Using Organisational Safeguards to Make Justifiable Privacy Decisions when Processing Personal Data

MARTIN S. OLIVIER

Department of Computer Science, University of Pretoria

`http://mo.co.za`

Privacy-enhancing technologies can be used to enhance the privacy of individuals who interact with information processing systems. This paper considers such technologies that can be used by the organisation to safeguard personal information it processes. The paper focuses on how access control could be used to protect the individual against misuse of personal data inside the organisation. More specifically the paper considers how such a privacy-enhancing technology can make a just choice when deciding whether an access request to personal data should be allowed or not.

Access control decisions in this paper are based on the regulations that govern the interaction, the organisational policies that apply and the individual's privacy preferences.

The proposed model forms part of the organisational safeguards layer of the Layered Privacy Architecture (LaPA) proposed earlier.

1. INTRODUCTION

Ever since machines were first used to store information about individuals in large databases, concerns were voiced about the possible negative impact this held for the individual's privacy [Johnson 2001]. In Europe and the United states this led to the introduction of laws to limit harm to the individual from the power gained by the use of such technologies [Johnson 2001; Froomkin 1996; Langenderfer and Cook 2002].

As a modern democracy, South Africa's first step was more basic — and therefore possibly further reaching — than focussing on technology: the individual's right to privacy was guaranteed in the constitution [Republic of South Africa 1996, §2.14(d)] — albeit in forms that are not always easy to apply specifically to the computerised processing of individuals' data. This was supported further by the South African *Electronic Communications and Transactions Act* [Republic of South Africa 2002a] that states in article 51 that "A data controller must have the express written permission of the data subject for the collection, collation, processing or disclosure of any personal information on that data subject unless he or she is permitted or required to do so by law." In addition, the South African *Regulation of Interception of Communications and Provision of Communication-related Information Act* [Republic of South Africa 2002b] initially protects an individual (in article 2): "Subject to this Act, no person may intentionally intercept or attempt to intercept, or authorise or procure any other person to intercept or attempt to intercept, at any place in the Republic, any communication in the course of its occurrence or transmission."

However, once the individual has been given such rights, it soon becomes clear that the individual's rights need to be balanced by, amongst others, the interests of society [Olivier 2003a]. In practise this means that the individual's interests need to be balanced with those of her employer, businesses with whom he interacts and the state. To illustrate, consider an example in the field of short-term insurance. If privacy rights imply that an insurer cannot investigate and identify attempted or actual insurance fraud, fraud is likely to increase and the costs of such an increase can only be borne by those who are insured. An increase in costs is, in turn, likely to lead to more fraud. And so the vicious circle continues until insurance can no longer be afforded by anyone other than the very rich. On the other hand, if a party entrusted with sensitive personal data can simply declare that the individual has no right to expect the protection of his or her data, the costs of errors on the side of the insurer, are likely to be borne by individuals. As has been argued elsewhere [Chadwick et al. 2003], due to the vulnerability of the individual, if errors regarding privacy are to be made, the situation should be such that the error should be in the direction of too much privacy, rather than too little privacy.

To deal with such requirements of society, laws are used to limit the extent of privacy. Perhaps the most widely known current example is the US Patriot Act that was introduced after the incidents of 11 September 2001 in New York, that limits the individual's privacy in a bid to improve the country's ability to identify possible terrorists. This law has been widely defended and criticised [Chadwick et al. 2003; Swire and Steinfeld 2002].

Rather than discussing the international case further, this paper will endeavour to demonstrate the local relevance of the research described below. The South African *Electronic Communications and Transactions Act,* 2002, has already been used above as an example of an act that holds the individual's privacy in high regard. However, in article 50(4), the effect of breach of the privacy principles referred to above ('express written consent') "are governed by the terms of any agreement between" the organisation and the individual whose information is being processed. In the case of the *Regulation of Interception of Communications and Provision of Communication-related Information Act* quoted from above, it soon becomes clear that the intention of the act is not primarily to protect the privacy of the individual, but to enable interception of messages, given certain condition. For example, already in article 5 interception is permitted "if one of the parties to the communication has given prior consent in writing to such interception, unless such communication is intercepted by such person for purposes of committing an offence." And some communications can be intercepted (article 6(2)(d)) if the organisation "has made all reasonable efforts to inform in advance a person, who intends to use the telecommunication system concerned, that indirect communications transmitted by means thereof may be intercepted or if such indirect communication is intercepted with the express or implied consent of the person who uses that telecommunication system."

In many cases, personal privacy in Information Technology therefore becomes a matter for an agreement between the individual and those parties who process his or her personal information. Often the 'agreement' may be prescribed with the organisation with which one deals. This probably explains the clause at the bottom of the recently received schedule to the author's short term insurance contract , that essentially states that, in order to eliminate insurance fraud, the insured has *no* right to privacy [Multinet Makelaars 2003]. (The full text of the clause is reproduced in the appendix of the paper to show the exact extent in which the insured's right to privacy is limited.) While an authoritative legal opinion is required to determine the enforceability of the clause, it will clearly — at the very least — complicate the individual's task to protect his or her privacy if any of the sensitive details about the individual ever ends up in the wrong hands by actions of the insurer, whether such actions were deliberate, negligent or accidental. Note that the insurer has details about the personal assets of the individual, as well as a detailed list of measures taken by the individual to protect such assets. This clearly constitutes sensitive personal information, that the insurer should protect as such. In fact very little policy information is actually required to limit fraud, and only then under very specific conditions. A one-sided general denunciation of privacy rights is neither necessary nor acceptable.

In a comparison between a just war and business, Rossouw [2003] points out that business differs significantly from war because "business has the possibility of engaging with those who might be affected by the foreseeable negative side-effects of its actions." He continues, "when the opportunity exists of engaging with those who might be harmed by one's actions, it is morally preferable, if not imperative, to involve those affected by one's action in the process of moral deliberation."

If such deliberation is to occur when private data is processed, questions arise on how decisions should be reached, represented and enforced. Against this backdrop, the goal of the current paper is to consider the basis for making just decisions in such an environment. The question is approached by modelling the decision process mathematically to gain insight in the decision-making process on the fine-grained level of accessing individual data fields about some individual in order to process the data for some specified purpose.

The modelled solution is of a form that can be implemented as part of the organisational safeguards layer in the Layered Privacy Architecture (LaPA) that we proposed elsewhere [Olivier 2003b]. As such a layer it is intended to authorise (or audit) any processing of personal data; its operation should in turn be subjected to external audit to ensure that it is properly implemented by the organisation. Stated differently, suitable mechanisms and procedures should be in place to ensure that the trust placed by individuals in it is indeed warranted. However, due to space limitations, issues of trust are not considered further in the current paper.

The remainder of the paper is structured as follows. Section 2 provides more information about LaPA to explain where the proposed work fits in. Section 3 then gives more information about the organisational safeguards layer, develops the required model and uses the model to consider decision-making. Section 4 considers a number of issues that stem from our work and compares our work to other work that fits into the organisational safeguards layer. Section 5 concludes the paper.

2. BACKGROUND

The Layered Privacy Architecture (LaPA) [Olivier 2003b] positions the various technologies that may be used to protect (or enhance) the privacy of individuals when they interact with another party. The other party will typically be an organisation (but need not be); for this reason we will refer below to the *individual* and the *organisation* as the first two parties involved in the interaction. Note that our primary concern is interaction that stretches over longer periods of time — as is typically the case when an individual frequently shops at a specific shop or uses a specific airline. This is especially important if the

Figure 1.　The Layered Privacy Architecture (LaPA)

individual subscribes to some loyalty program of the organisation, but also applies when the individual is seemingly almost anonymous: in the real world, use of a credit card may link one transaction to the next; on the Internet, an IP address or cookie may be used for the same purpose. Additionally, the architecture applies when the interaction occurs only once and only for a brief time.

The architecture is depicted in Figure 1. It consists of four layers (1–4) and identifies six viable combinations of technology (a–f). It has been shown that the four layers are fully ordered in the sense that the lower layers inform the higher layers, while the higher layers control the lower layers [Olivier 2003b].

The remainder of this section briefly reviews existing privacy-enhancing technologies that fit into three of the layers of the architecture. It draws heavily on our review of the same technologies in another recent paper [Olivier 2003b]. For alternative reviews of privacy-enhancing technologies see the OECD inventory of such technologies [OECD 2002], the progress overview by Goldberg et al. [1997], an overview from a technical and legal perspective by Froomkin [1996] or the review paper by Seničar et al. [2003].

Private communication — layer 1 of LaPA — is inherently an aspect of the right to privacy and is explicitly enshrined in the South African Constitution [Republic of South Africa 1996, §2.14(d)]. Encryption is clearly one well-established technology to ensure privacy of communications, with steganography currently receiving some renewed interest. Note, in addition, that private communication extends beyond the channel: The old Hush-a-Phone [Stamper 1994, p.4] was a mechanical device that fitted over a telephone handset to enable the 'sender' not to be overheard when talking. Reweb-ber (previously Janus) [Rieke and Demuth 2001] is one technology that ensures that the user's surfing habits cannot be established from the logs that clearly falls outside the traditional communications channel.

Layer 3 of LaPA is the anonymity layer. Various schemes to ensure anonymity (or pseudonymity) have been proposed (see, for example, [Reiter and Rubin 1999; Gabber et al. 1999; Caloyannides 2000; Goldschlag et al. 1999; Rieke and Demuth 2001]). Most of these schemes are based on Chaum's so-called *mix* [Chaum 1981] — using public key encryption — or, alternatively, based on the notion of a proxy.

Personal control — LaPA's layer 4 — refers to the use of technology to ensure that an individual's personal information is only used in a manner commensurate with the individual's privacy policy. The goal is usually to compare the individual's privacy policy to that of the organisation the individual is dealing with, and only to release private information about the individual to the organisation if the two policies are compatible (or can be negotiated to a level of agreement). The best-known example in this category is P3P [Reagle and Cranor 1999].

Organisational safeguards — layer 2 of LaPA — refer to the use of technology to ensure that the organisation complies with its own privacy policy as well as the preferences of the individual. Since the contribution made by this paper fits into layer 2, work previously done that can also be categorised in this layer will be considered in section 4 below.

3.　THE ORGANISATIONAL SAFEGUARDS LAYER

The organisational safeguards layer contains those technologies that are used by the organisation to enhance the privacy of the individual whose information it has collected. To do this properly, this layer has to

(1)　Be informed about the individual's preferences;

(2)　Make justifiable decisions given the individual's preferences, applicable laws and regulations, and the organisation's goals;

(3) Be in a position to enforce its decisions (or, at least, be in a position to notice and flag actions that are counter to its own decisions, if it cannot proactively enforce them);

(4) Produce an audit trail that can be used by external parties to verify that the technology works as claimed; and

(5) Communicate the four preceding abilities to individuals in a reliable manner that allows them to use this communication to make decisions about their interaction with the organisation.

This paper only considers the second of these requirements — making justifiable access control decisions.

It is clear that decisions in the organisational safeguards layer need a representation of the three factors on which such decisions are to be based (preferences, policy and regulation). In actual fact, the situation is somewhat more complex. Firstly, organisational policies are not immune to change, and the version of the policy against which preferences were expressed by the individual, might not be the most recent policy. Moreover, when data is processed in an international context, multiple regulatory contexts might be applicable. Finally, given the fact that the database schema is unlikely to be forever static, the question needs to be posed to what extent an individual's preferences could be (validly) derived.

At first glance it seems that the appropriate solution to the first two problems is maintaining different versions of policies and regulations and linking each data record to the appropriate policy version and regulation. This, unfortunately, already causes a problem when a new data field is inserted in the schema: should the new data records now be dealt with according to the latest policy (even though the individual might not be aware of the existence of a new policy) or should the old policy be used (even though that policy might not cater for fields such as the new field)? This clearly demonstrates the requirement that policies should apply to data fields rather than entire data records.

On the other hand, associating each and every field with its own policy would clearly be inefficient. This is similar to an old problem in database security and is easily solved by considering the data record as a hierarchy [Rabitti et al. 1991]. Various policies can then be associated with different nodes in the hierarchy, and be inherited by nodes lower in the hierarchy, until inheritance is overridden by a node with which a different policy is associated.

A similar approach may be followed in the case of regulations: It is possible that new regulations only apply to data collected after a certain date. It is also possible that a regulation depends on the location where the user was when he or she supplied the data. As yet another example, a given regulation may only apply to certain categories of data. In all such cases, associating a policy with a node in the hierarchy addresses the immediate problem.

Viewing a record about an individual as a hierarchy also addresses the third problem we have alluded to: inferring a user's preferences when the schema changes. To illustrate this, consider a record where the individual can opt-in or opt-out of receiving various forms of communication from the organisation. If a new form of communication is added (and a corresponding field is added in each record, to record the individual's preferences regarding such communication), it might be possible to infer a preference for the field if the individual has already indicated a preference on a higher layer of the hierarchy: If, for example, the individual has already opted-out of receiving *any* information from the organisation, it clearly applies to the new form of communication (at least, until explicit instructions about the new form of communication are received from the individual). Does the same apply for the case where the individual has opted-in to receive communications from the organisation and a new form of communication is added? The answer depends on exactly what the individual has opted in to. If the individual opted-in to receive communications in general, the new form of communication is a special case of what has already been assented to and the preference can be inherited (until specifically modified); if the individual has, however, opted in to receive various other forms of communication, the preferences of the individual for such forms will be siblings of the preference of the individual for the new form and, as is standard practice, inheritance does not occur from siblings. The hierarchical structure of data therefore has definite implications for both the formulation of organisational policies and the expression of personal preferences. We do not discuss such implications in detail in the current paper, however.

Of more interest for the current discussion is the possible interaction between such various policies, preferences and regulations. To formalise the discussion let $D = \{Y, y, N, n, uc, c, s\}$ be the set of decisions that can be reached when a specific action is to be performed on an individual's data. Here Y indicates a strong *yes*, while y indicates a weak *yes*. The difference between the strong and weak positive decisions will be discussed below. Similarly, N and n, respectively, denote a strong and weak *no*. uc indicates that the *user's choice* should be honoured (or obtained and then honoured). c denotes that the specific policy, regulation or user preference explicitly allows a *choice*. Finally, s indicates that the policy, regulation or user policy is *silent* on the matter.

Next we need to consider the factors that are to be taken into account when an access request is submitted. We contend that the purpose for which the data is to be accessed is one of these factors. Let P be the set of purposes for which data can be accessed. This implies that a list of legitimate reasons are defined a priori. One may question whether it is possible to define such a set. We assume it is, based on the fact that ontologies are currently being developed in many subject areas. As one specific example, consider the ICD-10 classification system for medical diagnoses [World Health Organisation 2002]. If something as complex as human health problems can be described with a finite set, it seems likely that the same can be done for valid purposes for using personal information. In the current paper we do not explore the details of P.

Let F be the set of (data) fields that the organisation stores (that is, the schema). Suppose a request is submitted to access a field $f \in F$ with purpose $p \in P$. When the purpose and field are known, we assume that it is possible to determine the access mode from this. The access mode can be the traditional *read*, *write* and related modes; or it can be a method in an object-oriented system [Olivier and von Solms 1994]. We therefore do not consider access mode as one of the explicit factors to be specified when requesting access.

In addition to F and P, two other factors will determine whether access should be granted or not. Clearly, whether a request should be granted or not, depends on the subject who makes the request. Let S be the set of subjects known in the system. Finally, access depends on whose information is to be accessed. There are valid reasons why a given subject $s \in S$ should (not) be allowed to access information about some specific individual [Dreyer and Olivier 1999]. Let I be the set of individuals about whom information is maintained. Note that both S and I can be structured; subjects (S) are typically structured in terms of roles, but other schemes may also be considered. Individuals could be structured in various groups; a regulation or policy may, for example, make provision for *vulnerable groups* or *children*, where the latter is a subgroup of the former. However, since this paper does not consider the detailed operation of the functions that determine whether a subject should, according to regulation or policy, (not) be allowed to access information of an individual, we do not elaborate here on such structuring.

Now suppose a subject s requests permission to access individual i's data-field f for purpose p. Represent this as the function

$$q : S \times I \times F \times P \to D - \{s\}$$

. The outcomes should be interpreted as follows:

. Y, y: Allow the request. (See below for the difference between the weak and strong results.)

. N, n: Deny the request. (Again. see below for the difference between the weak and strong results.)

. uc: Get the individual's preference before proceeding. (See below how to deal with cases when it is not practical to (timeously) determine the individual's preference.)

. c: The organisation is free to decide whether the request should be allowed or not.

The access request function q will be defined by considering three functions q_{op}, q_r, q_{ip} that will, respectively, determine what the organisational policy, regulations and the individual preference say about the request. This will be useful to model inheritance below.

In order to formally define the operation of q_{op}, q_{ip}, q_r and, eventually, q, it is useful to introduce the notion of an ancestor pair.

Definition 1. If n is a node of a tree with a parent p, then the ancestor pair of n will be denoted by $\langle \pi, n \rangle$, where π is the ancestor pair of p. If n is the root node of a tree, its ancestor pair will be denoted by $\langle \epsilon, n \rangle$.

Consider the set of all data fields, F, maintained in the system. Let $\mathbb{F}$ be the set of all ancestor pairs defined over the hierarchy of data field nodes.

Next we define a 'helper' function, q'_{op}, to aid in the definition of the original q_{op}. This function determines the organisational policy that applies at a sequence of nodes working up from the node towards the root, until a policy is found for the node under consideration. If a given node is *silent* on the matter, the function is applied recursively up the tree until a policy is found.

$$q'_{op} : S \times I \times \mathbb{F} \cup \{\epsilon\} \times P \to D$$

Let $s \in S, i \in I, f \in F, p \in P$. Then

$$\begin{aligned}
q'_{op}(s, i, \langle a, f \rangle, p) &= q_{op}(s, i, f, p) \text{ if } q_{op}(s, i, f, p) \neq s \\
q'_{op}(s, i, \langle a, f \rangle, p) &= q'_{op}(s, a, f, p) \text{ if } q_{op}(s, i, f, p) = s \\
q'_{op}(s, i, \epsilon, p) &= s
\end{aligned}$$

where $< a, f > \in \mathbb{F}$.

The intention of the definition above is to traverse the tree from the node concerned towards the root of the hierarchy until a node is found that has a policy associated with it that makes a specific statement about the request. The first such node found will be the lowest node in the hierarchy that applies to the node concerned — and hence the most specific policy applicable to the node. Note that an s result differs significantly from c: The former means that no definite answer has been found and it is therefore necessary to traverse the tree further to find an answer; the latter explicitly allows a choice that implies the decision should be based on other issues (regulation and/or individual preference).

The operation of

$$q'_r : S \times I \times \mathbb{F} \cup \{\epsilon\} \times P \to D$$

is similar to q'_{op} above, and we omit the formal definition. Also

$$q'_{ip} : S \times I \times \mathbb{F} \cup \{\epsilon\} \times P \to D - \{uc\}$$

operates in a similar manner (but clearly a result of uc is not meaningful in this case).

In order to define q (the function that yields the final answer whether access should be allowed or not in light of the applicable regulations, policies and preferences) we introduce yet another helper function, d, below. q is then defined as follows, using d.

$$q(s, i, f, p) = d \left(\; q'_r(s, i, \langle a, f \rangle, p), \; q'_{op}(s, i, \langle a, f \rangle, p), \; q'_{ip}(s, i, \langle a, f \rangle, p) \; \right)$$

where $\langle a, f \rangle$ is the ancestor pair for f. Clearly, the three functions determine whether the requested operation should be allowed, based on regulation, policy and individual preference, respectively. It is then the purpose of d to combine the three answers into a single answer.

To consider $d : D \times D \times D - \{uc\} \to D - \{s\}$, it is necessary to consider the $|D| \times |D| \times |D - \{uc\}| = 7 \times 7 \times 6 = 294$ individual cases. Fortunately, the majority of the individual cases are relatively easy to deal with, leaving one with a handful of special cases. All 294 cases are represented in Figure 2. The principles used to compile the tables were (1) strong regulations (Y and N) should be complied with, and (2) the user's choice is seen as relatively important — especially if the regulations indicate that user choice should form the basis of the decision. A Y entry indicates that the operation should be permitted. A y entry indicates that the operation should probably be permitted, but an argument may be made in some cases not to allow it. These are typically cases where the organisation can choose not to allow it, but in such cases the exception should be noted in the privacy policy of the organisation. An N entry indicates that the operation should not be allowed. An n entry indicates that the operation should probably not be allowed, but again with the same provisos given for the y case above. A uc entry indicates that the individual's preference should be established before proceeding. A ? entry indicates that the case is far from clear, and should be referred to an appropriate party for a decision. (Note that such a party may sometimes also have to make the decision when the result of the calculation is uc, but a decision is required before the individual's preference can be established.) Dealing with such special cases are considered below. A c entry indicates that the organisation is free to proceed in whichever manner it prefers in the particular case.

As one of the special cases, consider $d(n, N, Y)$. Here regulation expresses a bias towards a negative response, policy expresses a definite negative response, but the individual has indicated a strong positive preference. In the majority of cases this is likely to indicate an incompatibility between the organisation's and the individual's goals. If that is indeed the case (and both parties have reconsidered their positions on the specific case) it probably indicates that ties between the individual and the organisation should be severed. The issue becomes more tricky if the ties cannot easily be severed, for example when the individual is forced by law to be connected with the organisation (as is the case with the national income tax authority) or if the organisation has a monopoly for the service(s) it offers. In such cases the long term solution for the issue is for the organisation to lobby to get regulations changed (to strong support) to override the individual's preference. At the same time the individual can try to garner support for his or her position and put pressure on the organisation to change its policy.[1] The individual can also attempt to get the regulatory authority to change the applicable regulation, but in this specific case this is unlikely since the authority has already expressed a negative bias. (Clearly lobbying and garnering support could also be considered when the organisation is only one of a series of providers of a service — but in such a case one at least has the option of choosing another service provider. Lobbying and garnering support should also be considered when all organisations in a given sector have the same policy in some respect.) The question, however, remains how should the issue be dealt with until the regulation and/or policy is changed or the ties between the individual and the organisation have been severed. Clearly, ignoring the individual's strong preference could render the individual vulnerable. On the other hand, going against the organisation's strong policy may prevent the organisation from protecting its legitimate interests. (It is possible that the individual has expressed a strong preference in opposition to the organisation's policy to protect his or her privacy; it is also possible that this strong opposition is a bid by the individual to mislead the organisation or even to commit fraud. Similarly, the organisational policy might be strong to protect its legitimate interests, or it might simply be caused by a profit motive. If the organisation and the individual continue to disagree on this point, neither of them would be able to arbitrate the matter.) The prudent approach for such cases is the appointment of a privacy ombudsman (by the appropriate regulatory authority) to whom such cases can be referred.

Other cases that pose problems similar to that discussed in the previous paragraph include $d(y, N, Y)$, $d(c, N, Y)$, $d(s, N, Y)$, $d(uc, N, y)$, $d(y, Y, N)$, and others.

---

[1]When expecting the individual to garner support for his or her privacy position, one should remember Jung's [2002, p.78] warning about collective means to support individual positions: "Anxiously we look round for collective measures, thereby reinforcing the very mass-mindedness we want to fight against. There is only one remedy for the levelling effect of all collective measures, and that is to emphasize and increase the value of the individual."
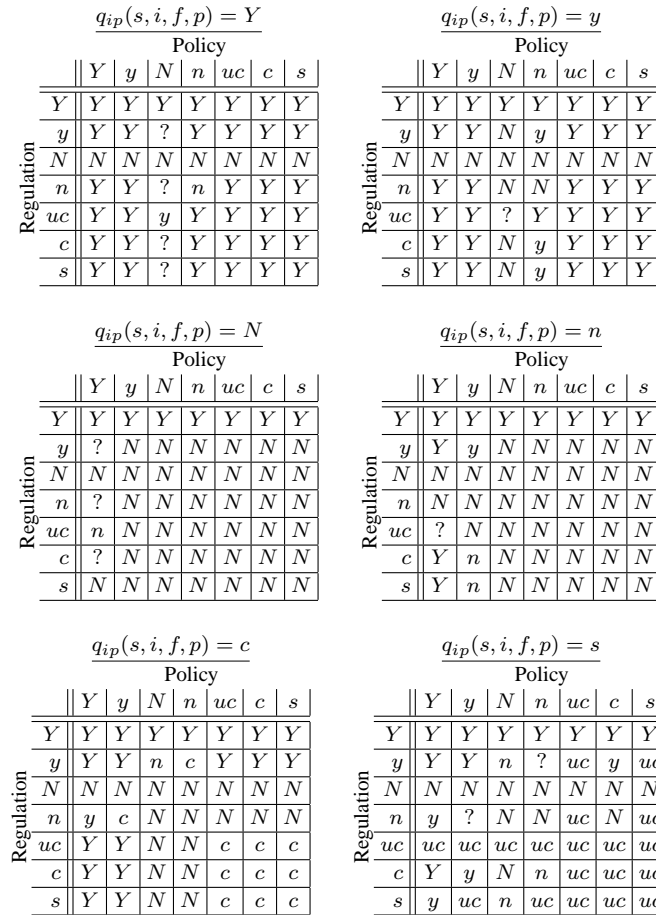
$$q_{ip}(s,i,f,p) = Y$$

Policy							
	Y	y	N	n	uc	c	s
Y	Y	Y	Y	Y	Y	Y	Y
y	Y	Y	?	Y	Y	Y	Y
N	N	N	N	N	N	N	N
n	Y	Y	?	n	Y	Y	Y
uc	Y	Y	y	Y	Y	Y	Y
c	Y	Y	?	Y	Y	Y	Y
s	Y	Y	?	Y	Y	Y	Y

(rows labelled Regulation)

$$q_{ip}(s,i,f,p) = y$$

Policy							
	Y	y	N	n	uc	c	s
Y	Y	Y	Y	Y	Y	Y	Y
y	Y	Y	N	y	Y	Y	Y
N	N	N	N	N	N	N	N
n	Y	Y	N	N	Y	Y	Y
uc	Y	Y	?	Y	Y	Y	Y
c	Y	Y	N	y	Y	Y	Y
s	Y	Y	N	y	Y	Y	Y

(rows labelled Regulation)

$$q_{ip}(s,i,f,p) = N$$

Policy							
	Y	y	N	n	uc	c	s
Y	Y	Y	Y	Y	Y	Y	Y
y	?	N	N	N	N	N	N
N	N	N	N	N	N	N	N
n	?	N	N	N	N	N	N
uc	n	N	N	N	N	N	N
c	?	N	N	N	N	N	N
s	N	N	N	N	N	N	N

(rows labelled Regulation)

$$q_{ip}(s,i,f,p) = n$$

Policy							
	Y	y	N	n	uc	c	s
Y	Y	Y	Y	Y	Y	Y	Y
y	Y	y	N	N	N	N	N
N	N	N	N	N	N	N	N
n	N	N	N	N	N	N	N
uc	?	N	N	N	N	N	N
c	Y	n	N	N	N	N	N
s	Y	n	N	N	N	N	N

(rows labelled Regulation)

$$q_{ip}(s,i,f,p) = c$$

Policy							
	Y	y	N	n	uc	c	s
Y	Y	Y	Y	Y	Y	Y	Y
y	Y	Y	n	c	Y	Y	Y
N	N	N	N	N	N	N	N
n	y	c	N	N	N	N	N
uc	Y	Y	N	N	c	c	c
c	Y	Y	N	N	c	c	c
s	Y	Y	N	N	c	c	c

(rows labelled Regulation)

$$q_{ip}(s,i,f,p) = s$$

Policy							
	Y	y	N	n	uc	c	s
Y	Y	Y	Y	Y	Y	Y	Y
y	Y	Y	n	?	uc	y	uc
N	N	N	N	N	N	N	N
n	y	?	N	N	uc	N	uc
uc	uc	uc	uc	uc	uc	uc	uc
c	Y	y	N	n	uc	uc	uc
s	y	uc	n	uc	uc	uc	uc

(rows labelled Regulation)

Figure 2. The values of d.

Referring cases where the individual has not yet expressed an opinion (for example $d(uc, Y, s)$) can, in some cases, also offer a solution. In the given example, regulation states that user choice should determine the decision. If it is impractical to obtain the individual's preference (such as when contact with the individual has been lost), arbitration by the ombudsman could be the only alternative.

In the case of $d(y, n, s)$ the situation is somewhat different. Here strong indicators are absent. One possibility is to accept the regulatory positive bias. Another possibility is to clearly state that the particular case is a policy exception — ie, whereas the policy is normally n, when the individual preference is s, the policy becomes N, leading to an answer without the aid of parties outside the organisation. A third possibility is referring the matter to the ombudsman. Clearly the best alternative is to establish the individual's preference.

While not ideal in all respects, the second alternative (using a policy exception to determine the decision within the organisation) provides a practical solution to another problem inherent in the proposed approach: the effort and cost (and, in general, the viability) associated with achieving the preference of the individual for each and every purpose for which a data field could be accessed could be prohibitively high. Using this alternative could be a cost-effective way of dealing with many day-to-day operations. The requirement that the case should be published as a policy exception is one aspect that justifies such an approach. If it is published, it should be communicated as part of the privacy negotiations when the personal control layer (such as P3P) is offered options. Further justification lies in the fact that, if a given situation is exploited by organisations, the regulatory authority has the option of changing the regulation to Y, N, or uc to force the organisation to deal with the matter differently. Note that policy exceptions could be highly specific: there are no reasons why an exception cannot be published for specific values (or sets) of s (subject), i (individual), f (data field) or p (purpose) — although exceptions for specific values of i should be questioned.

4. DISCUSSION

This section considers a number of aspects that need to be addressed before the paper can be concluded. Specifically, we need to consider possible alternatives to the d that was specified in Figure 2; we also need to consider the remarks in the previous section that policies and regulations are actually forests, rather than one permanent document each; finally the

work needs to be compared to other work that fits on the organisational safeguards layer.

To what extent is the definition of d in Figure 2 set in stone? While it was argued that in most cases the value of d is obvious, given the principles, it is very likely that another party may come to a different decision for a specific case. There are two ways to deal with such variations. The first is to accept the definition in Figure 2 as the baseline, and simply treat (and publish) specific differences as policy exceptions. Another alternative is to enumerate meaningful definitions of d, with the definition in Figure 2 referred to as, say, OSL1. If such names are assigned by the ombudsman, a simple name would uniquely define the scheme used by a particular organisation. Since not too many variations are likely, enumeration should not present serious problems.

In the previous section it was pointed out that policies and regulations are typically forests. The strategy described in the previous section explicitly traversed the data hierarchy until an applicable policy or regulation is found. The question is whether the policy (or regulation) hierarchy should be traversed before or after (or even simultaneous with) traversing the data hierarchy. We contend that, when traversing the data hierarchy and a node is found that is associated with a policy (or regulation) the policy (or regulation) hierarchy should be traversed at that point. If an appropriate policy (or regulation) is found, the traversal stops and the policy (or regulation) decision is determinable. If no appropriate match is found in the policy (or regulation) hierarchy, traversal of the data hierarchy proceeds. The process continues until a match is found during the traversals, or the root of the data hierarchy is reached. If this happens without a match, the decision is s (silent). The logic behind this order is that children nodes in policy (and regulation) hierarchies represent amendments to the parent nodes. If an amendment is associated with some data item, it is logical that the amended policy should also apply (with proper application of the amendments). If P' amends policy P, but individual i has previously agreed to policy Q', one either has to get i to agree to both P and P', or produce an amendment Q' that applies to i's policy, or produce a new policy R and get i to accept R, or to let i continue under Q. Therefore policy (and regulation) traversal is merely intended to find the appropriate item in a sequence of policy (or regulation) amendments. Hence traversal has to happen at the data node that identifies the policy (or regulation).

Finally, we have to compare the work of the previous section to other work that deals with organisational safeguards.

Agrawal et al. [2002] use the analogy of the Hippocratic oath to explain how databases should treat private data entrusted to them. A set of principles is identified that such a database should adhere to. These principles are then translated to a 'strawman' design. Similar to our approach, purpose for which data was collected plays a central role and is used to label all database fields, but no hierarchical ordering of fields (and hence no inheritance) is assumed. A validator is used to compare the organisation's policy and the individual's preferences. (Regulations are not explicitly considered.) Details of the decision process (the heart of the current paper) are not given, but it seems that individual preferences always take precedence. The paper raises many relevant questions that still need to be answered.

E-P3P [Karjoth et al. 2003; Ashley et al. 2003] uses P3P as its analogy, but focuses on using technology to backup the policy presented to the user. While an architecture to implement this functionality is presented, the major emphasis is the presentation of a comprehensive language that can be used to represent appropriate privacy policies, with a formal semantics for the language. Similar to our approach, policies are associated with data on a 'per-person and per-record basis' — something that Karjoth et al. [2003] refer to as a *sticky policy paradigm*. Policies are, however, not associated with nodes as such, but with data categories (such as *financial data* and where the data categories form a hierarchy). While using data categories is inherently different from using the instantiated structure of the data as we have done, both approaches have their advantages and disadvantages, which are not discussed further in this paper.

Apart from the different major goals, the most significant difference between E-P3P and our approach is the manner in which policies, regulations and preferences are combined to make decisions. In our approach the three factors are orthogonal. In E-P3P, regulations are considered, but play no direct role (as was the case for Hippocratic databases). Personal preferences in E-P3P only become relevant where the policy allows the individual explicit options; based on the individual's choice, the policy then allows or disallows an access request. While this approach simplifies matters, it suffers from two drawbacks:

(1) If policies are based on regulation and regulation changes, the policies need to be updated. Since this is orthogonal in our approach, just the policies need to be updated. Furthermore, our model caters for different policies to apply to different data elements, and for replacement of regulations on parts of the data hierarchy (which will be useful if a new regulation applies to data collected after a certain date, for example).

(2) If a user expresses a preference for an aspect that the organisational policy does not explicitly cater for, our approach will be able to deal with it, while E-P3P will require the organisation to ignore the individual's preference or to change its policy. (Note that our approach does not require the organisation to invite individual preferences for all data stored, but it does cater for cases where individuals on their own initiative express preferences.)

In the case of inconsistent rules in E-P3P, the request is rejected. (We contend that rejection is not always to the individual's benefit: Suppose that permission is sought from a user to sell his or her address to a mailing list; if the request to notify the user is denied due to inconsistent rules, but the sale proceeds, the user is denied making the final decision.)

An interesting facet of the presentation of E-P3P is the explicit distinction that is made between (security) access control and privacy (access) control.

Another interesting facet of E-P3P is support for obligations: accessing private data is not merely a matter of allowing or preventing access; accessing personal information may oblige the accessor to perform some action.

The organisational safeguards layer has also seen some initiatives on the business front. IBM's Tivoli Privacy Manager [Tivoli Software 2002] uses P3P as basis and associates P3P preferences with data at the collection point. Policies are then enforced when data is accessed, and an audit trail is created. PrivacyRight [PrivacyRight 2001] argues that letting the individual control the use of personal information in the organisation's database makes economic sense. Their TrustFilter permissions management system enables consumers to verify the integrity of their data and record permissions for the use of their data. The PrivacyWall family of products [IDcide 2001] enables organisations to monitor and audit compliance of their web practices with their privacy policies.

Note that trust seals (such as TRUSTe [Benassi 1999]) provide some of the listed requirements of the organisational safeguards layer listed earlier (see item 5 in section 3). However, we contend that privacy seals will be able to offer much more specific guarantees if the details of the organisational safeguards layer are fully developed and implemented.

5. CONCLUSION

This paper extended the LaPA privacy architecture [Olivier 2003b] by considering how just privacy decisions should be made on the organisational safeguards layer. This was done by modelling the appropriate aspects on that layer and by considering the various scenarios that can occur. It was shown that most cases are not complex to deal with, and strategies for dealing with complex cases were suggested.

As noted in section 3, decision-making is just one aspect of the requirements to build a comprehensive organisational safeguards layer. The remaining aspects identified in section 3 are left for further research.

Appendix

This appendix contains the full text of the insurance policy clause [Multinet Makelaars 2003] referred to in the introduction of this paper. The clause is quoted verbatim, as received, in Afrikaans, to avoid changing meaning when translating.[2] Note that the emphasis in the final statement does not occur in the original. Since that statement forms the crux of the matter, the English translation of that statement is given after the quotation for the benefit of those readers who do not understand Afrikaans.

Uitruil van Inligting

"Die Polishouer erken dat die uitruil van eise- en onderskrywingsinligting (met inbegrip van kredietinligting) deur Versekeraars noodsaaklik is om die versekeringsbedryf in staat te stel om polisse behoorlik te onderskryf en risiko regverdig te evalueer en om in openbare belang en met die doel om premies te beperk, die voorkoms van bedrieglike eise te verminder. Die Polishouer doen hiermee namens homself en namens enige persoon wat hy hierin verteenwoordig afstand van enige reg op privaatheid van enige eise-inligting wat deur hom of namens hom verskaf is ten opsigte van enige versekeringseis deur hom ingestel en verleen hiermee toestemming dat sodanige inligting aan enige ander versekeringsmaatskappy of sy agent geopenbaar mag word. Die Polishouer erken ook dat die inligting deur hom verskaf deur enige ander wettige bronne of databasisse gestaaf kan word. *Die Polishouer doen hiermee afstand van enige reg op privaatheid en verleen toestemming dat enige inligting van toepassing op enige versekeringspolis of -eis wat op hom betrekking het, geopenbaar mag word.*"

"The Policy Holder herewith waves any right to privacy and grants permission that any information applicable to any insurance policy or claim relating to him may be disclosed."

REFERENCES

AGRAWAL, R., KIERNAN, J., SRIKANT, R., AND XU, Y. 2002. Hippocratic databases. In *28th Int'l Conf. on Very Large Databases (VLDB), Hong Kong*.

ASHLEY, P., HADA, S., KARJOTH, G., AND SCHUNTER, M. 2003. E-P3P privacy policies and privacy authorization. In *Proceedings of the ACM workshop on Privacy in the Electronic Society*. ACM Press, 103–109.

BENASSI, P. 1999. TRUSTe: an online privacy seal program. *Communications of the ACM 42*, 2, 56–59.

CALOYANNIDES, M. A. 2000. Encryption wars: Shifting tactics. *IEEE Spectrum 37*, 5, 46–51.

CHADWICK, D., OLIVIER, M. S., SAMARATI, P., SHARPSTON, E., AND THURAISHINGHAM, B. 2003. Privacy and civil liberties: A panel discussion. In *Database and Application Security XVI*. Kluwer. In press.

CHAUM, D. 1981. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM 24*, 2, 84–88.

---

[2]The fact that the clause was communicated to the author in his home language is to be commended; too often such important information is communicated to individuals in a language they have not fully mastered.

DREYER, L. C. J. AND OLIVIER, M. S. 1999. An information-flow model for privacy (InfoPriv). In *Database Security XII: Status and Prospects*, S. Jajodia, Ed. Kluwer, 77–90.

FROOMKIN, A. M. 1996. Flood control on the information ocean: Living with anonymity, digital cash, and distributed databases. *University of Pittsburgh Journal of Law and Commerce 395*, 15. http://www.law.miami.edu/~froomkin/articles/oceanno.htm.

GABBER, E., GIBBONS, P. B., KRISTOL, D. M., MATIAS, Y., AND MAYER, A. 1999. Consistent, yet anonymous, web access with LPWA. *Communications of the ACM 42*, 2 (February), 42–47.

GOLDBERG, I., WAGNER, D., AND BREWER, E. A. 1997. Privacy-enhancing technologies for the Internet. In *IEEE COMPCON '97*. IEEE, 103–109.

GOLDSCHLAG, D. M., REED, M. G., AND SYVERSON, P. F. 1999. Onion routing. *Communications of the ACM 42*, 2 (February), 39–41.

IDCIDE. 2001. IDcide introduces corporate privacy compliance software. Press release. http://www.idcide.com/pages/press_releas.htm#6.

JOHNSON, D. G. 2001. *Computer Ethics*, Third ed. Prentice Hall.

JUNG, C. G. 2002. *Flying Saucers — A modern myth of things seen in the sky*. Routledge.

KARJOTH, G., SCHUNTER, M., AND WAIDNER, M. 2003. Platform for Enterprise Privacy Practices: Privacy-enabled management of customer data. In *Privacy Enhancing Technologies: Second International Workshop, PET 2002, San Francisco, CA, USA, April 14-15, 2002, Revised Papers*, R. Dingledine and P. Syverson, Eds. Springer.

LANGENDERFER, J. AND COOK, D. L. 2002. Oh, what a tangled web we weave — the state of privacy protection in the information economy and recommendations for governance. *Journal of Business Research 5749*, 1–14.

MULTINET MAKELAARS. 2003. Versekeringskedule. Persoonlike korrespondensie.

OECD. 2002. Inventory of privacy-enhancing technologies (PETs). Report DSTI/ICCP/REG(2001)1/FINAL, Working Party on Information Security and Privacy, Organisation for Economic Co-operation and Development.

OLIVIER, M. S. 2003a. Database privacy. *SIGKDD Explorations 4*, 2, 20–27.

OLIVIER, M. S. 2003b. A layered architecture for privacy-enhancing technologies. In *Proceedings of the Third Annual Information Security South Africa Conference (ISSA2003)*, J. H. P. Eloff, H. S. Venter, L. Labuschagne, and M. M. Eloff, Eds. Sandton, South Africa, 113–126.

OLIVIER, M. S. AND VON SOLMS, S. H. 1994. A taxonomy for secure object-oriented databases. *ACM Transactions on Database Systems 19*, 1, 3–46.

PRIVACYRIGHT. 2001. Control of personal information — the economic benefits of adopting an enterprise-wide permissions management platform. White Paper. http://www.privacyright.com/info/economic.html.

RABITTI, F., BERTINO, E., KIM, W., AND WOELK, D. 1991. A model of authorization for next-generation database systems. *ACM Transactions on Database Systems 16*, 1, 88–131.

REAGLE, J. AND CRANOR, L. F. 1999. The platform for privacy preferences. *Communications of the ACM 42*, 2 (February), 48–55.

REITER, M. K. AND RUBIN, A. D. 1999. Anonymous web transactions with Crowds. *Communications of the ACM 42*, 2 (February), 32–48.

Republic of South Africa 1996. Constitution of the Republic of South Africa. Act 108 of 1996.

Republic of South Africa 2002a. Electronic communications and transactions act. Act 25 of 2002.

Republic of South Africa 2002b. Regulation of interception of communications and provision of communication-related information act. Act 70 of 2002.

RIEKE, A. AND DEMUTH, T. 2001. JANUS: Server anonymity in the world-wide web. In *Conference Proceedings EICAR International Conference*, U. E. Gattiker, Ed. 195–208.

ROSSOUW, G. J. 2003. Business is not just war — transferring the principle of double effect from war to business. Colloquium lecture, Department of Philosophy, RAU University. May.

SENIČAR, V., JERMAN-BLAŽIČ, B., AND KLOBUČAR, T. 2003. Privacy-enhancing technologies — approaches and development. *Computer Standards & Interfaces 25*, 147–158.

STAMPER, D. A. 1994. *Business Data Communications*, Fourth ed. Benjamin/Cummings.

SWIRE, P. AND STEINFELD, L. 2002. Security and privacy after September 11: the health care example. In *Proceedings of the 12th Annual Conference on Computers, Freedom and Privacy*. ACM Press, 1–13.

TIVOLI SOFTWARE. 2002. Enable your applications for privacy with IBM Tivoli Privacy Manager for e-business. Technical discussion, IBM. July.

WORLD HEALTH ORGANISATION. 2002. The international statistical classification of diseases and related health problems, tenth revision (web page). http://www.who.int/whosis/icd10/.