Privacy Enhancing Technologies: A Review

Yun Shen, Siani Pearson

HP Laboratories
HPL-2011-113

Keyword(s):

privacy; privacy enhancing technology;

Abstract:

Organisations handle employees', customers' and third parties' Personally Identifiable Information (PII) in a number of ways and for a variety of reasons; when doing this, it is important that privacy is taken into account. Privacy Enhancing Technologies (PETs) provide a mechanism that helps with this, and can be used in conjunction with higher level policy definition, human processes, training, etc. In this paper we conduct a brief survey of Privacy Enhancing Technologies (PETs) in recent years and show how these may help address different types of privacy harm to employees, customers and, more generally, to data subjects.

External Posting Date: August 6, 2011 [Fulltext]   Approved for External Publication
Internal Posting Date: August 6, 2011 [Fulltext]

Noname manuscript No.
(will be inserted by the editor)

Privacy Enhancing Technologies: A Review

Yun Shen · Siani Pearson

14th June 2011

Abstract Organisations handle employees', customers' and third parties' Personally Identifiable Information (PII) in a number of ways and for a variety of reasons; when doing this, it is important that privacy is taken into account. Privacy Enhancing Technologies (PETs) provide a mechanism that helps with this, and can be used in conjunction with higher level policy definition, human processes, training, etc. In this paper we conduct a brief survey of Privacy Enhancing Technologies (PETs) in recent years and show how these may help address different types of privacy harm to employees, customers and, more generally, to the data subjects (Institute, 2009).

Keywords Privacy · Privacy Enhancing Technologies · PETs

Yun Shen
HP Cloud Services Bristol
Stoke Gifford
Bristol BS34 8QZ
United Kingdom
E-mail: Yun.Shen@hp.com

Siani Pearson
Cloud & Security Lab, HP Labs Bristol
Stoke Gifford
Bristol BS34 8QZ
United Kingdom
E-mail: Siani.Pearson@hp.com

1 Introduction

Privacy and identity are inexorably bound, in that privacy is not an issue if there is no handling of PII. There is no commonly accepted definition of Privacy Enhancing Technologies (PETs), although a good description is provided by UK Information Commissioner's Office as "... any technologies that protect or enhance an individual's privacy, including facilitating individual's access to their rights under the Data Protection Act 1998" (ICO, 2007) and "these design information and communication systems and services in a way that minimises the collection and use of personal data and facilitates compliance with data protection rules making breaches more difficult and/or helping to detect them" (EU, 2007).

In this paper, we conduct a survey of state-of-the-art PETs and consider how they may help resolve organisational privacy concerns. Although there have been a few reviews on PETs carried out by Goldberg et al. (1997); Goldberg (2002, 2007) and other researchers (Argyrakis et al., 2003; Harbird), our work provides more detailed reviews with regard to state-of-the art PETs, categorises them according to their technical contributions and links them to Solove's widely accepted privacy taxonomy (Solove, 2006) which discussed privacy issues from a social science perspective. In particular, our paper also discusses research that focuses on impact and usability of PETs to society that are considered as an enhancement of PETs research and that can enable organisations to adapt PETs in a more sensible way.

The rest of the paper is organised as follows. In Section 2, we categorise state-of-the-art PETs according to their technical contributions. In Section 3, we discuss how these PETs can be linked to Solove's taxonomy as "counter-measures" that prevent privacy violations. In Section 4 we discuss research that focuses on impact and usability of PETs to the society and draw conclusions in Section 5.

2 Review of Privacy Enhancing Technologies

In this section we categorise different types of PETs and give examples, concentrating on the main areas: due to space considerations, it is not possible to survey all PETs in a comprehensive manner.

2.1 PETs for Anonymisation

One promising technology for preserving privacy is anonymity (Goldschlag et al., 1996; Reiter and Rubin, 1998; Dingledine et al., 2004), which provides data minimisation and user identity protection, aiming at preserving privacy at different levels. PETs can provide - in certain contexts (email, payment, web browsing, etc.) - users with complete anonymity or else pseudonymity (i.e. anonymity that is reversible if needed, for example in case of fraud). We focus our discussion on anonymous communication techniques.

2.1.1 Anonymous Communication Techniques

Communication anonymisation targets protection of some of the user's PII, particularly, network addresses of the communicating parties. Several different techniques have been used: trusted infomediaries that remove PII, mix networks to obfuscate the source of a communication, addition of additional traffic or data to make the 'real' data more difficult to mine, etc.

The core of these technologies is hiding correlation between input and output data in order to protect the identity of the end user (data subject). Over the past few years, various technologies, first generation Onion Routing (Goldschlag et al., 1996), Hordes (Levine and Shields, 2002), Crowds (Reiter and Rubin, 1998), Anonymizer[1], and private authentication protocols for mobile scenarios (Abadi, 2003), have been proposed to keep users anonymous. Following these efforts, Dingledine et al. (2004) introduced

[1] http://www.anonymizer.com

Tor, a circuit-based low-latency anonymous communication service, to address limitations in the original Onion routing design by adding perfect forward secrecy, congestion control, directory servers, integrity checking, configurable exit policies, and a practical design for location-hidden services via rendezvous points.. In practical use, there is one alternative to TOR in wide use, called AN.ON/JAP.[2] Díaz et al. (2003) proposed a generalised framework and a new mix design - binomial mix - for expressing batching strategies of a mix. (Carbunar et al., 2007) proposed mechanisms for prevent sensor networks from leaking client interests to the servers when querying. George Danezis (George Danezis and Mathewson, 2003) proposed a type III (Mixminion) anonymous remailer protocol to handle pseudonymity. Morphmix (Rennhard and Plattner, 2002) and Tarzan (Freedman et al., 2002) can provide layered encryption mechanism to further protect packets in the onion routing nodes. Recently there has been interesting research in accountable anonymity (Tsang et al., 2007, 2008).

2.1.2 Various Anonymisation Techniques

Examples include anonymisation of records and logs (Flegel et al., 2002), cookie removal software (Kristol, 2001) and trusted infomediaries that remove PII. Various tools have also been proposed as "countermeasures to surveillance" to preserve online privacy in different scenarios, such as Free Heaven (Dingledine et al., 2000), Bugnosis (Alsaid and Martin, 2003), remailers (e.g. Premail[3], mixmaster[4]), Pretty Good Privacy (PGP)[5], Dephormation[6] etc.

2.2 PETs to Protect Network Invasion

In this category we review PETs relating to attacks on the established systems. These research efforts are treated as assessments on the stability of current PETs.

[2] http://anon.inf.tu-dresden.de/index.en.html
[3] http://www.mirrors.wiretapped.net/security/cryptography/ apps/mail/premail/
[4] http://mixmaster.sourceforge.net/
[5] http://www.pgpi.org/
[6] https://www.dephormation.org.uk/

Demuth (2003) developed LSI (latent semantic indexing) to index standard users' access data set (ADS) in order to identify them. Ollmann (2004) reviewed phishing history and summarised current phishing threats, then proposed defence mechanisms to phishing at the client, server and enterprise side. Avoine (2004) also discussed several attacks to the Juels-Pappu banknote protection scheme in RFID. Jackson et al. (2006) discussed various web browser privacy attacks. Arvind Narayanan et al. have done some interesting work (Narayanan and Shmatikov, 2006, 2008, 2009) in de-anonymising data sets and social networks. In addition to the methods mentioned above, other attacks can be protected against using security mechanisms and methods that increase k-anonymity and/or l-diversity.

2.3 PETs for Identity Management

Identity management deals with identifying individuals and controlling access to resources in a system. There are several major approaches in the marketplace in this area: in particular, Liberty Alliance's federated approach[7], OpenID[8] authentication (a decentralised approach), Identity Metasystem Architecture (Cameron and Jones, 2007) and Generic Bootstrapping Architecture (GBA) (telecommunication focused). PETs associated with identity management aim at identity verification with minimum identity disclosure, and protection against identity theft. For an organisation, this corresponds to release of PII only necessary for the business purpose, such that different third parties may know different (minimised) information about the user, and this process is governed by user choice and consent. This is different to anonymity since some PII and even sensitive information may have to be revealed in a given case.

Considerable work has been done to protect and manage users' identity in past years: privacy enhancing service architectures in mobile domains (Alamäki et al., 2003), online trust negotiation processes (Seamons et al., 2003), anonymising user location data (Gidófalvi et al., 2007; Solanas et al., 2008; Domingo-Ferrer, 2008), privacy en-

[7] http://www.projectliberty.org/
[8] http://openid.net/

hanced claim URIs (Gevers et al., 2007) (now integrated with Microsoft Cardspace). Quite a few research projects, e.g. *PRIME*[9], *FIDIS*[10], *PrimeLife*[11], *Picos*[12], focus on privacy issues for identity management.

2.3.1 Credential Systems

From a privacy point of view, PETs for identity management need to be able to provide authentication (and authorisation) without identification. Credential systems can allow this, by providing only the PII necessary for the transaction or a proof of entitlement. Work on management of attribute credentials linked to identity certificates has been done by the Internet Engineering Task Force (IETF) Public-Key Infrastructure (PKIX) Working Group[13], but that solution is complex in terms of reliance on multiple trusted third parties without fully addressing privacy and anonymity issues. Various anonymous credential schemes have been proposed, most notably those of Chaum (1992) and Brands (2002) and credential identity management (Herzberg and Mass, 2004a; Chaum, 1992, 1986)

2.3.2 Trust Management

A closely related area is trust management, where the information released depends upon an assessment of the recipient's trustworthiness. PETs that may be used to do this include reputation management (Crane and Mont, 2006), integrity checking of remote trusted platforms (Mont and Tomasi, 2001) and other trust management techniques (Pearson, 2005; Herzberg and Mass, 2004b).

2.4 PETs for Data Processing

There are a number of different techniques related to database privacy.

[9] https://www.prime-project.eu/

[10] http://www.fidis.net/

[11] http://www.primelife.eu/

[12] http://www.picos-project.eu/

[13] http://www.ietf.org/html.charters/pkix-charter.html

2.4.1 Privacy Preserving Data Mining

An important issue in preserving privacy when data processing is avoiding leakage of private information by information aggregation. Aggregation differs from identity management as aggregation cannot generate an explicit link to people in their day-to-day life but only in certain situations.

One of the most representative information aggregation technologies is data mining, a process of extracting hidden patterns/information from data. The current trend is especially in favour of Web search log mining, user behaviour mining etc. However, according to a recent survey (Tsai et al., 2006), 70.3% of people surveyed are privacy fundamentalists that "feel companies should not be able to acquire personal information for their organisational needs". It is vital for organisations to carefully use data mining technologies to obtain hidden information without any potential intrusion to the consumers' privacy. A wide range of methods relating to privacy preserving data mining[14] have been proposed to minimise access to customers' privacy: additive data perturbation (Agrawal and Srikant, 2000; Evfimevski et al., 2003); multiplicative data perturbation (Chen and Liu, 2005; Kargupta et al., 2003); data anonymisation (Machanavajjhala et al., 2006; M. Atzori and Pedreschi, 2005; Sweeney, 2002); secure multi-party computation (Pinkas, 2002); privacy preserving multivariate statistical analysis (Du et al., 2004; Yang et al., 2004); probabilistic automaton (Jacquemont et al., 2006); privacy preserving formal methods (Matwin et al., 2005); sampling-based method (Cuzzocrea et al., 2008); k-anonymization classification (Fung et al., 2007); privacy in graph data (Zheleva and Getoor, 2007); statistical disclosure control (Domingo-Ferrer, 2007).

2.4.2 Privacy Management in Data Repositories

Privacy management for data repositories has been developed to ensure that stored data is accessed in a privacy compliant way so that there is no collection of contact in-

[14] A dedicated bibliography is available at http://www.csee.umbc.edu/∼kunliu1/research/privacy_review.html

formation, no collection of long term person characteristics, and k-anonymity (Sweeney, 2002) with large value of k or l-diversity (Machanavajjhala et al., 2006) with a large value of l (to avoid identification or derival of profile characteristics from DBs).

Some mechanisms and solutions have been built to encrypt confidential data when it is stored in data repositories, for example solutions using Translucent Databases (Wayner, 2009). Most of these solutions focus on confidentiality and access control aspects, and have little flexibility in providing policy-driven mechanisms encompassing aspects beyond authentication and authorisation. Miklau and Suciu (2003); Bertino and Ferrari (2002) describe access control policy-based encryption mechanisms for XML (Extensible Markup Language) documents. (Bertino and Ferrari, 2002) describes mechanisms for fine-grained encryption of parts of XML documents, in which decryption keys can either be granted to data receivers or collected from LDAP servers, based on data receivers' credentials. Miklau and Suciu (2003) focuses on related cryptographic mechanisms. Hippocratic Databases (Agrawal et al., 2002) include mechanisms for preserving the privacy of the data they manage. Their proposed architecture is based on the concept of associating privacy metadata (i.e. privacy policies) to data stored in data repositories, along with mechanisms to enforce privacy. Although now withdrawn from production, IBM Tivoli Privacy Manager[15] provided mechanisms for defining fine-grained privacy policies and associating them with data. The privacy policies contain authorisation constraints along with constraints on contextual information and intent. This approach addressed the privacy management problem purely from an access control perspective within a single enterprise. An alternative approach is based on an adaptive privacy management system where data are retrieved from standard data repositories, and parts of these data are encrypted and associated with privacy policies (Mont and Pearson, 2005).

[15] http://www-01.ibm.com/software/tivoli/products/privacy-mgr-e-bus/

Privacy policies enable users and other entities to specify how they would like their personal data to be treated by other parties while limiting access to unauthorised persons. These policies can be taken into account before disclosure of PII, and can govern other ways in which PII is treated. PETs can be used to help with the following: privacy policy creation, use within decision making, and policy enforcement.

There are a number of PETs for privacy management that compare service-side polices about the handling of personal data with preferences expressed by users (for example W3C P3P[16], and PRIME[17]).

Other work has focused on extension of access control and related privacy-enhanced policies. There has been a great deal of work done on defining access control privacy polices: policy specification, modelling and verification tools include EPAL[18], OASIS XACML[19], W3C P3P, Datalog with constraints (Li and Mitchell, 2003), Ponder (Damianou et al., 2001), Platform for Enterprise Privacy Practices (E-P3P) (Karjoth et al., 2003), trust management policies (Keromytis, 1999; Blaze et al., 1996; Chu et al., 1997), RBAC (role based access control) privacy polices (Peleg et al., 2008; Ni et al., 2009), and privacy access control in shared social networks (Carminati and Ferrari, 2008; Carminati et al., 2009). De Capitani di Vimercati et al. (2007) summarised the main desiderata for access control systems and illustrated the main characteristics of access control solutions.

Further work has allowed user consent (Encore, 2009), user assurance requirements (Elahi and Pearson, 2007) or ongoing obligations (Mont, 2005) to be checked within a workflow, even independently of access control. A related technique is cryptographic binding of policies - that describe how personal data should be handled - to the data itself (sticky policies'). Karjoth et al. (2003) introduced the "sticky policy" paradigm and mechanisms for enterprise privacy enforcement. Similar ideas can be found in

[16] http://www.w3.org/P3P/
[17] https://www.prime-project.eu/
[18] http://www.zurich.ibm.com/security/enterprise-privacy/epal
[19] http://www.oasis-open.org/committees/tc_home.php?wg _abbrev=xacml

(Mont et al., 2003). Ardagna et al. (2006) introduced data handling policy to define how the personal information should be dealt with at the receiving party. Schunter and Waidner (2007) introduced a unified policy model towards privacy management. Salim et al. (2007) extended SITDRM to enforce P3P policies.

In addition, there is another category, which relates to checking that the business operation is legally compliant. This category includes forms of automated privacy risk management and compliance checking, and in addition decision support tools that highlight privacy requirements and controls even in a global environment (Pearson et al., 2009).

2.6 Summary

In the above sections we have surveyed a number of areas in which PETs are available or are being developed. Organisations are most likely to be most interested in deploying PETs that help them meet their legal and regulatory responsibilities (in particular, compliance systems, decision support systems and education and training tools), as well as technologies that help them meet their own internal privacy and security policies (such as enhanced access controls). Usage of some of these PETs may render others unnecessary: in particular, anonymising PETs used to decrease linkability of individuals to PII (cf. 'privacy by architecture, as described in (Spiekermann and Cranor, 2009), render PETs that enforce policies about the usage of PII unnecessary (cf. 'privacy by policy' (Spiekermann and Cranor, 2009). On the other hand, there has been little business takeup of the former: they may be more costly, complex and there can be some internal opposition (in particular, from marketing) to PETs that limit collection and selling of PII, as the data could have a business value. Creation of good business cases for PETs is an area that is still relatively undeveloped.

3 PETs and Solove's Taxonomy

Solove's privacy taxonomy (Solove, 2006) outlined the following structure to identify privacy problems:

- **Information Collection**

  Harms: *Surveillance, Interrogation*
- **Information Processing**

  Harms: *Aggregation, Identification, Insecurity, Secondary Use, Exclusion*
- **Information Dissemination**

  Harms: *Breach of Confidentiality, Disclosure, Exposure, Increased Accessibility, Blackmail, Appropriation, Distortion*
- **Invasion**

  Harms: *Intrusion, Decisional Interference*

Solove's taxonomy is the most comprehensive privacy taxonomy to date, with a focus on characterising privacy harms to the end user. In this section we use Solove's taxonomy as a framework for considering current PETs that organisations can use to reduce privacy-related harm to their employees, customers and partners. This can form part of a practical approach to addressing privacy, in that an organisation could carry out a privacy risk assessment to highlight the privacy risks to end users, as part of a PIA of a business process. Then from the analysis in this section it may be deduced which PETs are worthy of consideration for addressing the harms involved.

We consider this mapping further in this section.

3.1 Information Collection

Information collection creates potential privacy violations based on the process of data gathering and in many instances, end users are not aware of the harms incurred by such processes. Especially, surveillance, as a form of information collection, has been viewed as problematic and as violating people's right to privacy. The anonymisation techniques discussed in Section 2.1 can be used to protect data from being collected by third parties. Notice, choice and transparency can be integrated into privacy design (in particular, via information presented to the user via good user interface design).

3.2 Information Processing

Information processing refers to the use, storage, and manipulation of data that has been collected. Privacy issues relating to information processing arise from how data that have already collected are handled. There are various ways to connect/aggregate data together from various sources and link it back to the people to whom it pertains. Identity management technologies discussed in Section 2.3 and anonymisation techniques in Section 2.1 can be used as "counter-measures" to harms caused by identification. Privacy preserving data mining technologies discussed in Section 2.4 and controlled data disclosure such as identity management, anonymisation techniques can solve privacy violation issues caused by aggregation. Privacy management techniques in data repositories discussed in Section 2.4.2 can be applied to fight insecurity and secondary use issues.

3.3 Information Dissemination

Privacy issues relating to information dissemination arise from "the revelation of personal data or the threat of spreading information" (Solove, 2006). An interesting example is personalisation versus privacy: personalisation is not achievable without processing personal data. Solove (2006) pointed out that this category is "one of the broadest groupings of privacy harms". PETs in this category normally take a systematic approach to preserving privacy because various aspects need to be considered such as laws, politics, policies, regulations, best practice, technologies etc. Privacy management techniques in data repositories discussed in Section 2.4.2 can be used to prevent privacy violations in Breach of Confidentiality, Disclosure, and Increased Accessibility. Privacy policies discussed in Section 2.5 enable users and other entities to specify how information can be disseminated.

3.4 Invasion

Invasion involves impingements directly on the individual. Its harms include intrusion and decisional inference. In an information system, PETs to protect invasion discussed in Section 2.2 can be used to protect people's "right to privacy". Privacy preserving data mining techniques discussed in Section 2.4 can be applied to prevent organisations from inferring and influencing people's decision process.

Although it is useful to consider which privacy harms PETs address, this is not a 1-1 mapping: for example, technologies categorised in the same section earlier in this paper can cover different harms in Solove's taxonomy, and some PETs address more than one harm. Furthermore, PETs alone do not mitigate all privacy harms. From our analysis above, it follows that PETs do play a role in resolving privacy concerns. However, they do not resolve all privacy concerns. Why is this? There are a number of reasons, including: vested interests in obtaining personal information (for example, for marketing); lack of regulatory powers, lack of user awareness of privacy risks and other factors that prevent development of effective economic models for organisational investment in PETs; a rate of technological change that is high enough to introduce new privacy risks at at least the rate that older risks can be addressed technologically and legally (for instance, new privacy risks have been introduced by RFID tags, social computing and cloud computing); the complexity of privacy requirements in a global environment and the contextual nature of privacy risk, making it difficult for people to understand the privacy requirements in a given case; the increasing distribution and ease of exposure of personal information in a global online environment. In order to encourage adoption of PETS; it would not be just a question of increased fines for non-compliance, but encouragement of a new mind set within organisations to do the right thing', in particular via expansion of the use of Privacy Impact Assessments (PIA) that help organisations assess the impact of their operations on personal privacy, and more generally encouragement of accountability both within organisations and to external stakeholders. PETs can help with this, and moreover will be needed to help with this due to the potential complexity of determining requirements and the need to track and enforce obligations.

4 Enhancement of Current PETs

There are a number of important aspects that PETs help address that are orthogonal to Solove's characterisation: in particular, the Fair Information Principles[20]. For example, transparency (which perhaps maps most closely to Solove's interrogation category, and which can be addressed via advanced usability techniques) and choice. These principles should be included within Privacy by Design, and also individually are addressed by selected PETs.

There are some areas that PETs can be further developed where present processes are manual, such as data subject access requests. Improvements could also be made in areas that current methods may not be sophisticated, for example, when synchronising updates to (and user preferences associated with) replicated PII within an organisation. In addition, research can play a useful role in encouraging take-up of mechanisms (both technical and non-technical) to prevent harm: for example, providing economic analysis and helping develop business models for the introduction of PETs.

In this section, we discuss research that focuses on impact and usability of PETs to society: usability addresses the harms of non-transparency and lack of choice, privacy by design focuses on holistic approaches to design privacy aware systems and economics of privacy deals with the harms of non-practicality and non-adoption. These research efforts should be considered as enhancement of PETs research and can enable organisations to adapt PETs in a more sensible way.

[20] http://www.ftc.gov/reports/privacy3/fairinfo.shtm

4.1 Usability

This is vital for the success of PETs since it is important to provide an intuitive and straightforward user interface as there are a broad range of people with different experience and skills.

AT&Ts Privacy Bird (Cranor et al., 2006) is a plug-in for Internet Explorer that monitors P3P policies for the user it has an easy to use interface, but with very limited options. An alternative approach is to ask a series of dynamic questions which the user can answer to inform agents about their privacy preferences, and by these means to set user policies see for example the approach of (Irwin and Yu, 2005). It is also worth considering the balance between flexibility in policy definition and usability: for example, a pre-defined set of natural language clauses might be used as the policies, and evidence could be provided by the system that these are satisfied on the back end (Elahi and Pearson, 2007). Patrick and Kenny (2003) described the HCI requirements of an effective privacy interface design. The PRIME project (Pettersson et al., 2005; Bergmann et al., 2005) used three UI paradigms - role-centred, relationship-centred and townmap-based paradigms - for privacy-enhanced identity management in the PRIME project. Andersson et al. (2005) discussed the socio-psychological factors and HCI aspects that influence end users' trust in privacy enhancing identity management. Hawkey and Inkpen (2006) examined the privacy comfort levels of participants if others can view traces of their web browsing activity. At the implementation level, Kobsa (2003) adopted a redundant component array (RAIC) architecture to personalised web systems so that they can dynamically adjust to the current prevailing privacy concerns and requirements without burdening the application with privacy management tasks. Iachello and Hong (2007) summarised previous research and proposed new research directions in privacy aware HCI. Work is currently being carried out in a number of projects related to how to visualise privacy to the user: MobiLife[21], VOME[22] ,

[21] http://www.ist-mobilife.org/

[22] http://www.vome.org.uk/

PRIMELife[23]. One technique is to use privacy labelling to convey privacy policies and preferences to data subjects (Kelley et al., 2009).

4.2 Privacy by Design

Technologies in this category protect (or regulate organisations to protect) users' private information from the system design stage. These technologies mostly take a holistic approach to design privacy aware systems because privacy by design is a methodology that can use PETs.

There is an increasing awareness for the need for design for privacy from both companies and governmental organisations (Microsoft, 2009). Hippocratic database (Agrawal et al., 2002) was one of the initial attempts to protect user privacy by design. Similar ideas were also discussed in IBM Tivoli Privacy Manager[24] and translucent databases (Wayner, 2009). Sion et al. (2007) introduced a set of layered mechanisms and various protocols for securely storing data items and providing full computational privacy. Domingo-Ferrer and Bras-Amorós (2008) proposed a type of combinatorial design to reduce the number of required keys in peer-to-peer private information retrieval scenario. Pearson et al. (2009) suggested a variety of guidelines and techniques for newly-emerged cloud computing to ensure that the risks to privacy are mitigated at the design stage. At the implementation level, Berghe and Schunter (2006) introduced Privacy Injector to allow the users to add privacy enforcement to existing application. EnCoRe (Encore, 2009), a TSB project, aims to design reliable consent and revocation mechanisms for private information data flow. Regulation is also indispensable as an essential technology to preserve privacy. The Information Commissioner's Office (ICO) in 2006 (Wood, 2006) proposed privacy impact assessment (PIA) and surveillance impact assessments (SIA) as possible solutions to surveillance-related privacy issues.

[23] http://www.primelife.eu/

[24] http://www-01.ibm.com/software/tivoli/products/privacy-mgr-e-bus/

4.3 Economics of Privacy

Privacy, as a multi-disciplinary issue, should be analysed from various perspectives including the point of view of economics. Privacy emerged as a major consumer issue in the mid 1990s with the development of the Internet. However, "for most consumers, the costs of exercising choice - although not high - are not worth the perceived benefits" (Muris, 2004). In other words, the Federal Trade Commission (FTC)'s Fair Information Practices (FIP) model[25] based upon enforcement is economically inefficient. Ironically, although 69% of adults agree that "consumers have lost all control over how personal information is collected and used by companies", surveys[26] reveal that most Americans are "privacy pragmatists", who care about privacy but are willing to share information when they see tangible benefits and they believe care is taken to protect that information. From the point view of companies, as well as the incentives to take up PETs that include increased user trust, it is advisable also to take into account the costs if they do not take up PETs, for example, being sued by customers if insufficient safeguards are taken to protect their personal data.

The early economic analysis of privacy - the "Free Market" approach (Posner, 1981) - concludes that markets for personal information would work as well as markets for conventional goods and services. However, Hermalin and Katz (2006) point out that such an approach may not provide an economically efficient outcome. Recently, various approaches have been proposed to discuss how the collection and use of personal information within a single market (Hermalin and Katz, 2006; Acquisti and Varian, 2005; Taylor, 2004a; Dodds, 2002) or across markets (Taylor, 2004b; Calzolari and Pavan, 2004) affects the efficiency of market outcomes. If the consumers are allowed to set their own value of personal information, it would be interesting to see how individuals' "willingness to accept" (WTA) the use of their information (assuming that consumers have property rights over their personal information) interacts with their "willingness

[25] http://www.ftc.gov/reports/privacy3/fairinfo.shtm

[26] Further details can be found at http://www.harrisinteractive.com/harris_poll/index .asp?PID=365

to pay" (WTP) to protect their information from being exploited. PVNets[27] is investigating this issue; however, future research is needed to expand these research efforts into both government and commercial context. George Danezis et al. (Danezis and Wittneben, 2006) discussed target selection strategies for maximising surveillance (or disruption) return based on data collected from a real social network. Kerschbaum (2008) discussed privacy-preserving benchmarking enterprise systems for the economic advantage of the service provider. McDonald and Cranor (2008) concluded that consumers' time would be worth $781 billion if they were to read the privacy policy for each site they visit just once a year.

5 Conclusion

We have surveyed PETs that have been proposed in the past few years, including recent research related to economics and usability of privacy. This provides evidence of a number of techniques that might be used by organisations in order to decrease harm to data subjects. Privacy is very contextual, in the sense that privacy harms will vary according to the context. To determine which techniques might be most appropriate in a business context, a privacy risk assessment needs to be carried out in order to determine the possible harms and the degree of risk. Therefore, in our paper we have analysed the technological methods that might be used by organisations to address privacy with respect to the categories discussed in Solove's privacy taxonomy (Solove, 2006).

References

Martín Abadi. Private authentication. In *Privacy Enhancing Technologies*, volume 2482 of *LNCS*, pages 27–40, 2003. ISBN 3-540-00565-X.

Alessandro Acquisti and Hal R. Varian. Conditioning prices on purchase history. *Marketing Science*, 24:367–381, 2005.

[27] http://www.pvnets.org/

Rakesh Agrawal and Ramakrishnan Srikant. Privacy-preserving data mining. In *Proc. of the ACM SIGMOD*, pages 439–450. ACM Press, May 2000. ISBN 1-581-13218-2.

Rakesh Agrawal, Jerry Kiernan, Ramakrishnan Srikant, and Yirong Xu. Hippocratic databases. In *VLDB*, pages 143–154. Morgan Kaufmann, 2002.

Tero Alamäki, Margareta Björksten, Péter Dornbach, Casper Gripenberg, Norbert Gyorbíró, Gábor Márton, Zoltán Németh, Timo Skyttä, and Mikko Tarkiainen. Privacy enhancing service architectures. In *Privacy Enhancing Technologies*, volume 2482 of *LNCS*, pages 99–109. Springer, 2003. ISBN 3-540-00565-X.

Adil Alsaid and David Martin. Detecting web bugs with bugnosis: Privacy advocacy through education. In *Privacy Enhancing Technologies*, volume 2482 of *Lecture Notes in Computer Science*, pages 13–26. Springer, 2003.

C. Andersson, J. Camenisch, S. Crane, S. Fischer-Hubner, R. Leenes, S. Pearsorr, J. S. Pettersson, and D. Sommer. Trust in prime. In *Signal Processing and Information Technology*, pages 552–559, 2005.

Claudio Agostino Ardagna, Sabrina De Capitani di Vimercati, and Pierangela Samarati. Enhancing user privacy through data handling policies. In *20th Annual IFIP WG 11.3 Working Conference on Data and Applications Security*, volume 4127 of *Lecture Notes in Computer Science*, pages 224–236, 2006. ISBN 3-540-36796-9.

John Argyrakis, Stefanos Gritzalis, and Chris Kioulafas. Privacy enhancing technologies: A review. In *EGOV*, pages 282–287, 2003.

Gildas Avoine. Privacy issues in RFID banknote protection schemes. In Jean-Jacques Quisquater, Pierre Paradinas, Yves Deswarte, and Anas Abou El Kadam, editors, *International Conference on Smart Card Research and Advanced Applications – CARDIS*, pages 33–48, Toulouse, France, August 2004. IFIP.

Chris Vanden Berghe and Matthias Schunter. Privacy injector - automated privacy enforcement through aspects. In *Privacy Enhancing Technologies*, volume 4258 of *LNCS*, pages 99–117. Springer, 2006. ISBN 3-540-68790-4.

Mike Bergmann, Martin Rost, and John Sren Pettersson. Exploring the feasibility of a spatial user interface paradigm for privacy-enhancing technologys. In *ISD '05*, pages

437–448. Springer Verlag, 2005.

Elisa Bertino and Elena Ferrari. Secure and selective dissemination of xml documents. *ACM Trans. Inf. Syst. Secur.*, 5(3):290–331, 2002. ISSN 1094-9224.

Matt Blaze, Joan Feigenbaum, and Jack Lacy. Decentralized trust management. Technical report, 1996.

Stefan Brands. A technical overview of digital credentials. *International Journal on Information Security*, 2002.

Giacomo Calzolari and Alessandro Pavan. On the optimality of privacy in sequential contracting. Discussion Papers 1394, Northwestern University, Center for Mathematical Studies in Economics and Management Science, July 2004.

Kim Cameron and Michael B. Jones. Design rationale behind the identity metasystem architecture. In *ISSE/SECURE 2007 Securing Electronic Business Processes*, 2007.

Bogdan Carbunar, Yang Yu, Weidong Shi, Michael Pearce, and Venu Vasudevan. Query privacy in wireless sensor networks. In *SECON*, pages 203–212, 2007.

Barbara Carminati and Elena Ferrari. Privacy-aware collaborative access control in web-based social networks. In *Proceeedings of the 22nd annual IFIP WG 11.3 working conference on Data and Applications Security*, pages 81–96, Berlin, Heidelberg, 2008. Springer-Verlag. ISBN 978-3-540-70566-6.

Barbara Carminati, Elena Ferrari, Raymond Heatherly, Murat Kantarcioglu, and Bhavani Thuraisingham. A semantic web based framework for social network access control. In *SACMAT '09*, pages 177–186, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-537-6. doi: http://doi.acm.org/10.1145/1542207.1542237.

D Chaum. Showing credentials without identification. signatures transferred between unconditionally unlinkable pseudonyms. In *EUROCRYPT '85*, pages 241–244, New York, NY, USA, 1986. Springer-Verlag New York, Inc. ISBN 0-387-16468-5.

David Chaum. Achieving electronic privacy. *Scientific American*, pages 91–101, 1992.

K. Chen and L. Liu. Privacy preserving data classification with rotation perturbation. In *ICDM '05*, pages 589–592, Houston, TX, November 2005.

Yang-Hua Chu, Joan Feigenbaum, Brian A. LaMacchia, Paul Resnick, and Martin Strauss. Referee: Trust management for web applications. *Computer Networks*, 29 (8-13):953–964, 1997.

Stephen Crane and Marco Casassa Mont. A customizable reputation-based privacy assurance system using active feedback. In *Securecomm and Workshops, 2006*, pages 1–8, 2006.

Lorrie Faith Cranor, Praveen Guduru, and Manjula Arjula. User interfaces for privacy agents. *ACM Trans. Comput.-Hum. Interact.*, 13(2):135–178, 2006. ISSN 1073-0516. doi: http://doi.acm.org/10.1145/1165734.1165735.

Alfredo Cuzzocrea, Vincenzo Russo, and Domenico Saccà. A robust sampling-based framework for privacy preserving olap. In *DaWaK '08*, pages 97–114, Berlin, Heidelberg, 2008. Springer-Verlag. ISBN 978-3-540-85835-5.

Nicodemos Damianou, Naranker Dulay, Emil Lupu, and Nicodemos Damianou Naranker Dulay. The ponder policy specification language. In *Lecture Notes in Computer Science*, pages 18–38. Springer-Verlag, 2001.

George Danezis and Bettina Wittneben. The economics of mass surveillance - and the questionable value of anonymous communications. In *WEIS 2006*, 2006.

S. De Capitani di Vimercati, S. Foresti, S. Jajodia, and P. Samarati. Access control policies and languages. *IJCSE*, 3(2), 2007.

Thomas Demuth. A passive attack on the privacy of web users using standard log information. In *Privacy Enhancing Technologies*, volume 2482 of *LNCS*, pages 179–193. Springer, 2003. ISBN 3-540-00565-X.

Claudia Díaz, Joris Claessens, and Bart Preneel. Apes - anonymity and privacy in electronic services. *Datenschutz und Datensicherheit*, 27(3), 2003.

Roger Dingledine, Michael J. Freedman, and David Molnar. The free haven project: Distributed anonymous storage service. In *In Proceedings of the Workshop on Design Issues in Anonymity and Unobservability*, pages 67–95, 2000.

Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The second-generation onion router. In *Proceedings of the 13th USENIX Security Symposium*, pages 303–

320, San Diego, CA, USA, August 2004.

Stefan Dodds. Privacy and endogenous monitoring choice when private information is a public good. Working Papers 1010, Queen's University, Department of Economics, September 2002.

Josep Domingo-Ferrer. A three-dimensional conceptual framework for database privacy. In *Secure Data Management*, pages 193–202, 2007.

Josep Domingo-Ferrer. Location privacy via unlinkability: an alternative to cloaking and perturbation. In *PAIS*, pages 1–2, 2008.

Josep Domingo-Ferrer and Maria Bras-Amorós. Peer-to-peer private information retrieval. In *Privacy in Statistical Databases*, pages 315–323, 2008.

W. Du, Y. S. Han, and S. Chen. Privacy-preserving multivariate statistical analysis: Linear regression and classification. In *SDM '04*, Lake Buena Vista, FL, April 2004.

Tariq Ehsan Elahi and Siani Pearson. Privacy assurance: Bridging the gap between preference and practice. In *TrustBus*, pages 65–74, 2007.

Encore. Ensuring consent and revocation, 2009. `http://www.encore-project.info/index.html`.

European Union EU. Press release: Privacy enhancing technologies (pets), 2007. `http://europa.eu/rapid/pressReleasesAction.do?reference=MEMO/07/159&format=HTML&aged=0&language=EN&guiLanguage=en.%20Reference:%20MEMO/07/159`.

A. Evfimevski, J. Gehrke, and R. Srikant. Limiting privacy breaches in privacy preserving data mining. In *Proceedings of the ACM SIGMOD/PODS Conference*, pages 211–222, San Diego, CA, June 2003.

Ulrich Flegel, Pseudonymizing Unix, and Log Files. Pseudonymizing unix log files. In *Infrastructure Security*, 2002.

Michael J. Freedman, Emil Sit, Josh Cates, and Robert Morris. Introducing tarzan, a peer-to-peer anonymizing network layer. In *IPTPS '01: Revised Papers from the First International Workshop on Peer-to-Peer Systems*, pages 121–129, London, UK, 2002. Springer-Verlag. ISBN 3-540-44179-4.

B. C. M. Fung, Ke Wang, and P. S. Yu. Anonymizing classification data for privacy preservation. *TKDE*, 19(5):711–725, May 2007.

Roger Dingledine George Danezis and Nick Mathewson. Mixminion: Design of a type iii anonymous remailer protocol. In *Privacy Enhancing Technologies*, Proceedings of the 2003 IEEE Symposium on Security and Privacy, pages 2 – 15. IEEE, 2003.

Steven Gevers, Verslype Verslype, and Bart De Decker. Enhancing privacy in identity management systems. In *WPES '07*, pages 60–63, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-883-1. doi: http://doi.acm.org/10.1145/1314333.1314344.

Győző Gidófalvi, Xuegang Huang, and Torben Bach Pedersen. Privacy-preserving data mining on moving object trajectories. In *MDM*, pages 60–68, 2007.

Ian Goldberg. Privacy-enhancing technologies for the internet, ii: Five years later. In *Privacy Enhancing Technologies*, pages 1–12, 2002.

Ian Goldberg. *Privacy-Enancing Technologies for the Internet III: Ten Years Later*, pages 3–18. 2007.

Ian Goldberg, David Wagner, and Eric A. Brewer. Privacy-enhancing technologies for the internet. In *Privacy Enhancing Technologies*, pages 103–109, 1997.

David M. Goldschlag, Michael G. Reed, and Paul F. Syverson. Hiding routing information. In *in Information Hiding*, pages 137–150. Springer-Verlag, 1996.

Rae Harbird. Privacy by design: An overview of privacy enhancing technologies. `http://www.ico.gov.uk/upload/documents/pdb_report_html/pbd_pets_paper.pdf`.

Kirstie Hawkey and Kori M. Inkpen. Examining the content and privacy of web browsing incidental information. In *WWW '06*, pages 123–132, New York, NY, USA, 2006. ACM. ISBN 1-59593-323-9. doi: http://doi.acm.org/10.1145/1135777.1135801.

Benjamin Hermalin and Michael Katz. Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quantitative Marketing and Economics*, 4(3):209–239, September 2006.

Amir Herzberg and Yosi Mass. Relying party credentials framework. *Electronic Commerce Research*, 4(1-2):23–39, 2004a. ISSN 1389-5753.

Amir Herzberg and Yosi Mass. Relying party credentials framework. *Electronic Commerce Research*, 4(1-2):23–39, 2004b.

Giovanni Iachello and Jason Hong. *End-User Privacy in Human-Computer Interaction.* Now Publishers Inc., Hanover, MA, USA, 2007. ISBN 1601980760, 9781601980762.

Information Commissioner's Office UK ICO. Data protection guidance note: Privacy enhancing technologies, 2007. `http://www.ico.gov.uk/upload/documents/library/data_protection/detailed_specialist_guides/privacy_enhancing_technologies_v2.pdf`.

The Franklin Institute. 2009 franklin institute awards benjamin franklin medal in electrical engineering, 2009. `http://www.fi.edu/franklinawards/09/bf_elecengineer.html`.

Keith Irwin and Ting Yu. Determining user privacy preferences by asking the right questions: an automated approach. In *WPES '05: Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pages 47–50, New York, NY, USA, 2005. ACM. ISBN 1-59593-228-3. doi: http://doi.acm.org/10.1145/1102199.1102209.

Collin Jackson, Andrew Bortz, Dan Boneh, and John C. Mitchell. Protecting browser state from web privacy attacks. In *WWW '06*, pages 737–744, New York, NY, USA, 2006. ACM. ISBN 1-59593-323-9.

Stephanie Jacquemont, Francois Jacquenet, and Marc Sebban. Sequence mining without sequences: A new way for privacy preserving. In *ICTAI '06*, pages 347–354. IEEE Computer Society, 2006. ISBN 0-7695-2728-0.

Hillol Kargupta, Kun Liu, and Jessica Ryan. Privacy sensitive distributed data mining from multi-party data. In *Proceedings of the first NSF/NIJ Symposium on Intelligence and Security Informatics*, Lecture Notes in Computer Science, pages 336–342, Tucson, AZ, June 2003. Springer Berlin/Heidelberg.

Günter Karjoth, Matthias Schunter, and Michael Waidner. Platform for enterprise privacy practices: Privacy-enabled management of customer data. In *Privacy Enhancing Technologies*, volume 2482 of *LNCS*, 2003. ISBN 3-540-00565-X.

Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. A "nutrition label" for privacy. In *SOUPS '09: Proceedings of the 5th Symposium on Usable Privacy and Security*, pages 1–12, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-736-3.

A. Keromytis. The keynote trust-management system, version 2. *IETF RFC*, 2704: 164–173, 1999.

Florian Kerschbaum. Building a privacy-preserving benchmarking enterprise system. *Enterp. Inf. Syst.*, 2(4):421–441, 2008. ISSN 1751-7575.

Alfred Kobsa. A component architecture for dynamically managing privacy constraints in personalized web-based systems. In *Privacy Enhancing Technologies, Third International Workshop, PET 2003, Dresden, Germany, March 26-28, 2003, Revised Papers*, volume 2760 of *LNCS*, pages 177–188. Springer, 2003.

David M. Kristol. Http cookies: Standards, privacy, and politics. *ACM Trans. Internet Technol.*, 1(2):151–198, 2001. ISSN 1533-5399. doi: http://doi.acm.org/10.1145/502152.502153.

Brian Neil Levine and Clay Shields. Hordes: a multicast based protocol for anonymity. *J. Comput. Secur.*, 10(3):213–240, 2002. ISSN 0926-227X.

Ninghui Li and John C. Mitchell. Datalog with constraints: A foundation for trust management languages. In *PADL 03*, pages 58–73. Springer-Verlag, 2003.

Dennis V. Lindley. The probability approach to the treatment of uncertainty in artificial intelligence and expert systems. *Statistical Science*, 2(1):17 – 24, 1987.

F. Giannotti M. Atzori, F. Bonchi and D. Pedreschi. Blocking anonymity threats raised by frequent itemset mining. In *ICDM '05*, November 2005.

Ashwin Machanavajjhala, Johannes Gehrke, Daniel Kifer, and Muthuramakrishnan Venkitasubramaniam. l-diversity: Privacy beyond k-anonymity. In *22nd IEEE International Conference on Data Engineering*, 2006.

Stan Matwin, Amy P. Felty, István T. Hernádvölgyi, and Venanzio Capretta. Privacy in data mining using formal methods. In *TLCA*, pages 278–292, 2005.

Aleecia M. McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. *I/S: A Journal of Law and Policy for the Information Society*, 2008.

Microsoft. Privacy guidelines for developing software products and services, 2009. `http://www.microsoft.com/downloads/details.aspx?FamilyId=C48CF80F-6E87-48F5-83EC-A18D1AD2FC1F&displaylang=en`.

Gerome Miklau and Dan Suciu. Controlling access to published data using cryptography. In *VLDB '2003: Proceedings of the 29th international conference on Very large data bases*, pages 898–909. VLDB Endowment, 2003. ISBN 0-12-722442-4.

M. Mont and L. Tomasi. A distributed system, adaptive to trust assessment, based on peer-to-peer e-records replication and storage. In *FTDCS '01: Proceedings of the 8th IEEE Workshop on Future Trends of Distributed Computing Systems*, page 89, Washington, DC, USA, 2001. IEEE Computer Society.

Marco Casassa Mont. Handling privacy obligations in enterprises: important aspects and technical approaches. *Comput. Syst. Sci. Eng.*, 20(6), 2005.

Marco Casassa Mont and Siani Pearson. An adaptive privacy management system for data repositories. In *TrustBus*, pages 236–245, 2005.

Marco Casassa Mont, Siani Pearson, and Pete Bramhall. Towards accountable management of identity and privacy: Sticky policies and enforceable tracing services. Technical report, HP Laboratories Bristol, 2003.

Timothy J. Muris. The federal trade commission and the future development of u.s. consumer protection policy. *George Mason Law & Economics Research Paper*, 2004.

Arvind Narayanan and Vitaly Shmatikov. How to break anonymity of the netflix prize dataset. *CoRR*, abs/cs/0610105, 2006.

Arvind Narayanan and Vitaly Shmatikov. Robust de-anonymization of large sparse datasets. In *IEEE Symposium on Security and Privacy*, pages 111–125, 2008.

Arvind Narayanan and Vitaly Shmatikov. De-anonymizing social networks. *CoRR*, abs/0903.3276, 2009.

Qun Ni, Elisa Bertino, Jorge Lobo, and Seraphin B. Calo. Privacy-aware role-based access control. *IEEE Security and Privacy*, 7(4):35–43, 2009. ISSN 1540-7993. doi:

http://doi.ieeecomputersociety.org/10.1109/MSP.2009.102.

Gunter Ollmann. The phishing guide. understanding & preventing phishing attacks. Technical report, Next Generation Security Software Ltd., 2004.

Andrew S. Patrick and Steve Kenny. From privacy legislation to interface design: Implementing information privacy in human-computer interactions. In *Privacy Enhancing Technologies*, volume 2760 of *LNCS*, pages 107–124. Springer, 2003.

Siani Pearson. Trusted computing: Strengths, weaknesses and further opportunities for enhancing privacy. In *iTrust*, pages 305–320, 2005.

Siani Pearson, Tomas Sander, and Rajneesh Sharma. A privacy management tool for global outsourcing. In *DPM*, 2009.

Mor Peleg, Dizza Beimel, Dov Dori, and Yaron Denekamp. Situation-based access control: Privacy management via modeling of patient data access scenarios. *J. of Biomedical Informatics*, 41(6):1028–1040, 2008. ISSN 1532-0464. doi: http://dx.doi.org/10.1016/j.jbi.2008.03.014.

John Sören Pettersson, Simone Fischer-Hübner, Ninni Danielsson, Jenny Nilsson, Mike Bergmann, Sebastian Clauss, Thomas Kriegelstein, and Henry Krasemann. Making prime usable. In *SOUPS '05*, pages 53–64, New York, NY, USA, 2005. ACM. ISBN 1-59593-178-3. doi: http://doi.acm.org/10.1145/1073001.1073007.

B. Pinkas. Cryptographic techniques for privacy preserving data mining. *SIGKDD Explorations*, 4(2):12–19, 2002.

Richard A. Posner. The economics of privacy. *American Economic Review*, 71(2): 405–409, 1981.

Michael K. Reiter and Aviel D. Rubin. Crowds: anonymity for web transactions. *ACM Trans. Inf. Syst. Secur.*, 1(1):66–92, 1998. ISSN 1094-9224.

Marc Rennhard and Bernhard Plattner. Introducing morphmix: peer-to-peer based anonymous internet usage with collusion detection. In *WPES '02: Proceedings of the 2002 ACM workshop on Privacy in the Electronic Society*, pages 91–102, New York, NY, USA, 2002. ACM. ISBN 1-58113-633-1. doi: http://doi.acm.org/10.1145/644527.644537.

Farzad Salim, Nicholas Paul Sheppard, and Reihaneh Safavi-Naini. Enforcing p3p policies using a digital rights management system. In *Privacy Enhancing Technologies*, volume 4776 of *LNCS*, pages 200–217. Springer, 2007. ISBN 978-3-540-75550-0.

Matthias Schunter and Michael Waidner. Simplified privacy controls for aggregated services - suspend and resume of personal data. In Nikita Borisov and Philippe Golle, editors, *Privacy Enhancing Technologies*, volume 4776 of *Lecture Notes in Computer Science*, pages 218–232. Springer, 2007.

Kent E. Seamons, Marianne Winslett, Ting Yu, Lina Yu, and Ryan Jarvis. Protecting privacy during on-line trust negotiation. In *Privacy Enhancing Technologies*, volume 2482 of *LNCS*, pages 129–143. Springer, 2003. ISBN 3-540-00565-X.

Radu Sion, Sumeet Bajaj, Bogdan Carbunar, and Stefan Katzenbeisser. Ns2: networked searchable store with correctness. In *VLDB '07*, pages 1342–1345. VLDB Endowment, 2007. ISBN 978-1-59593-649-3.

Agusti Solanas, Josep Domingo-Ferrer, and Antoni Martínez-Ballesté. Location privacy in location-based services: Beyond ttp-based schemes. In *PiLBA*, 2008.

Daniel J. Solove. A Taxonomy of Privacy. *GWU Law School Public Law Research Paper No. 129*, 2006.

Sarah Spiekermann and Lorrie Faith Cranor. Engineering privacy. *IEEE Trans. Softw. Eng.*, 35(1):67–82, 2009. ISSN 0098-5589. doi: http://dx.doi.org/10.1109/TSE.2008.88.

Latanya Sweeney. k-anonymity: a model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570, 2002. ISSN 0218-4885. doi: http://dx.doi.org/10.1142/S0218488502001648.

Curt Taylor. Privacy and information acquisition in competitive markets. Berkeley Olin Program in Law & Economics, Working Paper Series 1165, Berkeley Olin Program in Law & Economics, May 2004a.

Curtis R. Taylor. Consumer privacy and the market for customer information. *RAND Journal of Economics*, 35(4):631–650, Winter 2004b.

Janice Tsai, Lorrie F. Cranor, Alessandro Acquisti, and Christina M. Fong. What's It To You? A Survey of Online Privacy Concerns and Risks. *SSRN eLibrary*, 2006.

Patrick P. Tsang, Man Ho Au, Apu Kapadia, and Sean W. Smith. Blacklistable anonymous credentials: blocking misbehaving users without ttps. In *ACM Conference on Computer and Communications Security*, pages 72–81, 2007.

Patrick P. Tsang, Man Ho Au, Apu Kapadia, and Sean W. Smith. Perea: towards practical ttp-free revocation in anonymous authentication. In *ACM Conference on Computer and Communications Security*, pages 333–344, 2008.

Peter Wayner. *Translucent Databases 2nd Edition: Confusion, misdirection, randomness, sharing, authentication and steganography to defend privacy.* CreateSpace, Paramount, CA, 2009. ISBN 1441421343, 9781441421340.

David Murakami Wood. A report on the surveillance society. Technical report, Information Commissioner's Office, 2006.

Zhiqiang Yang, Rebecca N. Wright, and Hiranmayee Subramaniam. Experimental analysis of privacy-preserving statistics computation. In *In Proc. of the VLDB Worshop on Secure Data Management*, pages 55–66, 2004.

Elena Zheleva and Lise Getoor. Preserving the privacy of sensitive relationships in graph data. In *Proceedings of the First International Workshop on Privacy, Security, and Trust in KDD*, pages 153–171, August 2007.