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

Abstract

As with all the major advances in information and communication technology, ubiquitous computing (ubicomp) introduces new risks to individual privacy. 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. We illustrate and compare the protection afforded by 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. Our model balances the user's privacy preferences against the applicable privacy regulations and incorporates five types of user controlled "noise" to protect location privacy by introducing ambiguities. We also incorporate an economics-based approach to assist users in balancing the trade-offs between giving up privacy and receiving ubicomp services. We conclude with a scenario and heuristic evaluation which suggests that regulation can have both positive and negative influences on privacy interfaces in ubicomp and that social translucence is an important heuristic for ubicomp privacy interface functionality.

Keywords: ubiquitous computing, privacy, legal, regulation, location-dependent and sensitive, pervasive computing

1. Introduction, Motivation and Scope

It is estimated that in 2005 there are some 2 billion mobile telephones in global use and RFID tags will be deployed in the billions by the end of the year (Lee, 2005). Mobile telephone handset capabilities range from being able to communicate with the network as a kind of dumb terminal using asynchronous text messages (GSM) to having embedded general purpose computers able to make video calls or transmit data at high speeds (GSM or UMTS SmartPhones). With mobile telephone penetration exceeding 100% in many areas of North America, Europe and the Far East, 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 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.

[1] Corresponding Author. Telephone +44 1908 653 701, Fax: +44 1908 652 140
E-mail addresses: [B.A.Price, K.A.Adam, B.Nuseibeh]@open.ac.uk

Risks such as fraud and identity theft are so great, and such a significant proportion of the planet's population is 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 (2005) identify the risks of location data becoming public – both for real-time data (location-based spam and stalking), and 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 users are aware of their 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 both in the US (P&AB, 2003; Taylor, 2003) and in Europe (Dawson et al., 2003) ranging from "unconcerned" to "privacy fundamentalist". 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).

Milberg et al. (Milberg et al., 1995) noted the general information systems relationships between nationality, cultural values, personal privacy concerns, and privacy regulation. In this paper we look at some of the differences in national/cultural values and the effect regulation has on ubicomp privacy interfaces. We are concerned with providing users 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 1). 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. 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 interfaces, by their very nature, must be able to accommodate seamless movement between different regulatory regimes. A ubicomp interface must ensure

that services comply not only with local laws, but also provide an appropriate level of privacy support based on the user's privacy preferences wherever the local law is weak or nonexistent. It must also warn the user when local law requires a disclosure which violates the user's privacy preferences. 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. We use the terminology of Gunter et al. (2004) who identify *holders* as the principals in an infrastructure that collects location data or *sightings.* These sightings might be generated from a mobile telephone network using signal triangulation, a GPS tracking system, or accesses to short-range wireless network equipment (WiFi or Bluetooth) connected to the Internet. A *subscriber* in this case is a system or service that uses data collected by holders, as opposed to the traditional definition of the person using (or pay the airtime bill for) a mobile telephone. Although some of the scenarios and related work rely 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 at this time.

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

The contribution of this paper is an analysis of the factors affecting user choices for control of ubicomp privacy. We consider both the architectural and legal implications and propose a model which incorporates them. Our model incorporates a number of types of "noise" to hide personal data and an economics-based model to help determine which data to reveal.

In the following sections we examine each of the layers of our model moving outward from the user. In section 2 we examine 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. In section 5 we examine 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 0 we present a model incorporating regulatory regimes, privacy protecting noise, and an economics-based approach to revealing data. 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 scenario and an evaluation incorporating a modified heuristic walkthrough.

2. What is Personally Identifying Information?

Personally identifying information (PII) is often subjective. There is usually some amount of information whose access requires control by their owners (subjects); PII can range from the identity of an individuals to their shopping habits. PII extends only those items that can be directly or indirectly linked to a single person (or in EU-speak a "natural person") and does not include aggregated anonymous data. 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 2.

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 are used: either *real-time* (where the implied timestamp is "now") or as a *historical* record. We note that dynamic/historical data are 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, 2005).

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 2 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 (2004)

US 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 relying on a patchwork of laws (described below).

As we will see in the next section, the different approaches taken between Europe and the US mirror 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, even within the organization, 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 (currently numbering 46); 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 are 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; self-regulation (e.g. TRUSTe (2004)) is promoted as an alternative, with virtually no mechanisms for redress of violations. Some US states have stronger privacy protection than others. When 145,000 consumer data profiles (including social security numbers) were stolen from the US data aggregator ChoicePoint in October 2004, they were only obligated to notify the 35,000 Californians on the list because California is the only state to require notification of security breaches (Claburn, 2005).

Unlike other Western countries, the US 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 regulations, 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 Toysmart (attempted), Microsoft (Office 98), and RealNetworks (RealJukebox) (Bronski et al., 2001), 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 (see Table 1 for a classification of the four regulatory regimes).

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 0) 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.

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, while 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 identity of the user;
ii. *hashing*: disguising the identity 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 and 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 mis-used, *avoidance* permits the release of PII but takes steps to try to prevent misuse; while *detection* is the process of sensing when mis-use 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 are 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 (2005) 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.

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 have 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. Using 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 cash machine). Some 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, some merchants may charge you a higher price because they know 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 (2004) 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 while the innermost layer is composed of users and their personal privacy policies (see Table 1). 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 et al. (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 are 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 (i.e., anonymising, hashing, blurring, cloaking, and lying) to be employed and removes computational load from the ubicomp device, thus allowing the user to have a richer privacy policy;
2. including a notify tag in the metadata allows enforcement of user feedback requirements (knowing you are being watched and discouraging overzealous spying);
3. 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 latter is stricter it can apply appropriate techniques before the user can accesses services;
    - warn the user if the regularatory regime *requires* release of PII in violation of the user's policy (e.g., E-911 or E-112);
4. digitally signed metadata attached to PII allows a broad range of enforcement techniques, including a community of peers.

The first two reasons above support the concept of Social Translucence (Erickson et al., 2002) upon which we elaborate below. The third reason supports our notion that the regulatory regime is tightly interconnected with the enforcement of the user's privacy preferences. The final reason is included to provide the user with a trustworthy mechanism for ensuring and enforcing privacy (encryption and digitally signed metadata).

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. 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 (2004) 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 users' experience or use that information against them in non-competitive pricing. In order to measure the probability of either of these events, 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 require immediate gratification, or where they need to decide whether or not to relax their privacy constraints to receive a service. The next sub-sections presents a scenario to illustrate our model followed by an evaluation using a modified heuristic walkthrough using the scenario.

6.1. Scenario

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

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 are 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.

6.2. Evaluation

The evaluation of ubicomp models and interfaces is often hampered by the lack of user experience of the technologies, combined with the fact that the technologies are not fully implemented. Privacy and security interface problems, especially for novices, are well documented (e.g., Whitten & Tygar, 1999). This has lead to a preference for heuristic evaluations allowing experts to take the role of users. We chose a heuristic walkthrough because some studies (e.g., Po et al., 2004) have shown that heuristic walkthroughs are more effective at evaluating ubicomp interfaces than traditional heuristic evaluations because the use of scenarios allows the inclusion of vital contextual information in scenarios.

6.2.1. Method

Our pool of evaluators included three senior HCI academics (including one co-author of a major HCI textbook) and two mature final-year HCI PhD students each with

several years of HCI experience in industry. All are long term EU residents (UK, The Netherlands, Greece) with complete fluency in English. We interviewed each asking about their use of store loyalty cards and e-commerce shopping activity to determine which of the three major privacy segmentations they fall into (unconcerned, pragmatist, or fundamentalist from the Westin/Harris Privacy Segmentation Model (Taylor, 2003)). One academic was clearly in the fundamentalist end of the spectrum while the others were privacy pragmatists, although one of the other academics had some fundamentalist leanings. We selected two of the academics, a fundamentalist and a pragmatist, to act as consultants on the study design and help choose the heuristics for the others to employ. This allowed us to get two independent perspectives covering a range of privacy perspectives. They settled on the following six heuristics:

- Visibility of system status
- User control and freedom
- Error prevention
- Flexibility and efficiency of use
- Help users recognize, diagnose, and recover from errors
- Social Translucence

The first five are chosen from Neilson's (1994) original ten. The last heuristic encapsulates the notion of *how well the interface supports social norms in communication.* We believe that the revelation (or hiding) of one's personal information to others in a ubicomp setting is simply a specialized case of CSCW. Erickson et al. (2002) coined the term *Social Translucence* to describe CSCW system design which allows users to draw on their social experience and expertise to structure interactions with each other.

In our evaluation, we sought to compare three models of ubicomp privacy control: the standard commercial implementation on mobile telephones in Europe used by dozens of providers to comply with EU law, the Confab model (Hong & Landay, 2004), and our extension to the Confab model. Since a standard heuristic walkthrough would be difficult without implanted or equivalent mocked up interfaces, we followed a modified version designed in consultation with our two experts. They noted that in contrast with a traditional heuristic walkthrough where working interfaces or mock ups were used, we wished to test how well the functionality of a model supported the ubicomp heuristics. The method chosen included a standard presentation of each of the three models to the evaluators, including the scenario presented above in subsection 6.1. The models were presented using a written description combined with some static images along with an animated demonstration using Topiary (Li et al., 2004), a tool for prototyping location-enhanced applications using hand-drawn mock-ups and story boards supporting Wizard of Oz techniques. We asked each evaluator to follow a think-aloud protocol as they worked through each scenario for each model and considered each heuristic. We took verbatim field notes for each. Once the evaluator had considered all three models and their questions were answered, we asked them to rank each model's compliance with the heuristics on a 1 (least) to 5 (most) Likert scale. The results appear in Table 4.

6.2.2. Results

Each of the evaluators' scores across each of the models and heuristics appears quite consistent. As one might expect, the simple GSM mobile telephone interface (which complies with EU privacy law and is in wide use commercially) was ranked at the bottom of the scale on all the heuristics. This model has a clunky asynchronous messaging interface providing an all or nothing option on revealing the user's location and no feedback to the users as to when or how often they were being "watched". Two of the three evaluators made explicit reference to this social translucence aspect in their qualitative comments (e.g. "I wouldn't know how many people were checking up on me").

The higher scores for Confab and our proposed extensions are also not surprising, although we were surprised at the lower score for error prevention in the Confab extension. The evaluators stated that the richer functionality and rules provided by the extensions also created a potentially confusing interface with complex rules which were prone to user error. Our model showed the highest score for social translucence and comments from evaluators indicated that this was likely to engender more trust in the technology by users ("it's less like a one-way mirror and more like a window").

6.2.3. Discussion

The simple design of the current standard interface in Europe for user control of privacy complies with the strong EU privacy regulations. EU legislators were clearly trying to strike a balance between privacy protection and compatibility with existing technology but the result is an interface with poor user support for privacy control. Strong regulation has, however, resulted in widespread deployment of the technology across all GSM providers. Similar clear legislation in Japan has resulted in a very wide deployment of location-based technologies while other countries without clearly defined regulation in this area lag behind.

The low score for error prevention in the Confab extension supports our earlier contention that effective ubicomp privacy policies for users are too complex for ordinary users and that a community of peers or consumer advocate groups will be the likely repositories of generic privacy polices that users may adopt as their default. The much higher score for social translucence in the Confab extension combined with the evaluators' comments suggest that allowing user feedback on location requests is important. This is supported by recent work investigating social context of location disclosure (Smith et al., 2005). While scenario based heuristic evaluations provided valuable initial data regarding the value of our model, additional usability studies employing members of the anticipated user community need to be performed in the future.

6.3. Future Work Implications

Our study focussed on European users and experts who are used to strong privacy and consumer protection laws. Most previous user studies of privacy segmentation (Sheehan, 2002; Taylor, 2003) and location disclosure (Consolvo et al., 2005) have focussed on US users where privacy protection is relatively weak. Since regulation can constrain privacy interfaces, our next study will attempt to extract the privacy

requirements of European users. This will allow us to compare US and European attitudes toward privacy to determine if strong regulation results in users perceiving a lower privacy risk and will give us an indication of the effects of Laws and Norms on user privacy requirements.

Our evaluation, while limited in scope and applicability, highlighted the importance of social translucence in ubicomp and it noted a significant interaction between regulation and functionality of a privacy interface in ubicomp. Strong privacy regulation combined with limited technology functionality has resulted in an interface with poor support for user control of privacy. Paradoxically, the legal protection of the user has resulted in the user having inadequate control over their privacy. However, with the expected widespread take-up of third generation mobile telephony (UTMS) in Europe in 2005 there is an opportunity to provide a richer privacy interface and it will be interesting to note if either the legislators or the interface designers take advantage of this.

The remaining societal forces affecting user privacy requirements are Market and Architecture. With the deployment of third generation mobile telephony (UMTS) becoming widespread in Europe (Lee, 2005) there is a unique opportunity to influence the Architecture and therefore the Market of the de facto ubicomp device for ordinary users. We will be testing our proposed model with real users using mock-ups and implemented prototypes using subsets of features on GSM platforms. These tests should allow us to refine our design to allow a UMTS implementation that is capable of supporting a powerful privacy interface.

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.

7. Summary

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: 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, thus supporting social translucence, 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 behavior 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.

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 and identified a number of rich areas for further exploration. 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 are taking an important step along this road.

Acknowledgements

We are grateful to Carolyn Brodie and Clare-Marie Karat for their critical reviews which have helped to improve this paper. We also wish to thank Marian Petre and Jennifer Rode for their helpful comments and advice.

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. (2004). *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.

Bronski, D., Chen, C., Rosenthal, M., & Pluscec, R. (2001). FTC VS. TOYSMART. *Duke Law and Technology Review, 0010*.

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

Claburn, T. (2005). Theft Of 145,000 Consumer Records From ChoicePoint Shows Law's Strengths, Weaknesses. *Security Pipeline*, 17 February. Available from: http://www.securitypipeline.com/60402129

Consolvo, S., Smith, I. E., Matthews, T., LaMarca, A., Tabert, J., & Powledge, P. (2005). *Location Disclosure to Social Relations: Why, When, & What People Want to Share.* Paper presented at the Conference on Human Factors in Computing Systems, CHI 2005, Portland, Oregon, USA, (to appear).

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

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. (2005). *A formal model of obfuscation and negotiation for location privacy.* Paper presented at the 3rd Int'l Conf on Pervasive Computing: Pervasive '05, Munich, Germany, 8-13 May, to appear.

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

Erickson, T., Halverson, C., Kellogg, W. A., Laff, M., & Wolf, T. (2002). Social Translucence: Designing Social Infrastructures that Make Collective Activity Visible. *Communications of the ACM, 45*(4), 40-44.

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). *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 Workshop on Privacy Enhancing Technologies, Toronto, Canada.

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., 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. available from: http://www.cs.berkeley.edu/projects/io/publications/privacy-techreport02.pdf

Lee, P. (2005). *TMT Trends: Predictions, 2005, a focus on the mobile and wireless sector*: Deloitte Research. available from: http://www.deloitte.com/dtt/cda/doc/content/Mobile%20wireless_FINAL_01FEB05_LR_FA_LOCKED.pdf

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

Li, Y., Hong, J. I., & Landay, J. A. (2004). *Topiary: A Tool for Prototyping Location-Enhanced Applications.* Paper presented at the UIST'04: Symposium on User Interface Software and Technology, Santa Fe, New Mexico, 24-27 October, 217-226.

Milberg, S. J., Burke, S. J., Smith, H. J., & Kallman, E. A. (1995). Values, personal information privacy, and regulatory approaches. *Communications of the ACM, 38*(12), 65-74.

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. available from: http://quixotic.cc.gt.atl.ga.us/~dnguyen/research/PrivacyMirrors.pdf

Nielsen, J. (1994). Heuristic Evaluation. In J. Nielsen & R. L. Mack (Eds.), *Usability Inspection Methods*: Wiley.

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

P&AB. (2003). Consumer Privacy Attitudes: A Major Shift Since 2000 and Why. *Privacy & American Business Newsletter, 10*(6).

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

Po, S., Howard, S., Vetere, F., & Skov, M. B. (2004). *Heuristic Evaluation and Mobile Usability: Bridging the Realism Gap.* Paper presented at the MobileHCI 2004, Glasgow, UK, 49-60.

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

Rosencrance, L. (2000). Amazon charging different prices on some DVDs. *Computerworld,* September 5. 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

Sheehan, K. (2002). Toward a Typology of Internet Users and Online Privacy Concerns. *The Information Society, 18*(1), 21-32.

Smith, I., Consolvo, S., Lamarca, A., Hightower, J., James Scott, Sohn, T., Hughes, J., Iachello, G., & Abowd, G. D. (2005). *Social Disclosure of Place: From Location Technology to Communication Practices.* Paper presented at the 3rd Int'l Conf on Pervasive Computing: Pervasive '05, Munich, Germany.

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

Taylor, H. (2003). Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits. *The Harris Poll*, *17*, 19 March. Available from: http://www.harrisinteractive.com/harris_poll/index.asp?PID=365

TRUSTe.org. (2004). *TRUSTe Home Page*. Retrieved 29 March, 2005, available from: www.truste.org

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

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.

Whitten, A., & Tygar, J. D. (1999). *Why Johnny can't encrypt: A usability evaluation of PGP 5.0.* Paper presented at the 8th USENIX Security Symposium, Washington, D.C., 169-184.

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.

Table 1: Four Layer Model of UbiComp Services

<table>
<tr><td colspan="5">User</td></tr>
<tr><td colspan="5">Data</td></tr>
<tr><td rowspan="2">Static</td><td colspan="2">Dynamic</td><td colspan="2">Derived</td></tr>
<tr><td>Historical</td><td>Real-time</td><td>Analysed</td><td>Composed</td></tr>
<tr><td rowspan="2"></td><td colspan="2">UbiComp Services</td><td colspan="2" rowspan="2"></td></tr>
<tr><td>Query</td><td>Interrupt</td></tr>
<tr><td colspan="5">Regulatory Regimes</td></tr>
<tr><td colspan="5">States with little to no privacy protection in law</td></tr>
<tr><td colspan="5">States with some protection (e.g. USA)</td></tr>
<tr><td colspan="5">States with strong privacy protection (e.g Canada, Australia)</td></tr>
<tr><td colspan="5">States with strong privacy protection including location aware (e.g. EU, Japan)</td></tr>
</table>

Table 2:Taxonomy of Data Types and Examples based on (Corby, 2002)

<table>
<tr><th colspan="3">Type of Data</th><th>Sub-Type & Example</th></tr>
<tr><td rowspan="4">Static</td><td rowspan="2">Identity</td><td>Offline</td><td>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. Relationships: child of, parent of, spouse of
6. Real Property Associations: home address, business address</td></tr>
<tr><td>Online</td><td>Digital ID: pseudonym, E-mail address, Username, IP address, Password</td></tr>
<tr><td rowspan="2">Assets</td><td>Tangible</td><td>Property: buildings, automobiles, boats, mobile phones
Personal Worth: credit balances, stock portfolios, debt balances</td></tr>
<tr><td>Intangible</td><td>Non-real property: insurance policies, employee agreements</td></tr>
<tr><td rowspan="2">Dynamic</td><td colspan="2">Historical</td><td>Low Resolution: Transactions: financial, travel, mobile phone records
High Resolution: UbiComp Sightings log (Time, Place)</td></tr>
<tr><td colspan="2">Real-Time</td><td>UbiComp Sightings ([Now], Place)</td></tr>
<tr><td rowspan="2">Derived</td><td colspan="2">Analyzed</td><td>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</td></tr>
<tr><td colspan="2">Composed</td><td>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</td></tr>
</table>

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, 2005)	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	Detection, Feedback, Noise (Anonymity)
8. (Lederer et al., 2002)	UI Metaphor: Situational faces metaphor – conceptualising end-user privacy preferences	Preventive	X	X	Matching Policies
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	Feedback

Table 4: Results of Modified Heuristic Walkthrough (actual and average score per model per heuristic)

Heuristic	Model 1 (EU GSM Phone)		Model 2 (Confab)		Model 3 (Extended Confab)	
	Actual	Avg	Actual	Avg	Actual	Avg
Visibility of system status	1 1 1	1	5 4 4	4.3	5 5 4	4.7
User control and freedom	1 1 1	1	4 4 4	4	4 5 5	4.7
Error prevention	2 2 2	2	4 4 4	4	4 3 3	3.3
Flexibility and efficiency of use	1 2 1	1.3	3 3 3	3	4 4 4	4
Help users recognize, diagnose, and recover from errors	2 2 2	2	3 2 4	3	4 4 4	4
Social Translucence	1 1 1	1	3 2 4	3	4 4 5	4.3

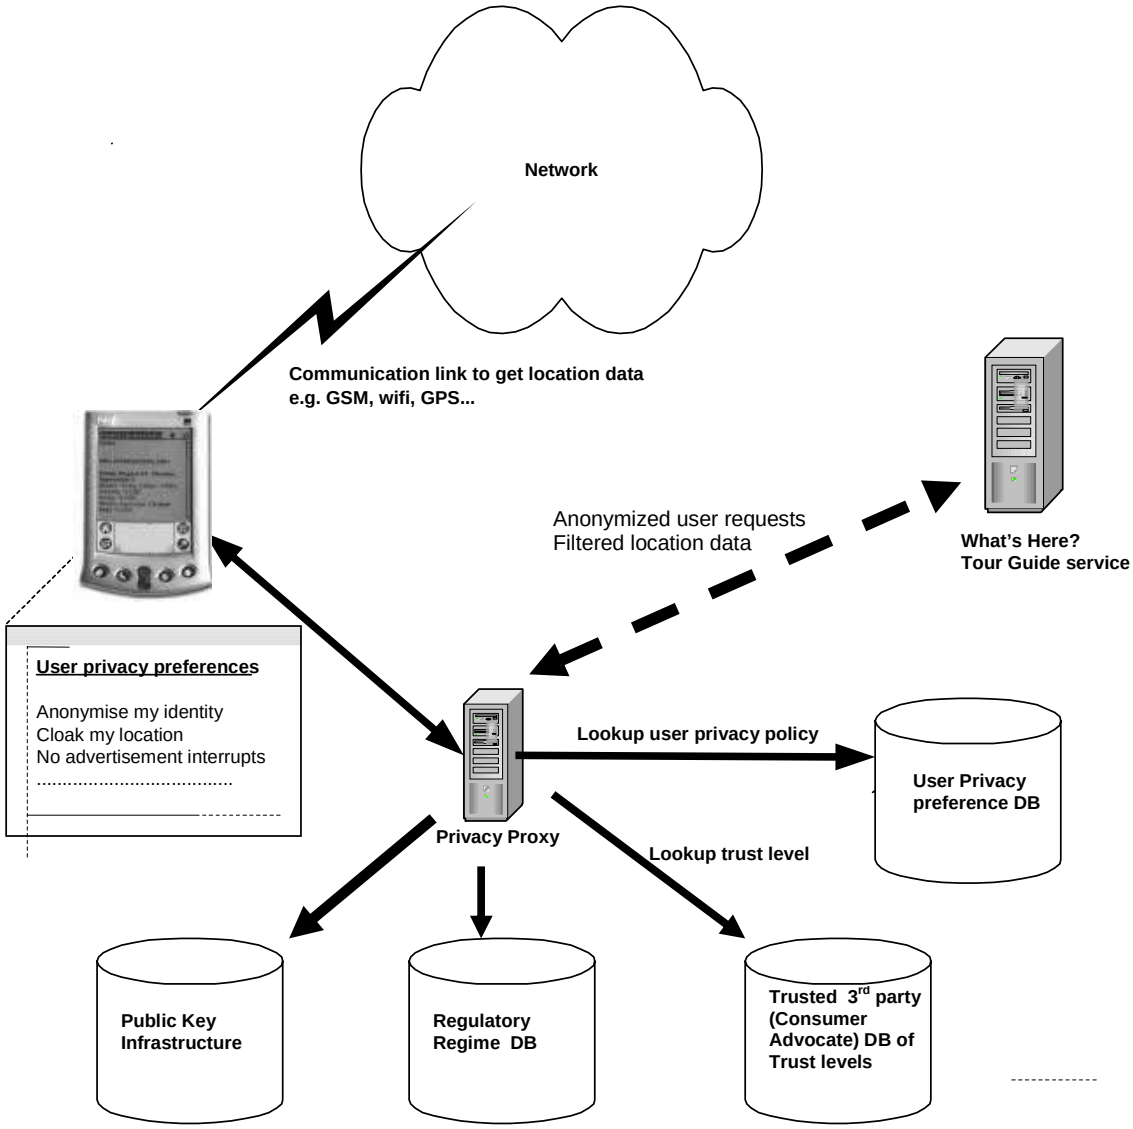

Figure 1: Privacy Model