

A Layered Architecture for Privacy-enhancing Technologies

Abstract

While a number of privacy-enhancing technologies have been proposed over the past quarter century, very little has been done to generalise the notion. Privacy-enhancing technologies have typically been discussed for specific applications (such as confidential and/or anonymous e-mail) or in specific contexts (such as on the Internet). This paper takes cognisance of existing privacy-enhancing technologies, abstracts from them to a more general environment, and structures the technologies in a general architecture, based on the relationships between the technologies.

The resulting architecture consists of four layers, viz the personal communications, anonymising, organisational safeguards and personal control layers. It is also argued that a strong ordering exists between the layers — in the order just given.

The proposed architecture suggests an approach to constructing integrated, comprehensive privacy solutions.

1 Introduction

The notion of using technology to enhance privacy is not new. At the time of writing this, a search on Google for the phrase “privacy-enhancing technologies” reported about 7200 hits. Despite all this activity on privacy-enhancing technologies (PETs), very little has been done to structure the various attempts to enhance privacy-enhancing technologies in a manner that (1) positions a specific technology in the context of the privacy problem it addresses and (2) allows one to see how various such technologies can be combined to address the privacy problem. In fact, much work on privacy-enhancing technologies present a specific privacy-enhancing technology as *the* solution to the privacy problem.

In order to discuss privacy-enhancing technologies, the privacy problem itself should be framed in a particular technological context. Clearly, masks, costumes and cosmetics that hide the identity of an individual in the real world may be construed as privacy-enhancing technologies. Similarly, devices to camouflage one’s voice when using the telephone, technology to distort a victim’s face beyond recognition when reporting about the victim on television, and even aspects of cosmetic surgery are forms of privacy-enhancing technologies. However, such technologies are not what this paper concerns itself with.

The technological environment that is of interest here, is the global IT infrastructure. We will henceforth refer to privacy problems introduced (or exacerbated) by this infrastructure as *the Privacy Problem*. Against this backdrop, one can ask the question: Which privacy-enhancing technologies exist (or can be created) to address the Privacy Problem? And the specific question considered by this paper is then: What are the relationships between such technologies?

Given our formulation of *the Privacy Problem*, it is still necessary to consider the notion of *privacy problems* on which our formulation is based. For our discussion we will assume that any act (or failure to act) based on information about an individual that renders the individual more vulnerable than prior to the act, constitutes *a* privacy problem. Note that *vulnerability* here is used in a sense that implies that the individual is not justly exposed: exposure of a crime committed by an individual does not render the individual *vulnerable*. We do not consider the notion of *just exposure* further in the current paper, but note that such exposure should be based on acceptable notions of justness, and be subject to appropriate legal and societal sanctions.

The explicit consideration of a technological environment may well explain the more limited focus of other work on privacy-enhancing technologies. If the concern is government's ability to collect information about individuals (ie, the concern is 'Big Brother') the solution is legislation, as embodied by the US Privacy Act of 1974 [14, p.114]. If the concern is the collection of click-stream data on the Internet, then anonymity is one of the major solutions [8]. Clearly, in the wider context all such solutions still have a role to play. The emphasis in this paper is on the relationships between such (technical) solutions.

The question of the relationships between various privacy-enhancing technologies is addressed in this paper by layering the various identified categories of privacy-enhancing technologies. This leads to a layered privacy architecture that structures the categories such that the viable combinations of technologies that can be used (from the individual's perspective) are identified. This enables us to discuss the combination — and hence the relationships — in a structured manner.

The paper is structured as follows. Section 2 briefly reviews privacy-enhancing technologies that have been proposed elsewhere. Section 3 develops the proposed individual privacy architecture. Section 4 concludes the paper.

2 Background

Various forms of privacy-enhancing technology have received research attention. In [17] it has been suggested that five such technologies have, to a greater or lesser extent, emerged, namely technologies to facilitate private communication, anonymity, personal control, organisational safeguards and inference control. This paper focusses on the first four; later work will revisit inference control to decide whether it should be added to the architecture or subsumed in one of the other categories. The four major categories are introduced briefly below.

Private communication is inherently an aspect of the right to privacy and is explicitly enshrined in the South African Constitution [21, §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 [23, p.4] was a mechanical device that fitted over a telephone handset to enable the ‘sender’ not to be overheard when talking.¹ Rewebber (previously Janus) [22] 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. Despite this remark, we will below often refer to the *channel* as if it included such aspects.

Various schemes to ensure anonymity (or pseudonymity) have been proposed (see, for example, [20, 10, 6, 11, 22]). Most of these schemes are based on Chaum’s so-called *mix* [7] — using public key encryption — or, alternatively, based on the notion of a proxy.

Personal control 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 [19].

Organisational safeguards 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. Keeping track of a user’s wishes to opt-in or opt-out of receiving unsolicited e-mail is one simple example. However, this category includes significantly safer technologies that are installed to double check that the organisation’s compliance with such policies and preferences. While some products have appeared in this category [18, 13, 24], academic research the topic is beginning to emerge — see E-P3P [15, 5], the notion of a Hippocratic Database [1] and making just decisions on this layer [2] for three examples.

Note that our proposed classification of privacy-enhancing technologies is not the first attempt to classify such technologies: elsewhere [16] privacy-enhancing technologies have been classified into the following categories: personal privacy-enhancing technologies, web-based technologies, information brokers and network-based technologies. Based on the source used here [16], we will refer to this classification as the OECD classification.

Below we repeat the OECD classification, but now include the examples of privacy-enhancing technologies they [16] list of each category, followed by our own classification of the particular technology in brackets. A lack of space here prevents a detailed motivation of each of our classifications.

- Personal privacy-enhancing technologies: Cookie managers or blockers (private communications), Ad blockers (personal choice), Encryption software (private communications)
- Web-based technologies: Anonymisers (anonymisers), Platform for Privacy Preferences Project (personal control), Privacy networks (anonymisers and personal control)
- Information brokers: Infomediaries (anonymisers and personal control)
- Network-based technologies: Proxies (anonymisers) and firewalls (see below), Privacy networks (anonymisers and personal control)

¹One cannot help but wonder whether re-introduction of such a device in the age of cellular phones would be beneficial — not for the privacy of the talker, but for the benefit of those in his or her vicinity.

In the case of firewalls, we contend that they are an auxiliary technology, aiding in protecting the integrity of the private communications channel, rather than being a proper privacy-enhancing technology itself. The possible exception to this statement is a personal firewall, that prevents rogue software on the user's machine to communicate out of band with other parties. In our case a personal firewall would be classified as a private communications technology. Given the OECD classification of firewalls as a network-based technology, it is unlikely that they specifically had personal firewalls in mind.

While the previous paragraph demonstrated that the two classification schemes are entirely different, we argue that that our approach holds (at least) three advantages over the OECD scheme:

1. The OECD classification does not make provision for our category of organisational controls (in the sense that it is not clear how any of the examples in our category should be treated in the OECD classification;
2. The OECD scheme is much closer tied to current technologies (such as the web) than our scheme; and
3. Our scheme will address the relationships between the categories and result in a coherent architecture.

Note that the fact that some of the OECD examples (such as infomediaries) extend over two of our categories does not detract from our approach: some solutions inherently use composite technologies and our approach helps one to identify such cases. (As an aside, note that privacy networks are also listed as examples of two categories in the OECD case.)

The IBM Enterprise Privacy Architecture [12] focuses on business processes and comprises a management reference model, a technical reference model and a privacy workflow framework. The technical reference model includes a technical architecture that seems to share some goals with the architecture proposed in this paper. Details about the technical architecture do not seem to be readily available; specifically a search of the major peer-reviewed literature returned no information, making a proper comparison with our approach impossible at this time.

3 The Privacy Architecture

3.1 Privacy of Communications

Given the four categories of privacy-enhancing technologies identified in the previous section, we contend that private communications is a fundamental category. Its fundamentality lies in the fact that its absence weakens almost all solutions in other categories. We consider this claim for each of the remaining three categories below.

In the case of organisational controls, it is clear that it will be of little use to the user who trusts organisational controls (to some extent) to enhance privacy, if the individual cannot assume that the communication between the organisation and the user is private. If this were not the case, an eavesdropper who does not fall under the organisational controls may be able to intercept

communications between the individual and the organisation (or even between this organisation and another about the individual, where this is allowed by the agreed upon privacy policy). Such interception clearly renders organisational controls useless in such cases.

In the case of personal control, the matter is somewhat less clear cut. When the individual decides to withhold information from another party, the presence or absence of a private communications channel is immaterial. However, the individual's choice is likely to depend on some representation made by the other party, such as a privacy policy. If such a privacy policy can be fabricated on the communications channel between the two parties, the user may well decide to disclose information based on the fabricated policy. While it is therefore true that the individual may use the availability of a private communications channel as the basis for exercising personal control, absence of such a channel will, in general, greatly restrict the options available to the individual wishing to exercise personal control.

In the case of anonymity, the fundamentality of private communication may be demonstrated by considering the reasons why someone would prefer to work anonymously. Often the reason will be that the individual prefers not to be associated with his or her actions. This may be the case when the individual votes in an election or when the individual wants to download, say, adult content from some server. In such cases it will clearly be of little use if the individual is anonymous to the server, but someone who can identify the individual can eavesdrop on the conversation and determine the individual's vote or the nature of content downloaded by the individual.

In all of these cases it is possible to identify situations where a private communication channel is apparently not required. One may argue, for example, that if one (anonymously) downloads an encrypted 'parcel' from a nondescript site, that a private communication channel is not required to make anonymity effective. However, in this case it is clear that the 'encrypted parcel' renders the channel (sufficiently) private. As a second example, one may argue that an individual who connects via a dial-up link directly to an organisation, and uses personal control based (partly) on organisational controls offered by the organisation, may consider the threat of eavesdropping minimal, and decide to proceed despite the absence of an inherently private channel. However, in this case again, the technology (the dial-up connection) provides a (sufficiently) private channel. As a third example, consider the individual who prefers to transact anonymously on a network — not because he or she wants to hide the existence of such transactions, but because he or she simply wants to prevent the organisation from constructing a profile. If this individual typically uses the network from a relatively large number of locations, it may indeed be difficult for an eavesdropper to construct a profile and the only real threat may be the organisation concerned. Hence even in this case there is an assumption of a private communications channel — provided by changing the point of access. We therefore contend that a private communication channel is assumed in practice, even though (1) it may not be a perfectly private channel and (2) the channel may indeed not be physically present.

In our layered architecture (see figure 1, a private communication channel is depicted as a layer that covers the breadth of solutions. Note that no ordering relation is implied by the positioning of the layers at this stage.

Note that the private communications layer in the architecture does not imply that all (or even many) privacy solutions should share the same private communications channel. (In a layered




Figure 1: Working version of the privacy architecture

network protocol, a protocol on a lower layer is typically shared by a significant proportion of higher layer protocols.) If a privacy solution does not use a private communications channel used by other privacy solutions, it still fits the proposed architecture as long as it addresses the problem of eavesdropping.

3.2 Personal Control

Having addressed the private communications layer, we now turn our attention to the personal control layer.

3.3 The Personal Control Layer

The personal control layer may be likened to two parties contracting. Both make certain representations and then, based on those representations, they continue to transact. Consider a concrete example: An organisation may warrant that it will not disclose the e-mail address to be supplied by the individual to third parties and the individual may warrant that he or she is in fact the owner of the e-mail to be supplied by him or her. If the two parties find these terms acceptable, the individual may decide to proceed and indeed supply the other party with his or her e-mail address.

The task at hand in this section is to consider the relationships between the personal control layer and other layers. Since its relationship with the private communications layer has already been considered, the relationships with the anonymity and organisational safeguards layers need to be considered.

The personal control and anonymity layers may be, but do not have to be used in combination. In some cases the individual may decide to interact with another party without any desire to hide

her or his identity. To demonstrate that personal control and anonymity can indeed be used in combination in a meaningful manner, is more complex. However, consider Froomkin's [9] observation that anonymity does not imply unlinkability. Chadwick [3] notes the extent (by using wigs, make-up and gloves) one has to go to to achieve true anonymity in the real world. In cyberspace one may, based on what is known (or believed) about the other party, decide to go to great lengths — beyond 'passive' anonymity (or even 'passive' pseudonymity) — to create such an alternative persona with the aim of presenting an obfuscating image to the other party. In such cases, the individual is likely to obtain information to serve as the basis of his or her decisions (as is the case in P3P), but obtain such information from third parties. Also, personal control here is not about deciding what information to entrust the other party with, but deciding how much obfuscating information (and behaviour) to present with the created persona to achieve one's privacy goals.

Presentation of an entirely new persona is not the only viable combination of personal control and anonymity. One may, for example, decide to use a pseudonym at an online shop, but use one's real postal address for deliveries. To illustrate this, consider someone who happens to be called Bill Gates. It is quite likely that this may draw unwarranted attention to his orders — especially when he orders software for his Apple computer. He may even frequently get nuisance e-mails enquiring whether he is the 'real' Bill Gates (as you, the reader, probably assumed when I introduced this example). Therefore this Bill Gates may decide to use a pseudonym for shopping, but continue to use his real nondescript postal address. However, in some cases our Bill Gates needs invoices in his real name to claim tax deductions. He is therefore faced with a choice each time to use his real identity, or his minimally modified 'persona'. Anonymity tools can aid in management of personas. Personal control tools can, based on information supplied by the second and/or third parties, as well as on personal issues, be used to decide what information to supply to the other party; in this context it can provide input as to which persona should be used. Note that it is also possible to use a chain of suppliers, where each has only some real information to process its part of the transaction [4]. We may refer to this as a chain of personas, each containing very little real information, that are anonymous when seen in isolation, but, when combined, yields the real identity. Privacy protection lies in the (hopefully) hard task of combining personas.

This leaves us with the question whether personal control tools, when used in combination with anonymity tools, depend on the availability of operational safeguards. The answer is yes and no: yes, because one may still entrust the other party with sensitive details such as an e-mail address (as highlighted in the previous example); no, because the individual may be presenting a persona that is so far from the truth that safekeeping of information supplied makes no difference.

The preceding discussion is reflected in figure 1 by positioning of the personal control layer such that it can be used without organisational safeguard or anonymity tools (a), only with organisational safeguard tools (b), only anonymity tools (d) or with both organisational safeguard and anonymity tools (c).

3.4 The Anonymity and Organisational Safeguards Layers

The relationships between the various layers have all been considered by now, with the exception of the relationship between the anonymity and organisational safeguards layers. This section considers that relationship.

As noted earlier, we will use the term *anonymity* to include notions of pseudonymity because the ultimate goal is to hide (aspects of) the individual's identity. As already noted in the previous section, management of anonymity may imply management of personas. As in the previous case, the persona may or may not communicate information that renders the individual vulnerable. It is also clear that anonymity can be used without even knowing whether the party with whom one is communicating has organisational controls in place. So it seems as if it is viable to use the two layers separately or in combination.

However, it is important to remember that we are considering an architecture for technical privacy solutions. Note that anonymity providers and organisational safeguards — almost by definition — have to execute on different platforms (since the anonymity provider has to anonymise the inputs it gets from the organisation, which is executing the safeguards. Given this fact, the question arises whether they, as part as a technical solution can interact in any meaningful way (in the absence of a personal control layer). The only viable possibility seems to be the case where the anonymiser queries the organisation about the existence of such controls, and adapts its actions based on the response it receives. However, it seems dangerous to let an anonymiser — without knowledge about the individuals preference's and motives as was the case in the personal control layer — make decisions about what about an individual should (not) be released to another party, only based on the presence or absence of organisational controls. We therefore omit this possibility from the proposed architecture.

Clearly an anonymising layer can be used (and is, in practice rather appropriate) if no organisational safeguard layer is present; this alternative has been labelled (e) in figure 1. The remaining question is whether an organisational safeguard layer can exist if neither the anonymity, nor the personal control layer is present. The answer is clearly positive: In many cases (such as tax returns) the individual does not have personal control and anonymity is not an option; the individual will, however, benefit from organisational controls at the processing party. In figure 1, this option has been labelled (f).

3.5 On the order of layers

Figure 1 was not intended to suggest an order of layers. This section considers whether such an order can be substantiated. For the purposes of this section, let o , a , p and c represent the organisational safeguards layer, anonymising layer, personal control layer and personal communications (channel) layer, respectively. Let $x > y$, with x and y two of these layers, mean x controls y . Layer x can control layer y by providing configuration parameters for y , or by choosing one of a number of available alternative solutions on layer y . Often this means that layer x has to be informed about the permitted ranges and available alternatives on layer y . However, the specific choice within the range, or the specific choice of an alternative, or even the option not to proceed, lies with x .

Consider the relationship between c and the other three layers. We contend that $o, p, a > c$. This follows from the observation all that other layers in practice prescribe the use of one or more communication channels:

- p : Personal control, by definition, includes ensuring that released data is not intercepted or modified by other parties. To do this properly, this implies choosing and/or configuring the channel. Hence $p > c$.
- a : Almost all anonymity solutions prescribe encryption techniques to be used on the various channels. From such prescription it follows that $a > c$.
- o : Organisational safeguards will typically ensure that such safeguards are not bypassed when communicating the safeguarded information internally. For this reason the organisational safeguards layer has to ensure that it commands the appropriate private communication channels internally, as well as when communicating the information with any other (authorised) party. (In particular, this has to be the case when it communicates with the individual.) Hence $o > c$.

From the definition of $>$ it is not clear that $>$ is not commutative, since it is possible that two layers can mutually control one another. In order for $>$ to be an ordering relation, we have to show that it is indeed not commutative. Since we now know that $a, p, o > c$, we next have to demonstrate that $c \not> x$, with $x \in \{a, p, o\}$. This, however, follows directly from the observation that none of these three layers (x) should be controlled by the private communications channel. All three warrant aspects of privacy to the individual (and to other layers). To keep such promises, they cannot be subjected to the control of the channel; at most they can be informed about the availability and/or capabilities of channels and make the decision on how the channel(s) should or should not be used.

Given $a, p, o > c$, we next claim that $a, p > o > c$. To substantiate this claim we have to show that $a, p > o$. In order to demonstrate the non-commutativity, we will also show that $o \not> a, p$. For the latter case, consider the two claims:

- $o \not> a$: Since the function of the anonymising layer is to hide information from the organisation, it is clear that the organisational safeguards layer should not be able to control the anonymising layer. It should be clear that the organisational safeguards layer may provide tolerance data to the anonymising layer: If the anonymising layer will create a persona for the individual, the anonymising layer needs to know, for example, what the restrictions are that the organisational safeguards layer places on user identifiers and passwords (such as lengths, case and other requirements). However, again it is up to the anonymising layer to choose the user identifiers within those tolerances. Hence $o \not> a$.
- $o \not> p$: It is quite clear that the organisational safeguards layer should not be able to control the personal control layer. Here the organisational safeguards layer should inform the personal control layer what its options are (such as opting in or out of receiving a regular newsletter by e-mail). It should also inform the personal control layer about the warranties it makes about processing of data. This then allows the personal control layer to make its decisions

about choosing various options and/or trusting the organisational safeguards layer with (specific) personal information. Hence $o \not> p$.

Now consider the two positive cases:

$a > o$: While it is possible to contrive examples to demonstrate that the anonymising layer does indeed control the organisational safeguards layer in some cases, we have been unable thus far to demonstrate this with realistic examples. We will therefore use the tenuous argument that, if a control relation exists between a and o , since we have established above that $o \not> a$, the relationship has to be $a > o$.

$p > o$: The organisational safeguards layer is the guardian of the individual's data and has to safeguard that data — at least in part — according to the user's wishes, as expressed in the personal control layer. Hence the personal control layer controls aspects of the organisational safeguards layer, and we have $p > o$.

This (sufficiently) confirms the claim that $a, p > o > c$.

Next we claim that $p > a > o > c$. The following two facts will support this:

$p > a$: Clearly, personal control also implies a choice over what information should be anonymised, whether a link should be maintained between the user's real identity and any created pseudonyms or aliases, etc. This constitutes control of the anonymising layer by the personal control layer by allowing the latter to configure the former and/or allows the latter to choose amongst available alternatives of the former. Hence $p > a$.

$a \not> p$: Again, the anonymising layer can inform the personal control layer of choices and alternative the latter has. If the anonymising layer were to control the personal control layer, it would remove the control function, inherent in the latter's name. Hence $a \not> p$.

This section demonstrated that $p > a > o > c$. The non-commutative nature of $>$ has also been demonstrated. Finally, by consideration of all combinations of p , a , o and c , the transitive nature of $>$ has been demonstrated. Hence the layers are fully ordered.

Figure 2 summarises the argument used in this section. The upper lefthand portion of the table illustrates the manners in which the upper layers of the layered architecture controls the lower layers. The bottom righthand triangle illustrates the way in which the lower layers of the architecture informs the upper layers about alternatives and options. Note that the lower layers need not inform the upper layers directly; other sources could also inform the upper layers about the capabilities of the lower layers.

Given the fact that a definite order has now been established for the architecture, figure 3 represents the layered version of the proposed privacy architecture. This answers the question posed at the beginning of this paper.

4 Conclusion

This paper considered the possibility of structuring privacy-enhancing technologies in a meaningful manner. A four layer privacy architecture was proposed. Interaction between layers was

		Target			
		Private communications	Organisational safeguards	Anonymity	Personal control
Source	Personal control	Choose, Configure	Specify personal preferences	Choose persona, Specify retention	N/A
	Anonymity	Choose, Configure	?	N/A	Alternatives, Configuration options
	Organisational safeguards	Choose, Configure	N/A	Tolerances	Tradeoffs, Warranties
	Private communications	N/A	Alternatives, Configuration options	Alternatives, Configuration options	Alternatives, Configuration options

Figure 2: Examples of how the upper layers of the architecture controls the lower layers and how the lower layers informs the upper layers




Figure 3: Layered version of the privacy architecture

considered and this resulted in the identification of six combinations of the considered technologies that can be used to provide a privacy solution with the characteristics required by the situation in which it is needed.

Furthermore, by considering relationships as either informing or controlling, a fully ordered relationship was established between the four layers.

This resulted in the proposal of a layered privacy architecture. For ease of reference, we will refer in subsequent work to this architecture as LaPA (Layered Privacy Architecture).

Since the nature of interaction between layers has now been established, it becomes possible to consider interoperation between layers — given a specific combination of layers — in detail. This may result in protocols that formalise interoperation between technologies that are currently used in isolation, which in turn will hopefully lead to the possibility to use such technologies in an integrated manner in practice. This, however, is left for future work.

References

- [1] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu. Hippocratic databases. In *28th Int’l Conf. on Very Large Databases (VLDB)*, Hong Kong, 2002.
- [2] Removed for blind refereeing.
- [3] Removed for blind refereeing.
- [4] Removed for blind refereeing.
- [5] P. Ashley, S. Hada, G. Karjoth, and M. Schunter. E-p3p privacy policies and privacy authorization. In *Proceeding of the ACM workshop on Privacy in the Electronic Society*, pages 103–109. ACM Press, 2003.

- [6] M. A. Caloyannides. Encryption wars: Shifting tactics. *IEEE Spectrum*, 37(5):46–51, 2000.
- [7] D. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of the ACM*, 24(2):84–88, 1981.
- [8] L. F. Cranor. Internet privacy. *Communications of the ACM*, 42(2):28–38, 1999.
- [9] A. M. Froomkin. Flood control on the information ocean: Living with anonymity, digital cash, and distributed databases. *University of Pittsburgh Journal of Law and Commerce*, 395(15), 1996.
<http://www.law.miami.edu/~froomkin/articles/oceanno.htm>.
- [10] E. Gabber, P. B. Gibbons, D. M. Kristol, Y. Matias, and A. Mayer. Consistent, yet anonymous, web access with LPWA. *Communications of the ACM*, 42(2):42–47, February 1999.
- [11] D. Goldschlag, M. Reed, and P. Syverson. Onion routing. *Communications of the ACM*, 42(2):39–41, February 1999.
- [12] IBM. Enterprise privacy architecture: Securing returns on e-business. Executive brief, IBM, 2001.
- [13] IDcide. Idcide introduces corporate privacy compliance software. Press release, February 2001.
http://www.idcide.com/pages/press_releas.htm#6.
- [14] D. G. Johnson. *Computer Ethics*. Prentice Hall, 3 edition, 2001.
- [15] G. Karjoth, M. Schunter, and M. Waidner. Platform for Enterprise Privacy Practices: Privacy-enabled management of customer data. In R. Dingledine and P. Syverson, editors, *Privacy Enhancing Technologies: Second International Workshop, PET 2002, San Francisco, CA, USA, April 14-15, 2002, Revised Papers*. Springer, 2003.
- [16] OECD. 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, 2002.
- [17] M. S. Olivier. Privacy under conditions of concurrent interaction with multiple parties, 2003. Paper submitted for review.
- [18] PrivacyRight. Control of personal information — the economic benefits of adopting an enterprise-wide permissions management platform. White Paper, 2001.
<http://www.privacyright.com/info/economic.html>.
- [19] J. Reagle and L. F. Cranor. The platform for privacy preferences. *Communications of the ACM*, 42(2):48–55, February 1999.

- [20] M. K. Reiter and A. D. Rubin. Anonymous web transactions with Crowds. *Communications of the ACM*, 42(2):32–48, February 1999.
- [21] Constitution of the Republic of South Africa, 1996. Act 108 of 1996.
- [22] A. Rieke and T. Demuth. JANUS: Server anonymity in the world-wide web. In U. E. Gattiker, editor, *Conference Proceedings EICAR International Conference*, pages 195–208, 2001.
- [23] D. A. Stamper. *Business Data Communications*. Benjamin/Cummings, fourth edition, 1994.
- [24] Tivoli. Tivoli secureway privacy manager — controlling access to consumer information. White paper, IBM, 2000.
<http://www.nasi.com/tivoliprivacymanager-wp.pdf>.