include this transaction have not yet consented to the new mandatory purpose.

When adding an association between a data user and a transaction, one has to specify the kind of access. A prerequisite is that only a privileged data user (for example the chief privacy officer), who is authorised to manage other data users, may perform this algorithm.

Deleting a column from a privacy policy entails the removal of the policy from the relevant table(s) and deleting from each privacy contract the privacy agreements associated with the specific policy. The data controller will not be able to perform any transactions on those privacy policies and corresponding privacy agreements that have been deleted. As the deletion of a privacy agreement does not affect the data subjects negatively, data subjects need not be informed that the policy has been deleted. Deletion of the associated privacy agreements from the privacy contracts of the data subjects affected, will only be logged in table *agreement_audit_log* for reference or auditing purposes.

Deleting a purpose induces several chain reactions. If the agreement for the transaction associated with this purpose is at level 3, then *ATP3* and *ATP2* agreements have to be deleted from those contracts. If the specific contract still has other purposes for the same transaction in *ATP3*, then the level stays at 3, otherwise the level must be changed to 1. If the agreement for the transaction associated with this purpose is at level 2, then *ATP2* agreements have to be deleted from those contracts. If the specific contract still has other purposes for the same transaction in *ATP2*, then the level stays at 2, otherwise the level must be changed to 1. *ATP2* agreements have to be deleted (if any) that include the purpose to be deleted. Following these steps, the transactional purposes containing this purpose must be deleted. Only then can the specific purpose be deleted from the privacy policy. While all the deletions listed above are taking place, logs must be kept for auditing purposes.

When deleting a transaction from the privacy policy of the data controller, all privacy contract agreements that include the specific transaction must be deleted and logged for auditing purposes. All associations between the specific transaction and related policy elements must be deleted and logged for auditing purposes. Only then can the transaction be deleted.

However, modification of a policy is not simple. If a mandatory transaction is modified all existing privacy contracts are affected, but new privacy contracts will be handled as if the policy always existed. As soon as an existing privacy contract is authenticated, it must be determined whether any policy associated with a mandatory transaction has changed. For every modified policy associated with a mandatory transaction, the data subject must be informed about the modification, and the data subject must consent to the modification. If the data subject does not consent to the modification, the privacy contract of the data subject must be frozen and no more transactions can be performed by the data subject.

The mandatory status of a transaction can be changed to mandatory or not mandatory (or optional). When changed to not mandatory, one has to log the change to the audit log of the privacy policies. However, when changed to mandatory, all affected contracts (those contracts that do not have an agreement for the specific transaction) must be frozen and logged for auditing purposes in the audit log of the privacy contracts.

To conclude, IT developers are not legal professionals and might find it difficult to interpret and implement all the legal and regulatory requirements required by law to protect personal information. Therefore, teams of privacy professionals and legal administrators should collaborate with IT developers to embed the most applicable privacy legislation and regulations in the definition of privacy policies and statements to be enforced through privacy rules into business applications.

3.3. MANAGE_PRIVACY_CONTRACTS

This section introduces a privacy contract composed of a number of formally signed (electronic, digital, or handwritten) privacy agreements between a data controller and a data subject. As the composition of the privacy agreements constitutes a contract, the privacy contract should be legally binding and enforceable. The data subject has control over what information may be used for what purpose when performing a specific transaction. Although the enforcement of the privacy contract resides with the data controller, data subjects exercise control over their personal information as they consent to specific transactions performed on their personal data for specific purposes. At any time data subjects can view all changes made against their privacy contracts as well as view all disclosures of their privacy agreements. This information can be retrieved from table *agreement_audit_log*.

The concepts discussed in the following subsections (see Section 3.3.1 to 3.3.8) mainly relate to privacy contracts. The algorithms that were developed for this package address issues of security and privacy, contracts, transactions, purposes, perceptions, agreement levels, authentication, and maintenance of privacy contracts. As these concepts are closely related, their discussion logically follows next.

3.3.1. Security and privacy

While security addresses the issue of protecting data against unauthorized disclosure, alteration, or destruction [22:504], privacy is generally accepted as the right to protect one against unwanted publicity or the right to be left alone.

In 1967, Alan Westin defined informational privacy (as cited by [23:205]) as ‘the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others’, whether the information might be electronically stored, verbally communicated, or in written or printed format. Significant from this definition are the issues of ‘determining for themselves, when, how, and to what extent’. The issue of ‘determining for themselves’ directly relates to the Principle of Consent [9] stating that data subjects should have a choice with regard to the use of their personal information. Inclusion of the word ‘when’ could relate to reasons why personal information are used, whether it will be to stored, processed, disclosed or shared. The words ‘what extent’ could relate to the Principle of Limited Use when information is revealed while privacy is preserved at the same time.

The idea or concept of a privacy contract was partially derived from a statement that Lau et al. [15] made when they stated that since no universal policy is appropriate for all users, designers must provide users with some means of specifying their own individual privacy policies. This means that data subjects not only have to consent or withhold consent to what is presented or promised to them, but data subjects must be able to customise their privacy preferences at an individual level (see Section 3.3.6). The concept of individual privacy policies, combined with the principles of Hippocratic databases, culminated in the concept of privacy contracts.

Privacy contracts are not the same as privacy policies, but are extensions of privacy policies. As an extension to privacy policies, each privacy contract is based on these privacy statements as defined by the data controller in the privacy policy of the organisation. It should be evident that before any privacy contract can be established, the privacy policy must first be defined and installed by the data controller.

Privacy contracting is based on database privacy where the definition of informational privacy is narrowed down as the personal and private information of individuals as electronically stored in databases. Personal information is information that is generally known about a person, for example the first name and surname, gender, address, and telephone number of a person. Private information is information that is not generally known about the person, for example the HIV status of the person, or monthly income.

As already stated, this paper is concerned with database privacy and it should, therefore, be clear that the security of the database must be intact – if security fails, database privacy would not be possible. A data user must be authenticated, then authorised, and then allowed to access personal and private information of individuals who saved their consent with certain conditions. While both security and privacy rules and procedures must be established and adhered to by the data controller, individuals whose information is stored in the database do not have any part in the security of the database and will have to rely on the measures implemented by the data controller, otherwise individuals will have to opt-out from the transaction.

In cases of both security and privacy, it is necessary that audit trails be kept of all accesses to personal and private information. In the case of security, an audit trail might indicate that someone has tampered with the database. In the case of privacy, an audit trail might indicate that the data controller or an authorised data user has not honoured the privacy preferences as indicated by a data subject. Typical information recorded for auditing purposes might include the transaction (or query) that was performed, the user that performed the query, date and time of the query, and the reason for the transaction.

3.3.2. Privacy contracts

As stated earlier, a privacy contract consists of a minimum number of mandatory privacy agreements. In order to perform a transaction on the data of a data subject, the data subject must consent to the data controller's use of the data associated with the specific transaction. This is done through the establishment of a privacy agreement associated with the specific transaction. As already stated - before any transaction can be performed, a privacy agreement associated with the specific transaction must exist, or otherwise a privacy agreement must first be established. Central to the concept of privacy contracts is the distinction between mandatory and optional transactions.

3.3.3. Transactions

As already stated, mandatory transactions are transactions without which the data controller cannot run his business. The data subject must conclude a privacy agreement for each mandatory transaction otherwise no optional transactions could be performed against the database. An optional transaction is a transaction that is not mandatory. It depends on every data subject whether to perform such an optional transaction or not. As soon as the data subject has consented to all mandatory agreements the subject can set privacy preferences for any optional transaction(s) through the closing of an optional privacy agreement associated with each of the specific optional transactions. However, closing of an optional privacy agreement depends on the data subject to consent to the optional transaction at any time, as long as consent is given before the optional transaction is performed, otherwise the optional transaction would not be possible.

3.3.4. Purposes

Every transaction has one or more purposes for which the transaction may be used, as defined in the privacy policy of the data controller. A purpose may be mandatory or optional. When consenting to any transaction, the data subject must consent to all mandatory purposes (if any) relating to the specific transaction otherwise the data subject is not allowed to perform the specific transaction. In contrast to mandatory purposes, an optional purpo**se** is a purpose that is not absolutely necessary for the day-to-day running of the business, but it might be to the benefit of the data controller or the data subject if the data related to a transaction could be used for the specific optional purpose. Therefore a data subject may or may not consent to an optional purpose. A data controller is allowed to use the data relating to a transaction only for the purposes that the data subject has consented to. These purposes consented to by each data subject are saved in the database.

3.3.5. Privacy perceptions

It has already been stated that people have different perceptions about privacy [13]. Some individuals might view a specific act as an invasion of their privacy, while other individuals view the same act as acceptable, and some individuals might not even be aware that they are eroding their own privacy. Based on these perceptions, one can argue that data subjects require their personal information to be treated in different ways. These perceptions induced the concept of agreements to be consented to at different levels of privacy as preferred by each data subject - for every optional transaction, every data subject can decide what level of privacy is preferred for the specific transaction to be performed on his/her data.

Before a data subject can request a specific transaction to be performed on his/her data, or before a data controller can perform a transaction on the data of the data subject, the data subject must already have entered into a privacy agreement associated with the specific transaction. As these agreements would be enforced (through privacy contracts) according to the privacy framework (see Section 3.3) it enables data subjects to exert control over their personal information.

3.3.6. Privacy agreement levels

As already stated, privacy agreement levels are defined at levels 0, 1, 2, or 3. Mandatory transactions are defined at privacy level 0 and their corresponding privacy agreements must be included in the privacy contract of every data subject. A data subject must consent to all mandatory transactions, otherwise a privacy contract cannot be concluded and no transactions can commence. Furthermore, every mandatory agreement must be consented to before any other optional transactions can be consented to and/or performed.

Data subjects must specify their own preferred privacy level for every optional transaction. Based on the different perceptions that data subjects have about privacy and their different levels of privacy awareness, optional transactions can be consented to at privacy levels 1, 2, or 3 and are stored in table *agreement*.

Privacy agreements at level 1 (standard level) aim to support the category of marginally concerned data subjects in protecting their personal information based on fair information practices. In addition to the mandatory transactions defined at privacy level 0, optional transactions can be consented to at privacy level 1. Optional transactions consented to at privacy level 1 are associated with mandatory purposes only. No opt-in or opt-out choices are available, except for the choice whether to allow the data controller to use or not to use the data of the optional transaction for the consented mandatory purpose(s) as specified in the privacy policy of the data controller. Thus, level 1 provides a take-it-as-it-is-or-not option for every optional transaction.

Privacy agreements at level 2 (medium level) are stored in table *ATP2* and aim to support the pragmatically concerned data subjects. In addition to mandatory purposes that might be defined for an optional transaction, optional transactions at this level provide data subjects with the choice of opting-in or opting-out of optional purposes relating to each optional transaction. The data controller may use the data on which the optional transaction is performed only for the purposes that the data subject has consented to, whether these purposes are mandatory or optional.

At level 3 (advanced level) the privacy model aims to support the privacy fundamentalists. These agreements are stored in table *ATP3*. In addition to the support that level 2 provides to data subjects, at level 3 data subjects could specify for each optional transaction and for each related optional purpose consented to, what data the data controller may or may not use for that specific optional purpose. Because the purpose is optional, it is the prerogative of the data subject to decide what data pertaining to the transaction may or may not be used.

3.3.7. Authentication

The management of privacy contracts mainly consists of authenticating privacy contracts, creating privacy contracts, updating privacy contracts, viewing and printing privacy contracts, and support for the auditing of privacy contracts.

Each time a data subject interacts with the data controller the first step is to authenticate the data subject - the objective is to verify that the data subject has a privacy contract and that the contract belongs to the specific data subject and not to anybody else. As soon as the data subject has been authenticated and it has been determined that the contract is active (see the attribute *active* in table *contract*), the data subject must choose the transaction to be performed. Active contracts are contracts against which database transactions (those agreed to) may be performed.

3.3.8. Maintaining privacy contracts

Creation of a privacy contract mainly consists of three steps: initialising the contract, consenting to mandatory agreements, and compiling the contract (including the privacy agreements consented to). During initialisation of the contract a contract number and password are assigned to the data subject and the status of the contract is set to '*active*'. Instead of consenting to a long list of policy statements that cover the whole spectrum of ways the personal information of the data subject will be protected, the data subject initially only has to consent to a small number of mandatory agreements when creating a privacy contract. When the data subject has consented to all the mandatory agreements, the contract is created. Thereafter, privacy agreements can be entered into with regard to the optional transactions.

As soon as the privacy contract has been created, a copy of the contract (including all mandatory agreements and optional agreements, if any) is compiled and signed. At any later stage the data subject has the opportunity to add, modify, view, or print the privacy agreements again.

Whenever a data subject requires a transaction to be performed, the first step is to verify that a privacy agreement related to the specific transaction does exist. If the agreement does not exist, an agreement could be added.

It might happen that a privacy agreement must be modified where the data subject wants to lower or increase the level of the privacy agreement. When the level is to be lowered, higher-level fine-grained detail must be deleted. When the level is to be increased, higher-level fine-grained detail must be added (see Section 3.3.6). However, it is also possible that the level stays the same, in which case optional purposes may be added or deleted from level 2, or data items added or deleted for a specific purpose at level 3.

When privacy agreements are to be deleted, care must be taken that no mandatory purposes are deleted at level 2, or no mandatory transactions are deleted from level 0. The Principle of Retention (data is kept as long as is necessary) is adhered to through the option to delete privacy agreements and the setting of the status of a privacy contract.

It must be noted that whenever the data controller makes any changes to the privacy policy, all affected privacy agreements stay as they are until consent has been given for the changed (new) agreements. The issue of changing privacy policies is handled by the component that manages privacy policies, and getting renewed consent is handled by the package that enforces privacy contracts (see next Section).

3.4. ENFORCE_POLICIES_AND_CONTRACTS

Conflict in privacy practices occurs when data subjects have to reveal personal information, while they want to preserve their privacy at the same time. Miller [24] asks how much information data subjects will be willing to sacrifice for the convenience of using the Internet and customising Internet-related services according to their personal preferences.

However, this trade-off between privacy and convenience is not only related to the world of the Internet, but also to the real world of brick-and-mortar companies and data being stored in relational databases. In order to open an account, you not only have to supply personal information like your home address, but also very private information like your banking details and gross income. This information is needed in order to approve or deny the application. Eventually, the data subject still has the choice either to disclose all the information or not to have the convenience of an account.

Another view of this problem of revealing information while protecting privacy is that of authorising data controllers to process personal information of data subjects for certain purposes or reasons. The problem is not really the issue of security, but that of privacy. Security measures will authenticate a data user and then authorise the user (through the granting of roles and accompanying privileges) to access the personal information of a group of data subjects. Privacy measures must ensure that an authorised data user has a legitimate reason or purpose to access certain personal information of a specific data subject. Alternatively, data users who have legitimate reasons to access personal information of a specific data subject should not be prevented access to this personal information.

The purpose of this package is to enforce the privacy policies of data controllers and the privacy contracts of data subjects. Distinction has to be made between performing business transactions (see Section 3.4.1) or personal transactions (see Section 3.4.2).

3.4.1. Business Transactions

Business transactions are transactions that may be performed for, but not requested by, an individual data subject, or for a group of data subjects according to the personal privacy agreements defined in the privacy contracts of the data subjects. These business transactions are performed by authorised data users that have been granted the right by the data controller to perform one or more specific business transactions. When a data user initiates a business transaction, the data user has to indicate whether the transaction type is a group or individual business transaction type, the code of the transaction, and the kind of access needed. The privacy policies are checked to ensure that the data user is allowed to perform the specific transaction. Authorisation to perform a business transaction and the kind of access authorised for the user to perform the specific transaction are stored in the associated policy table *transactional_user*. The kinds of access (read, write, modify, insert, delete) for which a data user may be authorised are stored in table *access_kind*. As soon as it has been verified that the data user is authorised, the transaction can then be performed on the data of those data subjects that concluded privacy agreements relating to the specific transaction. If the specific agreement does exist, the transaction can take place according to the agreement concluded. Otherwise a list of authorised transactions with related kinds of access that the data user is granted, is displayed.

Firstly, group business transactions are performed for a group of data subjects for one purpose at a time. Only data of those data subjects that have consented to the specific transaction to be used for the specific purpose is included in the transaction. Group business transactions performed are not logged for every individual data subject whose data was part of the group processing, but the transaction must be logged to indicate that the data user performed such a transaction. The transaction is logged for auditing purposes (see table *transaction_audit_log*).

A check must be performed to verify that the purpose belongs to the transaction to be performed. The date when the purpose was added to the policy, and whether the purpose is mandatory or not, is retrieved from the policy (see attributes *tp_mandatory* and *purpose_ date* in associated table *transactional_purpose*). The date and the mandatory purpose indicator is used to determine why some of the privacy contracts might be frozen. Only contracts that are not frozen are allowed to be part of the transaction. The indication whether the transaction itself is mandatory or not must also be retrieved from the policy. If the transaction to be performed is a mandatory transaction, then only the data belonging to data subjects with unfrozen privacy contracts may be processed. Furthermore, if it is a mandatory transaction, then the agreements were concluded at privacy level 0 and all purposes associated with the transaction are acceptable. Levels 1 to 3 are not relevant and are not taken into account.

For transactions that are not mandatory, the algorithm has to determine whether the purpose for which the optional transaction is to be performed is mandatory or not. For mandatory purposes, there is a possibility that this purpose could have been added after a data subject consented to the transaction. Therefore, the transaction can only be performed on unfrozen contracts. Furthermore, privacy agreements at level 3 are not allowed for mandatory purposes – only for optional purposes. Lastly, as this is an optional transaction that is to be performed on a group of data subjects, the algorithm will only retrieve the contract numbers of those data subjects on which the optional transaction can be performed for the mandatory purpose (those privacy contracts that are not frozen and of which the privacy agreement level is 1 or 2 for the specific transaction).

Transactions that are not mandatory and of which the purpose is not mandatory, can be performed on frozen or unfrozen contracts – therefore, the frozen status of the privacy contract of the individual is not applicable. Furthermore, as this is an optional transaction that is to be performed on a group of data subjects, the algorithm will only retrieve the contract numbers of those data subjects on which the optional transaction can be performed for the optional purpose (those privacy contracts of which the privacy agreement level is 1-3 for the specific transaction). For privacy agreements at level 2, only those agreements that include the specific purpose may be used. However, because the purpose comes into play at privacy level 3 when columns are consented to, agreements at level 3 complicate matters in the sense that not every user might have consented to the same data columns to be taken into account for the specific transaction. It might easily be accomplished by performing a join operation between tables *contract*, *ATP2* (agreements consented at privacy level 2) and *ATP3* (agreements consented at privacy level 3). Eventually the transaction is logged in table *transaction_audit_log* for auditing purposes.

Secondly, individual business transactions are transactions that are performed by an authorised data user on the data of one individual data subject. Although the data subject has not requested the transaction, the data subject however consented to the transaction. As soon as the data user chooses a transaction to be performed, the privacy contract of the data subject is checked to see whether the subject has concluded a privacy agreement that relates to the specific transaction. If the specific agreement does exist, the transaction can take place according to the agreement consented to.

Once again the data user first has to check whether the contract of the data subject is frozen or not. If the contract is not frozen, the individual business transaction can be performed normally. This means that the transaction has to be performed at the privacy agreement level consented to by the data subject. At levels 0 and 1 (standard level), the purposes (table *tal_purpose*) are listed why the individual transaction may be performed and the data user must indicate the purpose the transaction is intended for. The data user has already been authorised to perform the transaction. The transaction code and purpose are logged for auditing purposes and normal processing of the transaction follows. At level 2 (intermediate level), there is no main difference in the processing, except the agreements are retrieved from table *ATP2*. The same applies to level 3 (advanced level) except that the data columns consented to for processing are retrieved from table *ATP3*.

When the contract of the data subject is frozen, one has to determine whether the contract has been frozen as a result of the policy of the current transaction being modified after the data subject concluded a privacy agreement on the current transaction. It might be that a mandatory transaction has been added after the data subject last updated his/her contract and the data subject has not yet consented to the newly added mandatory transaction - therefore, the transaction cannot be performed on the data of the data subject. However, it might be that the transaction that caused the freezing is not mandatory. This means that a mandatory purpose might have been added to the transaction and the data subject has not yet consented to this mandatory purpose. The data user must check to see whether the mandatory purpose that was added, and that caused the freezing, belongs to this transaction. If so, the data user cannot perform this transaction. If not, the transaction can be performed normally (see previous paragraph).

3.4.2. Personal transactions

In contrast to business transactions, personal transactions allow a data subject to perform transactions according to personal privacy agreements contained in the privacy contract of the data subject. Personal transactions are performed on behalf of an individual data subject, but only on request of the data subject, for example ‘Purchase Medicine on Account’ or ‘Pay Account’.

As soon as the data subject has been authenticated and his/her contract is active and not frozen, the data subject may choose a transaction to be performed. The privacy contract of the data subject is checked to see whether the subject has concluded a privacy agreement that relates to the specific transaction. If the specific agreement does exist, the transaction can take place according to the agreement concluded. If the specific agreement does not exist, a new agreement might be established and added to the privacy contract of the data subject.

The data user does not really need to be authorised to perform a specific personal transaction because a contract number and password is needed to perform a personal transaction – the password of the contract is only known to the data subject and provided when the data subject requests the specific transaction to be performed by the data user.

The kind of access a data user is allowed to perform on a personal transaction is also not applicable as a data subject can request any personal transaction to be performed as long as a privacy agreement exists between the data subject and the data controller for the transaction to be performed.

4. Summary

Organisations collect large amounts of personal data about their clients and process this data into supposedly useful information. The term 'supposedly' is used, because some of this personal information might be of a very sensitive nature. If such sensitive information is used for purposes not intended by the client (defined as misuse), it might lead to unintentional or malicious disclosure [25].

More often than not, organisations do not realise the potential risk associated with the storage and processing of sensitive personal information until an audit reveals who has accessed this information [26]. In addition to audits performed, literature surveys reveal numerous cases where personal information of individuals was misused [27, 28, 29, 30]. It is therefore not unexpected that there is a growing concern among individuals about the ever-increasing list of privacy violations that occur and the realisation that individuals do not trust those that collect and handle their personal information [31].

When party A has to disclose personal or private information to party B the question that arises is whether A can trust B. Normally when A does not trust B, A is not going to do business with B. Consequently, individuals that do not trust the organisation might provide incomplete, inaccurate, or misleading information [32] or might not even be willing to disclose any of their personal information to the organisation [33], in which case the organisation might lose client value. However, individuals might be willing to share their personal information if they trust the organisation and perceive reasonable benefits in return [34].

In an attempt to instil trust, organisations publish their privacy policies to disclose how they will handle the data of their clients. In the context of this article, a privacy policy is one or more statements that a data controller (for example an organisation, enterprise, or any institution or public body that handles data) makes to a data subject (the individual or person whose data will be handled) stating how the data controller will handle the personal information of the data subject.

At the outset of this paper the question was raised how the personal information of an individual could be protected against privacy violations in a relational database, while at the same time allowing business processes to be performed on the same personal information? Secondly, based on the different perceptions that individuals have about privacy, how could clients exercise such privacy preferences and as such have control over their personal information?

As a solution to the above-mentioned problems, a privacy protection framework based on privacy contracts was develop to protect the personal information stored in relational databases and to allow the individual the freedom (although not totally) to decide what could be done with his/her personal information stored in the relational database. The framework was introduced through three packages. The first of these packages entails the management of privacy policies requiring the data controller to develop and install its privacy policies as relational tables. The second package focuses on the fundamental concept of privacy contracts that consist of agreements, transactions, and purposes. A data subject must have entered into a privacy contract with the data controller, otherwise no transactions can take place between the two parties. The third package enforces the privacy policies and privacy contracts as the data controller may only use the transactional and personal information as consented to by the data subject through the privacy agreements concluded between the two parties. As these agreements can be concluded at different levels of privacy, the data subject has more control over his/her information.

This article paved the way for an alternative approach to privacy protection and supports the tenets on which Hippocratic databases are based as this work can be integrated with such databases.

Although it might seem that privacy contracts involve less work for or interaction from the end user, the emphasis should be on the fact that the end user controls only the data directly related to the transactions he or she wishes to perform. No lengthy agreements have to be completed prior to interaction – only a small set of mandatory transactions has to be consented to.

With regard to implicit consent, nothing is left to chance – the concept of mandatory transactions supports this notion. The data controller can only perform those transactions consented to by the data subject. Should new regulatory measures be introduced or privacy laws be enacted, the privacy contracts of those individuals that have not yet consented to these changes could be frozen until they have consented to these changes or their data may not be used.

Lastly, privacy contracting (as mentioned in the title) is an extension of privacy policies. As long as the data controller or his/her employees act or intend to act maliciously, there could be no guarantee that even privacy contracting will solve the privacy problem effectively.

As this article introduced privacy contracting at a conceptual level, algorithms have already been developed at a lower level (based on a relational database scheme) to guide the implementation of the privacy contracts in relational database management systems. Through the algorithms it became clear that with little effort, privacy contracts could also be implemented in existing relational databases. The algorithms developed could also be adapted to implement privacy contracts when developing web applications.

Future work anticipated by the authors is similar to that of LeFevre et al. [10]. They present a practical approach based on the principle of 'Limited Disclosure'. Our research will focus on the development of a hierarchical privacy sensitive filtering technique where data subjects have the opportunity to specify how sensitive they are about individual items of their personal data. Depending on the level of sensitivity data will then be disclosed or not at different levels of sensitivity.

5. References

[1] REPUBLIC OF ARGENTINA. 2000. Personal Data Protection Act 25.326 [Online]. Privacy International. Available from: http://www.privacyinternational.org/countries/argentina/argentine-dpa.html [Accessed: 11/11/2004].

[2] KINGDOM OF THE NETHERLANDS. 2000. Personal Data Protection Act: unofficial translation: Upper House of the Dutch Parliament: session 1999-2000, nr. 92: 25 892 rules for the protection of personal data [Online]. Planet Internet. Available from: http://home.planet.nl/~privacy1/wbp_en_rev.htm [Accessed: 8/8/2004].

[3] REPUBLIC OF SOUTH AFRICA. 2002. *Electronic Communications and Transactions Act, No. 25 of 2002* [Online]. Available from: http://www.acts.co.za/ect_act/index.htm [Accessed: 27/07/2004].

[4] AUSTRALIAN PRIVACY CHARTER COUNCIL. 1994. *Australian Privacy Charter* [Online]. Available from: http://www.anu.edu.au/people/Roger.Clarke/DV/PrivacyCharter.html [Accessed: 07/06/2004.

[5] KARJOTH, G., SCHUNTER, M. & WAIDNER, M. 2003. Platform for enterprise privacy practices: privacy-enabled management of customer data. In: *2nd Workshop on Privacy Enhancing Technologies (PET 2002), Lecture Notes in Computer Science.* S.l.: Springer:69-84.

[6] NYKAMP, M. & MCEACHERN, C. 2001. Customer relationship report: privacy, CRM and ROI. *DM Review*[Online], Feb. 2001. Available from: http://www.dmreview.com/editorial/dmreview/print_action.cfm?articleid=3014 [Accessed: 22/07/2004].

[7] VILA, T., GREENSTADT, R. & MOLNAR, D. 2003. Why we can't be bothered to read privacy policies: models of privacy economics as a lemons market. In: P*roceedings of the 5th International Conference on Electronic Commerce, held in Pittsburgh, PA.* New York, N.Y.: ACM:403-407. (ACM International Conference Proceedings Series, vol. 50).

[8] FAVILLA, E. 2004. Companies and consumers clash on privacy issues. *DestinationCRM* [Online], Jan. 29. Available from: http://www.destinationcrm.com/articles/default.asp?articleid=3825. [Accessed: 18/02/2004].

[9] AGRAWAL, R., KIERNAN, J., SRIKANT, R. & XU, Y. 2002. Hippocratic databases. In: *Proceedings of the 28th VLDB Conference, 20-23 August 2002, Hong Kong, China.*

[10] LEFEVRE, K., AGRAWAL, R., ERCEGOVAC, V., RAMAKRISHNAN, R., XU, Y. & DEWITT, D. 2004. Limiting disclosure in Hippocratic Databases. In: *Proceedings of the 30th VLDB Conference, Toronto, Canada.*

[11] MASSACCI, F., MYLOPOULOS, J. & ZANNONE, N. Minimal disclosure in hierarchical Hippocratic databases with delegation. [Online]. Available from: http://dit.unitn.it/~zannone/publication/abstract/mass-mylo-zann-05-ESORICS.html [Accessed: 14/08/2005].

[12] OBERHOLZER, H.J.G. & OLIVIER, M.S. 2005a. Privacy contracts as en extension of privacy policies. In:

Proceedings of the International Workshop on Privacy Data Management, PDM 2005, April 9 2005, Tokyo, Japan, in conjunction with IEEE ICDE 2005. S.l.:s.n.:11-19.

[13] ISHITANI, L., ALMEIDA, V. & MEIRA, W. 2003. Masks: bringing anonymity and personalization together. *IEEE Security & Privacy,* 1(3), May/June:18-23.

[14] OBERHOLZER, H.J.G. & OLIVIER, M.S. 2005b. Principles of privacy protection: a revised hippocratic approach supported by an entity relationship diagram. [Pretoria: TUT]. (Technical Report; FRC/29/08/05).

[15] LAU, T., ETZIONI, O. & WELD, D.S. 1999. Privacy interfaces for information management. *Communications of the ACM,* 42(10), Oct.:89-94.

[16] IBM Almaden Research Center. 2004. *Intelligent information systems: Hippocratic database* [Online]. IBM Almaden Research Center. Available from:
http://www.almaden.ibm.com/software/quest/Projects/hippodb/activeenf/
[Accessed: 19/01/2004].

[17] ASHLEY, P., POWERS, C. & SCHUNTER, M. 2002. From privacy promises to privacy management: a new approach for enforcing privacy throughout an enterprise. In: *New Security Paradigms Workshop '02, 23-26 September, Virginia Beach, Virginia.* New York, NY: ACM Press:43-50.

[18] WESTIN, A.F. 1991. Harris-Equifax consumer privacy survey 1991. *GA:Equifax Inc*, Atlanta.

[19] HILL, K. 2004. Online privacy policies misleading. *CRMDaily* [Online], Jan. 21. Available from: http://www.crm-daily.newsfactor.com/perl/story/23036.html [Accessed: 22/01/2004].

[20] IRVINE, U.C. 2004. Clear privacy practices boost trust and online sales. *ScienceBlog* [Online], Aug. 30. Available from: http://www.scienceblog.com/community/article-print-3871.html [Accessed: 03/09/2004].

[21] ALEXANDER, Z. 2004. A closer look at the fine print in privacy statements. *InformIT* [Online], Jun. 11. Available from: http://www.informit.com/articles/printerfriendly.asp?p=174302 [Accessed: 08/06/2004].

[22] DATE, C.J. 2000. *An introduction to database systems*. 7th ed. Reading, Mass: Addison Wesley Longman.

[23] GARFINKEL, S. & SPAFFORD, G. 2001. *Web security, privacy & commerce*. 2nd ed. Cambridge, Mass.: O'Reilly.

[24] MILLER, M. 2002. Absolute PC security & privacy: defend your computer against outside intruders. San Francisco, Calif.: Sybex.

[25] BUECKER, A. & WATANABE, Y. 2003. *Design considerations for privacy-preserving database access* [Online]. Austin, Tex.: IBM:1-30. Available from: http://www.redbooks.ibm.com/redpapers/pdfs/redp3720.pdf [Accessed: 21/01/2004].

[26] NEVINS, S.C. 2003. Database security: protecting sensitive and critical information. *DM Direct Newsletter* [Online], *Jan. 3. Available from:*
http://www.dmreview.com/editorial/newsletter_article.cfm?nl=dmdirect&articleId=6310&issue=351 [Accessed: 7/8/2004].

[27] BOWERS, F. 2004. Dept admits privacy breach. *IrishHealth* [Online], Jun. 6. Available from: http://www.irishhealth.com/?level=4&id=5992. [Accessed: 27/10/2004].

[28] JAHNKE, A. 2003. Three things you don't want to know about your personal information. *Darwin*, [Online], July 9. Available from: http://www.darwinmag.com/connect/opinion/column.html?ArticleID=831 [Accessed: 27/10/2004].

[29] STOLEN data: yet another financial privacy breach. 2004. *SignOnSanDiego* [Online], Apr. 17. Available from: http://www.signonsandiego.com/uniontrib/20040417/news_lz1ed17middle.html [Accessed: 27/10/2004].

[30] SENATORS want to investigate privacy breach by jetblue airlines. 2004. *NewsTargetNetwork*, [Online], Oct. 26. Available from: http://www.NewsTarget.com/000251.html [Accessed: 27/10/2004].

[31] CLABURN, T. 2004. Report: people don't trust government to protect privacy. *InformationWeek* [Online], Feb. 13. Available from: http://www.informationweek.com/story/showArticle.jhtml?articleID=17700220. [Accessed: 24/02/2004].

[32] GROSZKRUGER, Z. *Patient privacy: practical tips to avoid being caught in a legal crossfire* [Online]. CAPPMT. Available from: http://www.cap-mpt.com/riskmanagement/patient-privacy.html. [Accessed: 19/07/2004].

[33] CAZIER, J.A., SHAO, B.B.M. & ST. LOUIS, R.D. 2003. Addressing e-business privacy concerns: the roles of trust and value compatibility. In: Proceedings of the 2003 ACM Symposium on Applied Accounting, held in Melbourne, Florida, USA. New York, N.Y.: ACM: 617-622.

[34] SINHA, A.K. 2003. *Mixing privacy and business intelligence*. *DMReview* [Online]. Available from: http://www.dmreview.com/editorial/dmreview/print_action.cfm?articleId=7268. [Accessed: 22/07/2004].

Citation information

H. J. G. Oberholzer and M. S. Olivier. “Privacy contracts incorporated in a privacy protection framework”. In: Computer Systems Science and Engineering 21.1 (2006), pp. 5–16