Keeping Ubiquitous Computing to Yourself: a practical model for user control of privacy

Blaine A. Price[1], Karim Adam, Bashar Nuseibeh
Computing Research Centre, The Open University, MK7 6AA, UK
[B.A.Price, K.A.Adam, B.Nuseibeh]@open.ac.uk

Abstract

As with all the major advances in information and communication technology ubiquitous computing introduces new risks to individual privacy. In this paper we identify specific new elements of personal privacy at risk through the widespread use of ubiquitous computing (ubicomp). Our analysis of privacy protection in ubicomp has identified four layers through which users must navigate: the regulatory regime they are currently in, the type of ubicomp service required, the type of data being disclosed, and their personal privacy policy. Case studies and scenarios are used to analyze a range of ubicomp interactions. We illustrate and compare the protection afforded by existing and proposed regulation and by some major models for user control of privacy. We identify the shortcomings of each and propose a model which allows user control of privacy levels in a ubicomp environment incorporating an economics-based approach in order to allow users to balance the trade-offs in privacy and service provision. We incorporate the concept of noise to protect privacy and the necessity of supporting a benign form of deception. We conclude with a case study illustrating an internationally-applicable model for privacy selection and control in ubicomp. This model extends existing work using an economics-based approach. It allows users to explicitly relax their privacy constraints and employs privacy protecting noise in order to receive services.

Keywords: ubiquitous computing, privacy,

1 Introduction, Motivation and Scope

It is estimated that there are some 2 billion mobile telephones in global use. An increasing proportion of these devices qualify as computers in their own right. Ubiquitous computing (ubicomp) has become a mainstream activity applicable to a sizeable population within developed countries.

Every major advance in information and communication technology since the late 19th Century has raised new concerns about individual privacy. The consequences of ignoring these concerns have ranged from receiving unsolicited e-mail or telephone

[1] Contact Author. Telephone +44 1908 653 701, Fax: +44 1908 652 140

calls during dinner; to the deaths of hundreds of thousands in extermination camps (Black, 2001). The former has prompted a patchwork of regulation or self-regulation while the later prompted many European countries to institute strong privacy (or data protection) laws.

The risks (such as fraud and identity theft) are so great, and such a significant proportion of the planet's population are potentially affected, that user control of privacy protection in ubicomp is essential.

Although data protection/privacy is not a new problem, ubicomp introduces a new privacy risk: timely and accurate location data for an individual (both real-time and historical) being made available. This paper concerns the new privacy risks created by this functionality and the risks of the release of personal information in a ubicomp setting offering Location Based Services (LBSs). Duckham and Kulik (to appear) identify the risks of location data becoming public - both for real-time data (location-based spam and stalking), as well as for historical data (intrusive inferences about personal life, political view, or health).

Others have demonstrated the need for explicit user control of privacy in ubiquitous computing (Bellotti & Sellen, 1993), and the complicated nature of user choice regarding what to disclose to whom in a networked world (Palen & Dourish, 2003). In this paper, we assume that the user is aware of his or her privacy needs. We also assume they know to whom they wish to disclose personal data, or hide personal data from. Privacy sensitivity is highly individual (Dawson et al., 2003) ranging from "naive and completely open" to "ultra paranoid and non-revealing". Although others have attempted to guide users on their privacy risks (Ackerman & Cranor, 1999; AT&T, 2003) or suggest interface metaphors that encapsulate privacy preferences between one user and another (Lederer et al., 2002), we assume that the user has already made these choices. A user's policy may have been generated in a number of ways; including choosing a representative template from a trusted third party (such as a consumer advocate) or from a community of peers providing suitable policies (Yee & Korba, 2005). We are concerned with providing the user with the necessary tools to protect their privacy in a global ubicomp environment.

Our model is based on an analysis of the layers of control afforded to the user, who is located at the centre of our model (shown in extended form in Table 2). A user is an individual, identifiable human being. The user will have a variety of attributes, including a great deal of personally identifying information (PII). In our model, the PII forms a discrete layer surrounding the user.

The *types of services* available to a ubicomp user with an explicit interface form the next layer outwards in our analysis of privacy in ubicomp. Since we are concerned with user control of privacy, we restrict our discussion to the class of active personal ubicomp devices which have an explicit user interface, such as a PDA or mobile telephone. Although we do not deny the importance of privacy for passive devices such as Active Badges or RFID tags, we do not discuss them here because of the lack of user control available. Since we are restricting our discussion to a subset of ubicomp devices, the available ubicomp services to consider is similarly constrained.

Ubicomp, by its very nature, must be able to accommodate seamless movement between different regulatory regimes. A ubicomp environment must ensure that services comply not only with local laws, but also provide an appropriate level of privacy support to the user wherever the local law is weak or nonexistent. Therefore the *regulatory regime* for a given jurisdiction provides the outermost layer in both the privacy protection and service constraint in a ubicomp environment. Lessig (1999) noted that privacy is dependent on four forces: law, market, norms, and architecture. The nature of ubicomp services is such that the architecture (the device in your hand) can remain constant, while the other factors may change depending upon where you are standing.

In this paper, we do not consider the infrastructure being used by a ubicomp device to take advantage of LBSs. Gunter et al. (2004) use the term *holders* to identify the principals in an infrastructure that collect location data or *sightings.* These sightings might be generated from a mobile telephone network using signal triangulation, a GPS tracking system, or accesses to a short range wireless network equipment (WiFi or Bluetooth) connected to the Internet. A *subscriber* is a system or service that uses data collected by holders. Although some of the scenarios and related work relies on a specific holder’s technical capability, we do not consider the detail of how a subscriber receives data.

In the context of this paper, private data refers to data in digital form. Langheinrich (2002) extends his model of privacy protection in ubicomp to non-digital sources such as CCTV cameras that use wireless privacy beacons to advise users when their privacy is at risk. Although we believe this is a worthy goal, we believe that the differences in regulatory regimes and the proliferation of dense CCTV coverage (especially in the UK) make addressing this issue impractical.

Many of the principles discussed here may also be applicable when the *user* is part of an organization. In this paper we restrict ourselves to the user who is a single identifiable human being. We regard organizational privacy as a security issue which can be regulated using contracts and agreements between institutions.

In the following sections we examine each of the layers of our model moving outward from the user. Section 2 examines types of personal data, both primary and derived, that are at risk in both conventional and ubicomp environments. Section 3 examines the next layer; we classify the types of ubicomp services a user may might request and illustrate these services using scenarios. The various regulatory regimes form the outermost layer in protecting the user and regulating available services; we analyze and classify these in section 4.

Section 5 examines existing models that attempt to tackle the problem of protecting privacy in ubicomp. We compare the provisions of these models and attempt to identify their shortcomings.

In section 6 we present an economics-based model. It assists a user in deciding which, if any, services to accept based on the appropriate regulatory regime, service, and type of data. We build on the economic utility model of Acquisti (2002). We discuss the use of privacy-protecting “noise” as an alternative to the release of personal information. We conclude by illustrating our model through a case study.

2 What is Personally Identifying Information?

Personally identifying information (private data) is often subjective. There is usually some amount of information whose access requires control by its owners (subjects); PII can range from the identity of an individual to that person's shopping habits. We use the term *attacker* to denote a person or organization who seeks to obtain PII without the consent of the owner. In order to consider what PII must be protected, we must first analyze the categories of data linked to an individual. Corby (2002) classifies private data into *static, dynamic, and derived* data. We present an extended version in Table 1.

Table 1: Taxonomy of Data Types and Examples based on (Corby, 2002)

Type of Data			Sub-Type & Example
Static	**Identity**	**Offline**	1. *Bio-identity*: fingerprints, race, colour, gender, height, weight, physical characteristics, retinal pattern, DNA 2. *Financial identity*: bank accounts, credit card numbers 3. *Legal identity*: government ID numbers (SSN, Passport #, Driver's Licence) 4. *Social identity*: membership in church, auto clubs, ethnicity 5. *Relationship*s: child of, parent of, spouse of 6. *Real Property Associations*: home address, business address
		Online	*Digital ID*: pseudonym, E-mail address, Username, IP address, Password
	Assets	**Tangible**	*Property*: buildings, automobiles, boats, mobile phones *Personal Worth*: credit balances, stock portfolios, debt balances
		Intangible	*Non-real property*: insurance policies, employee agreements
Dynamic	**Historical**		*Low Resolution: Transactions*: financial, travel, mobile phone records *High Resolution: UbiComp Sightings log (Time, Place)*
	Real-Time		*UbiComp Sightings ([Now], Place)*
Derived	**Analyzed**		**Data derived by analyzing trends over time** ***Financial behaviour*** 1. *Trends and changes*: month-to-month variance against baseline 2. *Perceived response to new offerings*: matched with experience ***Social behaviour*** *Behaviour statistics*: drug use, violations of law, family traits ***Tastes*** *Buying patterns*: purchase of item in a certain class suggests desire to buy other items in same class
	Composed		**Linking Data about person to other data** 1. *DNA analysis:* DNA linked to human genome database infers tendency to disease, psychological behaviour 2. *Multi-Data linking*: e.g. knowing a device with a given MAC address was seen at a given place/time and knowing that the number is registered to a person infers person was at place/time

As the table shows, ubicomp *sightings* occupy the *Dynamic* slot; adding one new data item composed of two parts: *timestamp* and *location*. This can be further divided by how data is used: either *real-time* (where the implied timestamp is “now”) or as a *historical* record. We note that dynamic/historical data is not a new privacy risk; it has been available through such mundane IT applications as credit card and telephone records. Ubicomp does, however, have the potential to provide far finer detail about one’s location with much greater temporal precision.

It should also be noted that ubicomp implicitly occupies parts of the *Derived* data category since analysis of location data over time can yield crucial PII to an attacker.

This classification motivates our examination of ubicomp services in the next section.

3 Classifying Ubicomp Services and Scenarios

Until recently, the lack of actual ubicomp services available to the general public has meant that much of the work in ubicomp privacy has used hypothetical scenarios analyzed as case studies. In this paper, we re-use some of the popular scenarios which represent the range of activities available to a ubicomp user of a device with an explicit user interface. We classify them according to the type of data and how the service affects the user. We only consider scenarios where there is a privacy risk from data processing taking place beyond the user’s control. Therefore we do not investigate ubicomp services achieved entirely by computation on the user’s device.

Gunter et al. (2004) present four scenarios similar to those found in other work: *FriendsInTown.com, Market Models, What’s Here?,* and *Travel Archive*.

1. *FriendsInTown.com* is an alerting service allowing two people to register an interest in being notified when they are close to one another. As soon as the criterion is satisfied both users are informed. Similar scenarios proposed in other work also involve being interrupted by a ubicomp device once a location-based criterion is satisfied. These might include advertising notifications where a user is alerted as they approach a product on sale, or a form of semi-automated check-in as one enters an airport.

2. *Market Models* provides historical information about characteristics of a group of users who satisfy a certain time/space criterion; such as the average income of everyone at Penn Station at noon on a given day.

3. *What's Here?* is typical of services which provide more detail to a user in response to a request about their present location. Examples include a list of forthcoming events in a building, tourist points of interest (e.g. (Hong & Landay, 2004) among others), or the route to the nearest sushi restaurant (Duckham & Kulik, to appear).

4. *Travel Archive* keeps a record of the datestamps and locations of people in order to answer queries like "where was I this time last year?" or "How many sales people did we have in the Birmingham area on Tuesday?"

According to the data breakdown in Table 1 in the *Dynamic* section it is clear that *FriendsInTown.com* and *WhatsHere?* are both examples of Real-Time data, while *MarketModelsI* and *Travel Archive* rely on historical data. Ubicomp does not bring many new issues with respect to Dynamic Historical data other than the possible increased resolution of sightings. Access to and analysis of the data does not require a ubicomp device. For the Real-Time scenarios, there are clearly two types of service: *Interupt-Based*, where the user is alerted once certain criteria are satisfied, and *Query-Based*, where the user asks for information based on their current location.

4 Regulatory Regimes

"After a while you learn that privacy is something you can sell, but you can't buy it back." – Bob Dylan (Dylan, 2004)

American legal commentators began to consider privacy ("the right to be let alone" (Warren & Brandeis, 1985)) as a "natural law" or residual right in the late 19th Century. Their discussions were prompted by the rise of the newspaper industry which had been invigorated by the widespread use of photography. Their consensus was that the right to privacy had always been there but never formally incorporated in statute. Later Supreme Court decisions would suggest that the 9th and to some extent 3rd, 4th, and 5th amendments to the United States Constitution provided personal privacy protection.

On the other side of the Atlantic, Article 8 of the European Convention on Human Rights and Fundamental Freedoms (ECHR) (1950) explicitly states that everyone has a right to privacy in private and family life (subject to some restrictions).

In the mid 20th Century IBM's Hollerith punch card technology was used to collect census data which was later used by the Nazis to identify Jews for transport to extermination camps (Black, 2001). In the post-war era European countries codified strict privacy protection using both international treaties and national legislation. Most Western countries have followed suit, following OECD Guidelines (OECD, 1980) which are often cited as Fair Information Practices (FIP). The United States is an exception to the strong legal protection of personal privacy; instead having strong guidelines, it relies on a patchwork of laws covering, among other things, children's use of websites, video tape rental records, health insurance and financial data.

As we will see in the next section, the different approaches taken between Europe and the US mirrors two of the approaches for the protection of privacy in ubicomp. The European approach has been to consider PII alongside intellectual property; historically in the US most PII has been considered to be in the public domain once it has been revealed to one institution.

Consider a situation in which an individual reveals a postal address to a business to process a specific request. The default position in the EU is that any other use is implicitly forbidden. In the US and other less-restrictive regimes, one institution can sell mailing lists to another without obtaining the permission of the people on the list. Such lists can be sold and re-sold many times over, including composed data from spending patterns. This problem is probably what Bob Dylan had in mind in the quote at the beginning of this section.

The ECHR was one of the first 'Bill of Rights' style documents to explicitly mention privacy as a fundamental human right. Several of the larger European countries were early adopters of the OECD guidelines on privacy which effectively influenced the development of European Community law on data protection and privacy.

The ECHR is only enforceable against signatory governments; two pieces of EC legislation extend privacy protection to cover individuals and non-governmental organisations.

The first is Directive 95/46/EC (1995) which ensures that users have access to all of the data held about them; that data is only collected with the individual's explicit consent, and that it is destroyed when it is no longer needed for the original purpose.

The directive has possible consequences for location-aware computing. For example, as a user enters an area offering a service to which they would like to subscribe, does the user have to give explicit permission for the use of personally identifiable data for each new instance of the service? It is possible that the law may protect users, but is insufficiently flexible to allow them to effectively utilise the advantages of a technology.

Fortunately, recent European law is anticipating some measure of technological change. The recent Directive 2002/58/EC (2002) is aimed at extending Directive 95/46/EC to the telecommunications sector and makes explicit mention of location-aware technology. Although the drafters of this directive were considering second- and third-generation mobile telephones, the directive prohibits the use of location information without explicit informed consent. Furthermore, the directive requires that equipment and service providers offer a simple free-of-charge method for users to temporarily hide their location information. This legislation also controls the use of cookies in web browsers. European privacy laws attempt to implement a kind of ‘transitive closure’ whereby data may only be exported to another country possessing an equal data protection regime, or where the exporter has a special data protection contract with the importer providing equivalent protection to the directive.

Japan is one of the countries with the greatest take-up of consumer-level ubiquitous computing (in the form of location-aware mobile telephones). It was one of the earliest countries to define privacy regulations for ubiquitous computing. This early level of market certainty resulted in increased business confidence and thus a wide proliferation of services. Similarly, thanks to well-established regulations consumer confidence in the new services was probably higher than it would be in a completely unregulated arena.

Canada and Australia have also instituted strong privacy laws although without explicit attention to location-aware computing. Like the EU and Japan, each have Information/Privacy commissioners with the power to take both punitive and retributive action against privacy violations.

In the US, a patchwork of legislation at both the state and national level provides privacy protection in certain narrow domains, including websites aimed at children (Children's Online Privacy Protection Act, 1998), financial sites (Gramm-Leach-Bliley Act, 1999), health insurance sites (Health Insurance Portability and Accountability Act, 1996), and certain baffling collections of data such as the records of videotape rentals (Video Privacy Protection Act, 1988). The data processing industry has provided heavy resistance to any form of privacy regulation; with self-regulation (e.g. TRUSTe (2004)) being promoted as an alternative, with virtually no mechanisms for redress of violations.

Unlike other Western countries, America does not possess a comprehensive national data protection law, and the closest equivalent to a national privacy commissioner is the Federal Trade Commission (FTC). The FTC can take action against a business that violates its posted privacy policy under unfair trading, but such violations are difficult to prove and the FTC has only acted in a very small number of cases. The most notable case was against GeoCities in 1998 (FTC, 1998) for misrepresenting the purpose for which it was collecting data from both adults and children. Despite several high profile violations of the TRUSTe standards by Microsoft and Lotus, their TRUSTe certificate has never been revoked. Given the weak standards set for simple online privacy protection, there is no immediate prospect of legislation in the US either affording any privacy protection or impediment for location-aware computing. However, the regulations requiring mobile telephone networks to provide location information to emergency services (E-911 in the US, E-112 in Europe) are likely to affect how privacy enhancing technologies can be applied. Table 2 shows the four layer model incorporating the four broad regulatory regimes identified.

Ubicomp services obviously need to be aware of the current regulatory regime so that they can comply with it. Certain regimes require very explicit notice and consent. This will constrain how services are delivered. Users need to know that the level of protection they require personally will be maintained as they cross regulatory borders, some of which will be invisible. The very nature of ubicomp suggests that moving between regulatory regimes will be a common enough occurrence that this requirement must be supported. The complexity of the legal differences between regimes is such that the user should neither be expected to understand them, nor keep up with them as regulations change. We suggest in our model (section 6) that an understanding of the relevant regulatory regime be coded into a privacy protecting proxy; users need only express their own privacy policy for the appropriate action to be taken in a given regulatory regime. In the next section we examine how others have approached the automation of privacy protection in ubicomp.

Table 2: Four Layer Model of UbiComp Services

<table>
<tr><th colspan="6">User</th></tr>
<tr><th colspan="6">Data</th></tr>
<tr><td rowspan="2">Static</td><td colspan="3">Dynamic</td><td colspan="2">Derived</td></tr>
<tr><td>Historical</td><td colspan="2">Real-time</td><td>Analysed</td><td>Composed</td></tr>
<tr><td rowspan="2"></td><th colspan="3">UbiComp Services</th><td colspan="2" rowspan="2"></td></tr>
<tr><td colspan="2">Query</td><td>Interrupt</td></tr>
<tr><th colspan="6">Regulatory Regimes</th></tr>
<tr><td colspan="6">States with little to no privacy protection in law</td></tr>
<tr><td colspan="6">States with some protection (e.g. USA)</td></tr>
<tr><td colspan="6">States with strong privacy protection (e.g Canada, Australia)</td></tr>
<tr><td colspan="6">States with strong privacy protection including location aware (e.g. EU, Japan)</td></tr>
</table>

5 Related Work

Before considering attempts to preserve user privacy in ubicomp, we first consider the simpler problem of privacy preserving mechanisms in traditional desktop computing. A common privacy risk in desktop computing is through the unintentional revelation of PII through a web browser. The only built-in protection for users in most web browsers is through restricting the automatic acceptance of cookies. Website privacy policies are written in natural language, making it difficult to perform automatic analysis of compliance with an individual's privacy policy. Some attempts have been made to codify site privacy policies using XML to perform some measure of automatic analysis. The Platform for Privacy Preferences Project (P3P) was developed by World Wide Web Consortium (W3C) to integrate machine-readable privacy policies into web browsers (Cranor, 2002). P3P enables web browsers to automatically read privacy policies of web sites possessing appropriate XML tags; not all browsers are able to parse these tags and most websites do not post P3P polices. The AT&T Privacy Bird (AT&T, 2003) is an example of a browser plug-in that automatically compares a website's P3P policy with the user's own privacy preferences; it indicates green for a match, red for non-match, and yellow when no P3P policy is present. P3P version 1.0 has been criticized for its lack of enforceability, lack of relationship to existing legislation, and for failing to reflect Fair Information Policies (Electronic Privacy Information Center, 2000). P3P makes an assumption that companies own the data collected from visitors and make non-binding promises about how it will be used.

The W3C has proposed a P3P Preference Exchange Language called APPEL so that users can own sets of policies for different situations and collect sets of complex policies from databases of trusted third parties (Cranor et al., 2002). Criticisms of P3P aside, the direction of this work is important because it acknowledges that individuals may require complex sets of privacy preferences covering a wide range of situations. Most people will rely on trusted third parties, such as consumer organizations, to suggest policy sets appropriate for them, a notion supported by Yee and Korba (2005) in their work on semi-automatic policy derivation and matching.

Another promising development in the automated analysis of privacy requirements is IBM's Enterprise Privacy Application Language (EPAL) (2003). EPAL is much more finely-grained than P3P and therefore has the potential to address some of P3P's shortcomings. P3P is designed to present an enterprise's very general privacy policy in machine-readable form to the outside world, whilst EPAL is designed to allow enterprise-internal relationships to be formalized and enforced.

5.1 Privacy Models in Ubiquitous Computing

Previous work aimed at helping ubicomp users protect their privacy, which generally means their location privacy, can be divided roughly into two groups:

1. *policy matching*: attempts to provide mechanisms for comparing a user's policy to that of the ubicomp service and notifies the user of mismatches, and;
2. *noise*: tries to hide or disguise a user's location or identity.

We divide noise into five types:

i. *anonimizing*: hiding the identify of the user;
ii. *hashing*: disguising the identify of the user
iii. *cloaking*: making the user invisible;
iv. *blurring*: decreasing the accuracy of the location (and possibly time); and
v. lying: giving intentionally false information about location or time.

The Internet Engineering Task Force Working Group on Geoprivacy recently released an Internet-Draft (Schulzrinne et al., 2004) defining an XML schema for rules that match user requirements to geo-location requests. The current draft supports policy matching through a rich set of rules that permit users to grant or deny access to their location information. The schema also supports noise in the form of blurring, by permitting a user to specify the resolution of their location information. However, it does not support any other forms of noise.

The complex nature of the XML schema underlines the suggestion of Yee and Korba above that effective user privacy policies can be extremely complicated and will require a great deal of support if they are to be at all manageable by consumers.

Support for user control over personal privacy policies is provided by Lederer et al. (2002). Here, the authors note that ubicomp users may need different personal privacy policies at the same time depending on the recipient of the data. They use the metaphor of *situational faces* to allow a user to show an anonymous "face", for example to retailers, while at the same time showing their "public" face to close friends (thus allowing a scenario like *FriendsInTown.com* to work).

However, the problem of defining each of the complex ubicomp privacy policies still remains.

Jiang et al. (2002) use an economics-based approach to analyze information flow in ubicomp. They have developed a model called *approximate information flow*. The model proposes a number of abstractions which try to minimize any imbalance between those who release their data and those who collect it. From an end-user perspective one of these abstractions classifies the methods of preserving privacy as *prevention, avoidance,* and *detection*. *Prevention* means not releasing PII if it could be misused, *avoidance* permits the release of PII but takes steps to try to prevent misuse; whilst *detection* is the process of sensing when misuse has occurred. We also use these terms when looking at related work, and summarize our analysis in Table 3.

One example of a commercial ubicomp system is AT&T's *Find People Nearby* service (AT&T, 2004) which uses the conventional GSM/GPRS mobile telephone network. It allows users to register friends they would like to locate and obtains consent from each of those individuals. Once consent is obtained, the user can send a query which returns the location of a friend. A registered user may elect to flag themselves as *unfindable* or *findable* to others. This is a real-time query ubicomp system employing cloaking for privacy protection as a preventative measure. Many similar network-independent services are available in Europe. European law requiring notification and consent constrains the interface; users must send a text message each time they wish to turn tracking on or off.

These systems illustrate a common concern about privacy problems in ubicomp: the departure from social norms. This information asymmetry was noted by Jiang et al. (2002); one person is allowed to know the location of another without the second person knowing that their personal information is being passed on. This is in contrast to a face-to-face interaction in which each person can see that they are being observed by the other.

Hong and Landay (2004) identify a number of privacy requirements for end users, including *simple and appropriate control and feedback.* They address this concern in their *Confab* architecture by adding digitally signed privacy tags to shared data items. Privacy tags contain attributes such as *TimeToLive* (specifying retention time), *MaxNumSightings* (how much history should be kept), and *Notify* (allowing the data owner to know who has been looking at their information). In the event that the retention time is exceeded, or if data is disclosed without permission, or if the tag's digital signature is invalid; then data can automatically be deleted or marked unreadable by the clients of an individual's peers.

Confab uses a privacy proxy to handle data requests and manage the user's privacy policy so the actual ubicomp client is insulated. The *Notify* field supports the feedback requirement; it is possible for a data subject to know who has been looking at their data and how often. This important feedback element was also identified by Nguyen and Mynatt (2002) in their Privacy Mirrors system..

Hong and Landay identify another end-user requirement in ubicomp privacy: *plausible deniability*, or, in plain English: 'lying'. There are many situations where people rely on "white lies" or benign deception to avoid social embarrassment or simply to surprise a loved one. In one study, 88% of respondents said that they believed it was acceptable to deceive a person if it was in that person's best interest (Sokol, 2004). Hong and Landay's Confab satisfies these desires by returning "Unknown" when a data request violates a user's set privacy policy or by returning a preset value if the user wishes to lie.

This approach corresponds to the European data protection model of data being licensed for a specific purpose and no other. The idea of combining data with metadata in Confab is similar to one form of Digital Rights Management (DRM) approach where music playback software enforces the number of licensed devices on which a piece of music may be played. Langheinrich (2002) proposed using metadata in his privacy awareness system (pawS). Like Confab, pawS makes use of a privacy proxy. It matches user privacy policies with those advertised by ubicomp services using P3P and APPEL and allows the user to accept or decline if there is a mismatch.

The DRM approach to privacy is typified by the work of Gunter et al. (2004) who combined a method using a formal access control matrix with Personal DRM (PDRM). Their PDRM system combines the features of P3P with the eXtensible rights Markup Language (XrML) (ContentGuard.com, 2005) to create digitally signed contracts licensing the use of personal data for specific purposes and for fixes periods of time. PDRM uses a geographic information server to enforce contracts in much the same way as Confab and pawS use a proxy to hide the real ubicomp user. PDRM requires prospective subscribers to submit digitally-signed privacy policies which are compared with individual users' policies and either accepted or rejected on an as-needed basis.

Other approaches to protecting privacy have focussed on using 'noise' to protect location information. Gruteser and Grunwald (2003) expand the uncertainty of the location of a single user to a cover an area that includes a number of other users, thus making them anonymous within the group. Duckham and Kulik (to appear) give a false, but nearby location, instead of the actual location of the user. Beresford and Stajano (2003) show how the identity of a user can be protected by hashing it to a frequently re-named pseudonym using a proxy.

Each of these methods is designed to balance the need between privacy protection and the quality of service provided to the user. Each is intended to prevent the subscriber from gathering too much private information about a subject, and to prevent an attacker from gleaning sufficient information to track subjects without their knowledge or consent.

Table 3 shows a comparison of the major ubicomp privacy models against our framework.

Table 3: Comparison of Privacy Protecting Models in UbiComp

Author(s)/ System Name	*Description*	*Type of privacy protection*	*Real-time*	*Historical*	*Method of protecting privacy*
1. (Duckham & Kulik, to appear)	Location blurring to nearby point	Preventive	X		Noise (Blurring)
2. (Gruteser & Grunwald, 2003)	k-anonymity	Preventive		X	Noise (Blurring/Anonymity)
3. (Beresford & Stajano, 2003)	Provides unlinkability between pseudonyms	Preventive		X	Noise (hashing)
4. (Hong & Landay, 2004) **Confab**	Privacy proxy handles digitally signed privacy metadata	Avoidance, Preventative	X	X	Matching Policies, Noise(Cloaking, Lying)
5. (Langheinrich, 2002) **Privacy Awareness System (pawS)**	Use of : • privacy proxy • privacy-aware database	Avoidance, Preventive	X	X	Matching policies
6. (Gunter et al., 2004) **AdLoc**	Combining formal access control with PDRM	Avoidance, Preventive	X	X	Matching policies/access control
7. (Jiang et al., 2002)	Model: Approximate Information Flows The Principle of Minimum Asymmetry	Prevention, Avoidance & Detection	X	X	
8. (Lederer et al., 2002)	UI Metaphor: Situational faces metaphor – conceptualising end-user privacy preferences	Preventive	X	X	
9. (AT&T, 2004) **Find People Nearby**	Friend finding application	Preventive	X		Noise (cloaking)
10. (Nguyen & Mynatt, 2002) **Privacy Mirror**	UI Metaphor: Privacy Interface (for feedback and detection)	Detection		X	

5.2 The Economics of Privacy

The systems discussed in the previous section use the techniques of *prevention* (refusing to use services that will release PII the user does not wish to release) and *avoidance* (using noise to minimize the risk of actual PII being released). In the case of a service which requires more PII than the user is willing to reveal, the service will be rejected by the privacy proxy. The ultra-paranoid user who chooses to reveal no PII to anyone will find few if any services available, making the ubicomp device almost useless to that user.

What is missing is a tool that helps a user analyze potential risks and balances them against their needs.

Economists studied privacy for some time (Posner, 1978) and have expanded the relatively simple concept that privacy protection represents a trade-off between the benefits of sharing PII and its associated costs. In terms of ubicomp, the benefit from releasing one's current location or other PII, is the receipt of a service. The value of the benefit may be outweighed by the present or future cost of unknown "attackers" being able to track you.

A Privacy Enhancing Technology (PET) will have some cost (monetary, in functionality, or perhaps a lower quality of service in the case of blurring) which must be subtracted from the benefit received from the service. The lack of a service may also have a cost in terms of convenience or necessity (e.g. if you need cash urgently and need to find the nearest ATM).

Many trade-offs will be less clear; for example, if you allow a merchant to see your buying patterns then it can send you highly targeted ads or offers, thus reducing the amount of irrelevant material you have to process. The merchant could also use this information to your detriment: if your shopping patterns show what you are willing to pay for certain items, a merchant may charge you a higher price because it knows you are likely to pay it (Rosencrance, 2000). Acquisti (2002) shows that this course of action is not in the best interests of the merchant, but it is impossible for users to know if they are the victim of discriminatory policies.

Acquisti's (2004b) analysis of consumer behaviour indicates that consumers are unlikely to act rationally (in a privacy sense); self-proclaimed privacy advocates are prepared to give up personal information for relatively small rewards. He shows how, with the economics of immediate gratification, even sophisticated users become "privacy myopic." Pre-set privacy policies can help prevent privacy myopia, but there is a clear need for tools to help users come to rational decisions about privacy.

6 A Practical Model for User Control of Privacy

In our four layer model of ubicomp privacy issues; the outermost layer represents the regulatory regime whilst the innermost layer is composed of users and their personal privacy policies (see Table 2). We assume that the user's policy is coded (perhaps in geopriv’s XML schema (Schulzrinne et al., 2004) or some other method), and has been defined for a variety of recipients (perhaps using a *faces* metaphor as suggested by Lederer, Markhoff, and Dey (2002)).

We assume that the ubicomp device gathers its own location information by some means (such as connecting to a network provider or from an integral GPS or some combination of methods). Location data is transmitted to a trusted privacy proxy. As with the Confab and pawS systems, the proxy handles all requests from subscribers and has access to each user’s current privacy policies. Figure 1 shows the layout of the model.

We chose to extend Hong and Landay’s (2004) Confab architecture in our model for a number of reasons:

1. using a proxy allows a broad range of noise to be employed and removes computational load from the ubicomp device;
2. since the proxy knows the user’s current regulatory regime it can:
    - apply appropriate regulations when accessing services such as notice and consent on behalf of the user;
    - balance the current regulatory protection with the user’s personal policy; if the former is stricter than the latter then the user can access services directly. If the user has a more strict policy than the regulatory regime, then the proxy will have to apply appropriate techniques before the user can access services;
3. digitally signed metadata attached to PII allows a broad range of enforcement techniques, including a community of peers;
4. including a notify tag in the metadata permits the enforcement of user feedback requirements (knowing you are being watched and discouraging overzealous spying).

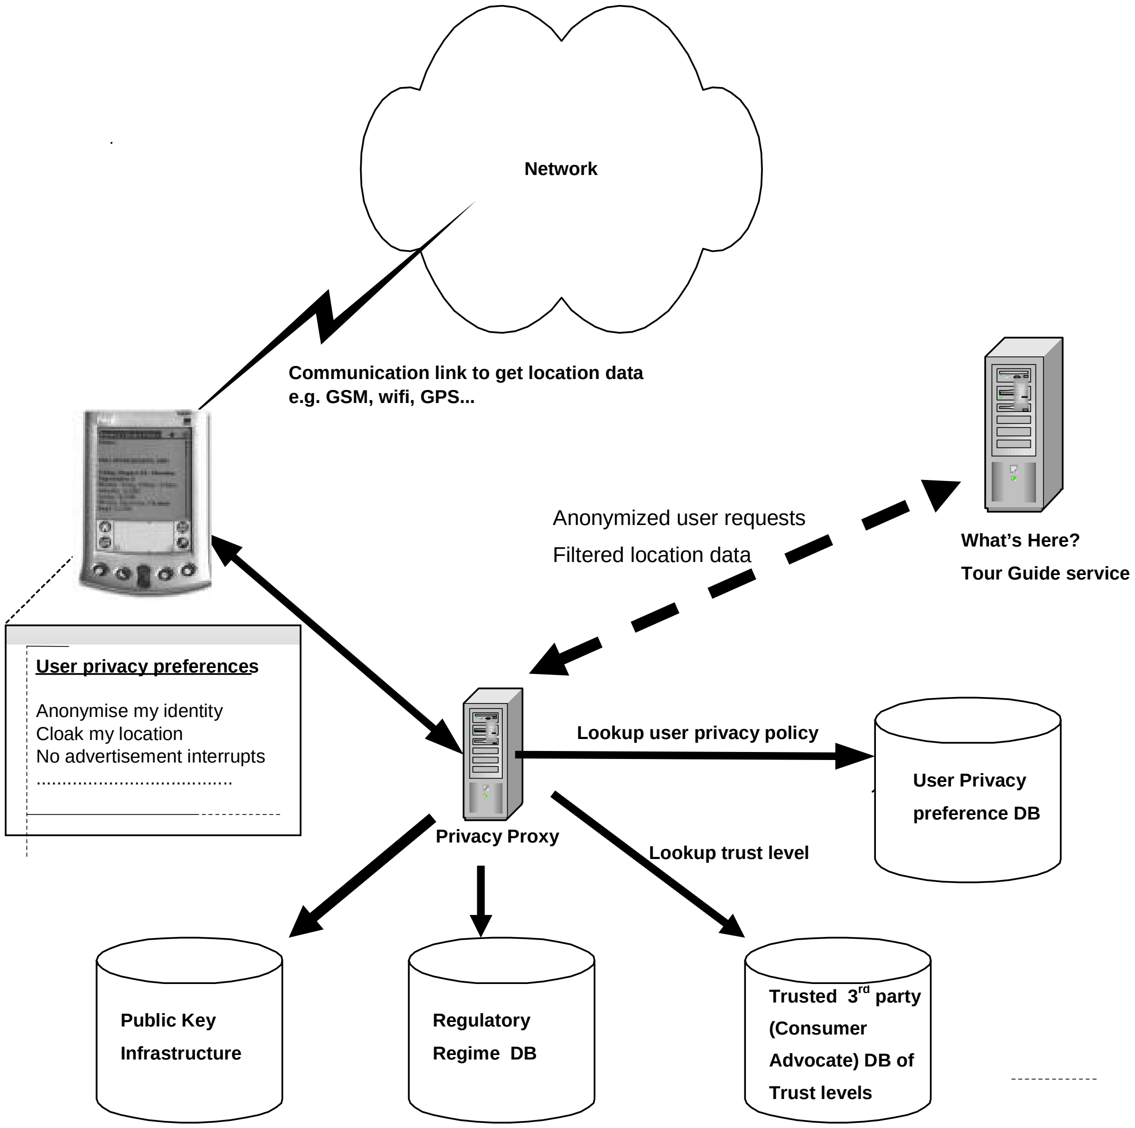

Figure 1: Outline of Interactions in Model for User Control of Privacy

In our model, the proxy not only acts on behalf of the user in sending (or not) location information to a subscriber, but it also acts on behalf of users when they access the features of the service. Some regulatory regimes have explicit requirements for how notice and consent is given (if at all). Since the proxy always knows the location of the user, it is able to apply the appropriate regulations and ensure interface compliance. The proxy's knowledge of the local regulations also allows it to compare the user's policy with local regulatory protection and either rely on this or provide additional protection through noise as necessary.

This proxy permits all five types of privacy-protecting noise to be applied in situations where the user does not wish to be interrupted by certain classes of person or organisation. In particular, it allows a user to lie to other subscribers according to their policy settings whilst still complying with local regulations.

We also adopt the PDRM approach of Gunter et al. (2004) which creates digitally-signed licenses or contracts for the use of data wherever possible (if the service provider allows and a public key infrastructure is present). In regulatory regimes lacking strong legal protection for privacy the user still has enforceable civil redress against privacy theft in the same manner that music companies have redress for copyright violations. This is compatible with Lessig's view of the influence of laws and norms on privacy, and mirrors his use of copyright law to license media in the Creative Commons (2005).

The final element of our model provides users with the tools necessary to adjust their privacy level in a rational way in the event of a conflict between their privacy policy, a regulation, and a required service. Our model incorporates Acquisti's (2004b) utility model for measuring the potential benefit of the release of PII against the possible costs. Acquisti's utility equation is a complex function of five variables, some of which are composed of multiple factors and some of which are probabilities (for example; data misuse). In our analysis above we noted that users cannot know in advance if a merchant will use their PII to enhance the user's experience or use that information against him in non-competitive pricing. In order to measure the probability of this occurrence, we incorporate a third-party database of *trust in organizations*. This could be provided by independent consumer advocates who are able to regulate a merchant's trust rating based on consumer reports and from their own investigations.

This model provides a tool for making rational decisions based on actual versus perceived risk. It would prove particularly useful in situations where consumers requires immediate gratification, or where they need to decide whether or not to relax their privacy constraints to receive a service.

In the next section we present a case study to illustrate our model.

7 Case Study

Section 3 introduced a number of typical ubicomp scenarios. Here we illustrate our model by following the travels of an imaginary ubicomp user, Bob, through these scenarios.

Bob has programmed his ubicomp device to upload his location to a trusted privacy proxy server at five minute intervals. This traffic is encrypted, so even if his ex-girlfriend Eve, was listening to network traffic she would be unable to decrypt his location data. Bob stores a number of privacy policies on the privacy proxy. These policies relate to individuals, classes of individuals and organizations. Many of these policies are triggered by his location and the time of day.

For example, during the working day Bob's policy provides location data to his partner, Alice, his work colleagues and his children's school. Each of them can send an explicit request to a service like *FriendsInTown.com* provided Bob has an account with the company and has previously authorized them to have access. When a request is sent from *FriendsInTown* to Bob's privacy proxy, the proxy applies a policy that is appropriate for the time of day and requestor.

Assuming the proxy approves the issue of data, Bob's information is tagged with metadata indicating an appropriate retention time. The data is then encrypted and transmitted to *FriendsInTown*. The entire transaction is then logged by the proxy for later examination and for legal purposes.

An attacker (stalker) may gain some measure of access to Bob's location data by stealing a private key belonging to one of Bob's friends. They could then make repeated requests to build up a profile of his movements. Bob would be informed of this when his proxy reports that a friend is taking an overkeen interest in his movements.

Bob is partially protected from accidental or intentional re-forwarding of his location information by an authorized recipient. Suppose Bob's daughter has been taken ill; the temporary secretary at his daughter's school sent a location request to find Bob and subsequently accidentally forwarded the data to a third party. The signed metadata would indicate that the data had expired and that the unauthorized recipient was not on the original recipient list. The final recipient's computer should either automatically delete the data or at least refuse to read it (in the same way that one person's purchased digital music cannot be played on another person's player).

Bob has control over the location data issued by his proxy, therefore he has roughly the same ability to commit benign deceptions as he did before his movement was monitored. By instructing the proxy to utilize a noise effect (such as blurring) he could choose to take an extra long lunch rather than visit a nearby client. Even if his boss used a *TravelArchive* service to look at Bob's location history it would simply indicate he was in the area. Bob might explicitly lie about his location if he wishes to surprise Alice and didn't want her to know he had been in a jewellery store. As with conventional deception there are risks, but anecdotal evidence as well as some ethics research (Sokol, 2004) indicates that people must be able to lie at times.

Most of Bob's privacy needs can be satisfied by a set of predefined policies that are activated by the time of day or his location; so, other than for secret trips to the jewellery store, he does not need to change his privacy profile.

When the work day is over, Bob's colleagues no longer have access to his location data but close friends might automatically be granted access. Bob may want to be advised when GadgetsRUs have a sale on accessories for his ubicomp widget. He can enable certain advertising interrupts that will be activated when he visits a shopping mall.

Upon entering the shopping mall he might be informed that *MarketModels* would like to collect information about his movement around the mall, in exchange they will offer him discounts at certain stores. How would he know if this would be worth doing?

Bob can ask his privacy proxy to make an assessment of the costs and benefits of completing the *MarketModels* survey. It will apply an economic utility model to Bob's situation. For example, the proxy's utility model might calculate that Bob was likely to save an additional $100 over the next few if he used the discounts. The proxy must then offset these savings against the risks involved; *MarketModels* privacy policy claims they will anonymize Bob's data after collection. The proxy then checks *MarketModels*' entry in the Online Consumers Association database to determine the probability of them honouring their policy. Finally, the proxy offsets the relevant risk calculations against the projected savings. With all of this information, the proxy can give informed advice to Bob; either that he should accept the offer and benefit from the projected savings, or, that he should decline since the risk from *MarketModels* exceeds Bob's comfort level.

Bob enjoys adventure holidays and decides to take his vacation in the recently democratized Republic of Elbonia, which claims to host eight of the seven wonders of the world, but has very little formal privacy legislation. The only travel advice application available to his ubicomp device is *ElboniaNow*, an equivalent to the *WhatsHere?* tourist advice application. Bob's location and other data are not at risk; his privacy proxy recognizes Elbonia's lax privacy regulation and restricts itself to sending anonymous information to *ElboniaNow*.

If Bob has to provide additional PII to take advantage of another service in Elbonia (even if the economic utility model advises him against it), then he will have some protection from ordinary civil contract law if his personal data was sent with a PDRM license attached to it.

8 Conclusions and Future Work

Many surveys have demonstrated consumer concerns about privacy in ordinary desktop computing. Ubicomp promises to bring many consumer benefits, but it magnifies existing privacy concerns. Widespread adoption at a societal level will require strict attention to the societal forces acting on privacy. Lessig (1998) sees these as Laws, Norms, Market, and Architecture. Our model addresses each of these factors. We encode relevant laws so that the privacy proxy can manage both protection and compliance. We address societal norms in two ways:

- by allowing five forms of privacy-protecting noise, we allow people to control when and how they are visible to others, and by supporting lying we ensure that existing societal behaviour patterns are not disallowed by technology; and
- by ensuring that access to location data is logged (so users can see who is watching them), we support feedback to correct the asymmetry introduced by making such location data available.

Finally, our model allows market forces to be applied to privacy by ensuring that services offering low privacy protection are not chosen by the proxy. If enough users adopt privacy-friendly policies, this should force the market to shift in favour of those offering privacy-protected services.

We have also indicated how Acquisti's utility model can be used to aid users in determining the net benefits involving certain privacy-sensitive decisions. However, because privacy intrusions are very specific and very much context-dependent, Acquisti's abstract model has to be calibrated for specific scenarios within our model (Acquisti, 2004a). We are currently investigating specific trade-offs for specific scenarios.

Our model has a number of technical challenges. One is an XML schema to encode the protection and compliance features of various regulations, which we are now working on. Another is the absence of a widely-deployed public key infrastructure to support signing and encryption of data and metadata. To address this second problem, we are looking at incorporating a mixed approach to allow encryption and signing support where possible.

An important issue that we are not addressing is the interface for user selection and control of complex privacy policies which is clearly crucial to consumer protection. Ease of use for this kind of interface is crucial if adequate consumer protection is to be achieved. One solution will be for trusted third parties to design generic privacy solutions and pre-package those in such a way as to make it easy for users to select a policy set appropriate for them. Another would be to provide a questionnaire-based tool to help users determine their privacy needs.

In this paper we have highlighted the additional privacy risks in ubicomp. We analyzed previous work in the context of our characterization of the data, services, and regulations affecting user privacy. From the analysis we combined and extended previous approaches in order to address the new privacy needs in a flexible and comprehensive way. We illustrated the model with a case study and noted where further work is needed to realize an individual user-centred solution to privacy protection in ubicomp.

Ultimately, ubicomp take-up by users depends on privacy protection being both trusted and usable. By providing a model that protects users from the invasiveness of both the technology and other people we believe we have taken an important step along this road.

9 References

Ackerman, M., & Cranor, L. (1999). *Privacy Critics: UI Components to Safeguard Users' Privacy.* Paper presented at the CHI'99, 258-259.

Acquisti, A. (2002). *Protecting Privacy with Economics: Economic Incentives for Preventive Technologies in Ubiquitous Computing Environments.* Paper presented at the Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing, UbiComp 2002.

Acquisti, A. (2004a). Personal Communication.

Acquisti, A. (2004b). *Privacy in electronic commerce and the economics of immediate gratification.* Paper presented at the 5th ACM conference on Electronic commerce, 21-29.

AT&T. (2003). *Privacy Bird*, available from: http://www.privacybird.com

AT&T. (2004). *Find People Nearby*. Retrieved 31 January, 2005, available from: http://www.attwireless.com/personal/features/organization/findfriends.jhtml

Bellotti, V., & Sellen, A. (1993). *A. Design for Privacy in Ubiquitous Computing Environments.* Paper presented at the 3rd European Conf. on Computer Supported Cooperative Work, (ECSCW 93), 77-92.

Beresford, A. R., & Stajano, F. (2003). Location Privacy in Pervasive Computing. *IEEE Pervasive Computing, 2*(1), 46-55.

Black, E. (2001). *IBM and the Holocaust: The Strategic Alliance Between Nazi Germany and America's Most Powerful Corporation*. New York: Crown.

Children's Online Privacy Protection Act (1998), 15 USC 6501-6505. Available from: http://www.ftc.gov/ogc/coppa1.htm

ContentGuard.com. (2005). *XrML Version 2.0*, available from: www.xrml.org

Convention for the Protection of Human Rights and Fundamental Freedoms (1950). Available from: http://conventions.coe.int/treaty/en/Treaties/Html/005.htm

Corby, J. M. (2002). The Case for Privacy. *Information systems security, 11*(2), 9 - 14.

Cranor, L. (2002). *Web Privacy with P3P*. Cambridge, MA: O'Reilly & Associates.

Cranor, L., Langheinrich, M., & Marchiori, M. (2002). *A P3P Preference Exchange Language*, available from: http://www.w3.org/TR/P3P-preferences/

CreativeCommons.org. (2005). *Creative Commons*, available from: http://creativecommons.org/

Dawson, L., Minocha, S., & Petre, M. (2003). *Social and Cultural Obstacles to the (B2C) E-Commerce Experience.* Paper presented at the People and Computers XVII - Designing for Society, 225-241.

Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data (1995), 95/46/EC. Available from: http://europa.eu.int/comm/internal_market/privacy/law_en.htm

Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (2002). Available from: http://europa.eu.int/comm/internal_market/privacy/law_en.htm

Duckham, M., & Kulik, L. (to appear). *A formal model of obfuscation and negotiation for location privacy.* Paper presented at the Pervasive 2005.

Dylan, B. (2004). *Chronicles: Volume One*. New York: Simon & Schuster.

Electronic Privacy Information Center. (2000). *Pretty Poor Privacy: An Assessment of P3P and Internet Privacy*, available from: http://www.epic.org/reports/prettypoorprivacy.html

FTC. (1998). *Internet Site Agrees to Settle FTC Charges of Deceptively Collecting Personal Information in Agency's First Internet Privacy Case*, 2004, available from: http://www.ftc.gov/opa/1998/08/geocitie.htm

Gramm-Leach-Bliley Act (1999), 15 USC, Subchapter I, Sec. 6801-6809. Available from: http://www.ftc.gov/privacy/glbact/glbsub1.htm

Gruteser, M., & Grunwald, D. (2003, May). *Anonymous usage of Location-Based Services Through Spatial and Temporal Cloaking.* Paper presented at the First International Conference on Mobile Systems, Applications, and Services.

Gunter, C. A., May, M. J., & Stubblebine, S. G. (2004). *A Formal Privacy System and its Application to Location Based Services.* Paper presented at the Privacy Enhancing Technologies.

Health Insurance Portability and Accountability Act (1996), 42 USC 201. Available from: http://aspe.hhs.gov/admnsimp/pl104191.htm

Hong, J. I., & Landay, J. A. (2004). *An Architecture for Privacy-Sensitive Ubiquitous Computing.* Paper presented at the Proceedings of the 2nd international conference on Mobile systems, applications, and services, Boston, MA, USA, 177 - 189.

IBM. (2003). *Enterprise Privacy Authorization Language*, available from: http://www.zurich.ibm.com/security/enterprise-privacy/epal/Specification/index.html

Jiang, X., Hong, J. I., & Landay, J. A. (2002). *Approximate Information Flows: Socially-based Modeling of Privacy in Ubiquitous Computing.* Paper presented at the Fourth International Conference on Ubiquitous Computing, Goteberg, Sweden.

Langheinrich, M. (2002). *A Privacy Awareness System for Ubiquitous Computing Environments.* Paper presented at the 4th International Conference on Ubiquitous Computing (Ubicomp 2002), 237-245.

Lederer, S. (2004, April 1, 2004). *Background Readings*, available from: http://guir.berkeley.edu/groups/privacy/readings.html

Lederer, S., Dey, A. K., & Mankoff, J. (2002). *A Conceptual Model and a Metaphor of Everyday Privacy in Ubiquitous Computing Environments* (Technical Report UCB/CSD-2-1188): Computer Science Division, University of California, Berkley.

Lessig, L. (1998). *The Architecture of Privacy.* Paper presented at the Taiwan Net'98, Taipei, Taiwan.

Lessig, L. (1999). *Code and Other Laws of Cyberspace*. New York: Basic Books.

Nguyen, D. H., & Mynatt, E. D. (2002). *Privacy Mirrors: Understanding and Shaping Socio-technical Ubiquitous Computing Systems* (GIT-GVU-02-16): Georgia Institute of Technology.

OECD. (1980). *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data*, available from: http://www1.oecd.org/publications/e-book/9302011E.PDF

Palen, L., & Dourish, P. (2003). Unpacking "Privacy" for a Networked World. *CHI Letters (CHI 2003), 5*(1), 129-136.

Posner, R. (1978). An economic theory of privacy. *Regulation*, 19-26.

Rosencrance, L. (2000, September 5). Amazon charging different prices on some DVDs. *Computerworld*. Available from: http://www.computerworld.com/industrytopics/retail/story/0,10801,49569,00.html

Schulzrinne, H., Tschofenig, H., Morris, J., Cuellar, J., & Polk, J. (2004, November 28). *A Document Format for Expressing Privacy Preferences for Location Information* [Internet Draft]. The Internet Society. Retrieved 26 January, 2005, available from: http://www.ietf.org/internet-drafts/draft-ietf-geopriv-policy-05.txt

Sokol, D. (2004). The Moral of Santa's story. *BBC News Magazine*. Available from: http://news.bbc.co.uk/1/hi/magazine/4121991.stm

Video Privacy Protection Act (1988), 18 USC 2710. Available from: http://www4.law.cornell.edu/uscode/18/2710.html

Warren, S., & Brandeis, L. (1985). The right to privacy, *Ethical issues in the use of computers* (pp. 172 - 183). Belmont, CA, USA: Wadsworth Publ. Co.

Yee, G., & Korba, L. (2005). Semiautomatic Derivation and Use of Personal Privacy Policies in E- Business. *International Journal of E-Business Research, 1*(1), 54-69.