

Maintaining privacy in pervasive computing — enabling acceptance of sensor-based services

A Soppera and T Burbridge

In the near future, everyday objects like cars and home appliances will connect the living environment to information networks. Pervasive computing devices will allow applications to gather and share a large amount of information. This may then open up a market for a large range of new services and applications. With a world densely populated by 'smart dust' [1] sensor devices, no single part of our life will be able to escape from digitisation. Soon, sensor networks will be able to track everything from our feelings to our behaviour. Besides the enormous potential value, we can foresee many undesirable uses. In a worst-case scenario, privacy implications, particularly the bad publicity around invasions of privacy, could block the incredible potential of pervasive computing. In this paper, we introduce the privacy issues found in the field of pervasive computing in two parts. The first part provides a brief look at the understanding of privacy, factors that can be used to control privacy, and the development of fair information practices and how they relate to the world of pervasive computing. The second part introduces technology that can provide a tool-set to support these fair information practices, and maintain the role of the data subject in the management of their private information.

1. Introduction

During the 1980s, Mark Weiser [2] predicted a world in which computing was so pervasive that devices embedded in the environment could sense their relationship to us and to each other. These tiny ubiquitous devices would continually feed information from the physical world into the information world. Twenty years ago, this vision was the exclusive territory of academic computer scientists and science fiction writers. Today this subject has become of interest to business, government and society. Governmental authorities exercise their power through the networked environment. Credit card databases maintain our credit history and decide whether we are allowed to rent a house or obtain a loan. Mobile telephones can locate us in real time so that we do not miss calls. Within another ten years, all sorts of devices will be connected through the network. Our fridge, our food, together with our health information, may all be networked for the purpose of maintaining diet and well-being. The Internet will move from being an infrastructure to connect computers, to being an infrastructure to connect everything [3, 4].

The development of pervasive computing will expose personal information to a host of applications. How will people maintain control of their personal information and enforce their privacy in this brave new world? This

paper presents the privacy and pervasive computing communities' efforts to develop technology, guidelines and models that can be used to manage privacy in the new world of pervasive computing.

This pervasive computing revolution has already started. A group of researchers at the University of California, Berkeley, have designed tiny sensor motes, using low-cost commercial components, which can automatically organise themselves into an *ad hoc* radio communications networks when dispersed into the environment. Each device contains an open source operating system known as TinyOS [5] that can fit in less than 8 kilobytes of memory, and can be configured or reprogrammed remotely. They can be used for an enormous range of applications including surveying natural environments and wildlife, monitoring buildings and structures and tracking objects.

Similar devices have already been deployed in the automotive sector. Sensors, black boxes and telemetry tools have been built into vehicles to improve their security, to notify when the engine needs to be serviced and to warn the driver of imminent danger. In future this information may also be used by insurance companies to create personal driving profiles and insurance quotes, or by highway authorities and law enforcement agencies.

Technology developers need to anticipate when the deployment of privacy-invading technology may generate resentment in end users and block the huge potential for the growth of beneficial applications. In such cases they need to be aware of the tools available to give control of personal data back to the users. Over the last quarter-century, principles for the treatment of personal data have been developed around the globe for IT systems and communications networks. Such principles can extend in scope to cover data collected from pervasive devices and sensors, but as we see in this paper, pervasive computing has its own challenges to develop solutions to support these principles.

In the following section we describe the opportunities and threats for pervasive computing, by introducing some of the emerging uses and the concerns for privacy that they are generating. We then discuss the social aspects of privacy, illustrating a model that can be used to analyse how people perceive privacy, and the factors that can be used to control and manage privacy. Following this we look at the development of fair information practices, concentrating on the OECD guidelines for privacy, and how these relate to pervasive computing. Finally we survey technology that can contribute to enhancing privacy and discuss these solutions in two parts. We first mention methods to control or minimise the release of sensitive information, before talking about how the flow and use of personal data can be managed. We focus our technology discussion towards techniques that can enhance the participation of the data subject, since we consider this to be the major hurdle in pervasive computing.

2. Emerging pervasive computing — opportunities and threats

Take a look into the future of a world in which minimal computing power devices are so cheap that they are embedded in the fabric of everyday life. Devices that do not look like 'real' computers will be able to disappear so effectively that end users will lose awareness of the devices' presence or purpose. The Internet will extend its presence to the physical world, and across it will flow large volumes of data that are analysed and correlated by powerful servers (Fig 1). We must discuss the consequences that this scenario introduces into everyday life before it becomes reality. Today, we can barely perceive the benefits that might be ultimately delivered, or the ingenious uses that it might be put to by malicious or indiscriminate parties. Nevertheless, worrying scenarios have already been described in books, journals and research articles.

Disappearing sensors are welcome because they hide complexity, but this also introduces some serious usability issues. If you cannot interact with the

computer, how can you tell what data is collected, where the data is flowing to, and more importantly, what are the consequences of your actions? The lack of a clear user interface introduces a tension between technology and human factors. Can we do something to maintain control or will we finally lose the ability to control our privacy?

Already private companies such as Wal-Mart and Gillette have deployed the first-generation of systems to automatically monitor their supply chains and increase the security of their assets. These systems are based on radio-frequency identification (RFID) tags [6]. These are small and commonly passive devices that transmit an identifier when scanned by a reader. Part of the appeal of this technology lies in the fact that these chips do not require line of sight to be read (i.e. they can be read with radio technology), can be scanned simultaneously, and can contain a global unique identifier for the item. The objective is to make them a powerful replacement for optical barcodes. This technology enables objects to be clearly identified, and thereby linked to an associated data record held on the Internet or in a remote database. While many companies are running trials for their supply chains and retail operations, users could also benefit from the tags by obtaining ingredient origins, dietary information, expiry dates, cooking instructions or example recipes — and that is just for foodstuffs.

Since RFID tags can be scanned unobtrusively from a distance, it is easy to realise the potential for privacy violations (Fig 2). Both consumer and manufacturer communities have already shown their concerns about this technology. For instance, Benetton made the headlines when a proposal to use RFID tags in their shops was misreported. Benetton proposed a solution to track clothes from the time they were produced until the time they were sold. The impression was that these tags would remain active even after the point of retail sale, so that they could be used to track returns or identify customers entering the shop wearing clothes previously bought from the same retailer. To protect the privacy of customers, checkout clerks can just 'kill' the tags, or alternatively the use of tags can be confined to disposable packing.

Security and privacy worries are not restricted to consumers or retail environments. Two further areas that are related to use of RFID inside organisations, and do not receive as much press are described below.

- End-user tracking

Individuals, or the employers they work for, could decide to have permanently active tags for use by authorised readers. However, other parties and

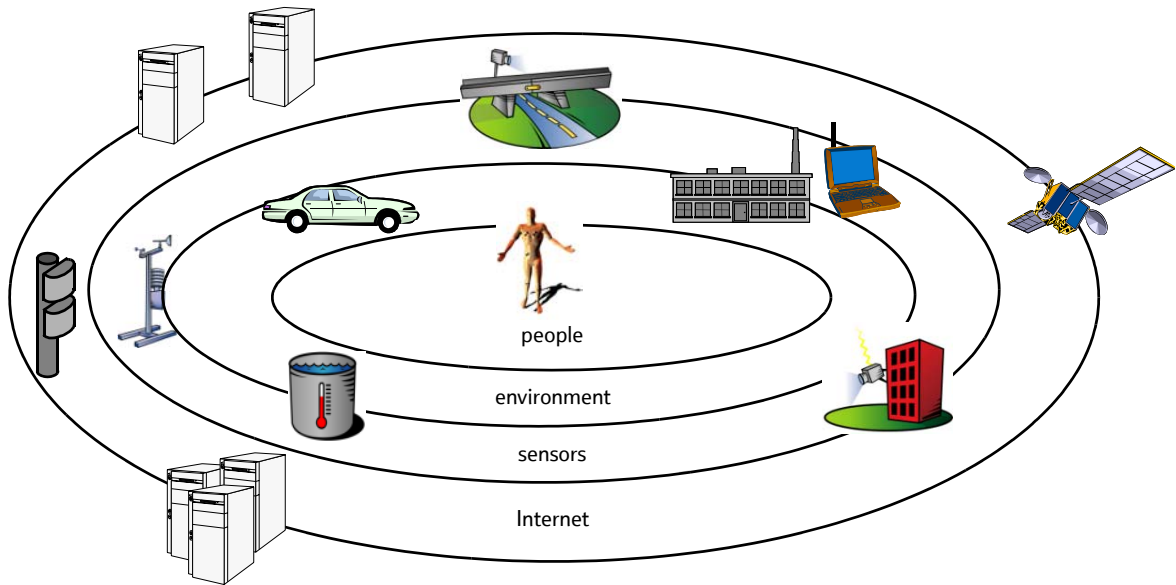


Fig 1 Connecting the physical world with the pervasive network.

applications than intended by the holder may exploit these tags. For example, remote access cards to enter premises may be used for security, and for building fire evacuation. They might also be used to clock working activities and hours, or read by third parties to identify someone who works for that company.

- Corporate espionage

The ease of monitoring competitor activities can lead to industrial espionage. For example, competitors can easily track the movement of pallets or trucks. Firms consider it unacceptable that the private information of their supply chain is visible to the outside world and can be exploited by competitors or other parties (such as thieves).

The privacy issues that a highly digitised world will face are far more complex than the consequences of simply associating an identity to each object with an RFID tag. With such tags, we can identify them and control access (through killing/activating the tag, encryption, pseudonyms or blocking technology). Other technologies may be far less noticeable, or be beyond our physical reach. Devices such as cameras or microphones read the physical world directly, rather than an associated electronic tag. We cannot encrypt or otherwise police access to the physical world. Technologists and system designers that implement and deploy pervasive computing or intrusive technologies should be concerned by the nature of networked environments and the vulnerabilities from increased connectivity of information systems.

RFID represents the first real effort to extend the Internet to global physical activity, where other sensor networks have generally been limited to closed physical environments (such as an area of forest or tidal flow).

The collection of digital information about the activities of individuals and assets, through sensor networks and aggregation with intentionally revealed information in the Internet such as purchasing or

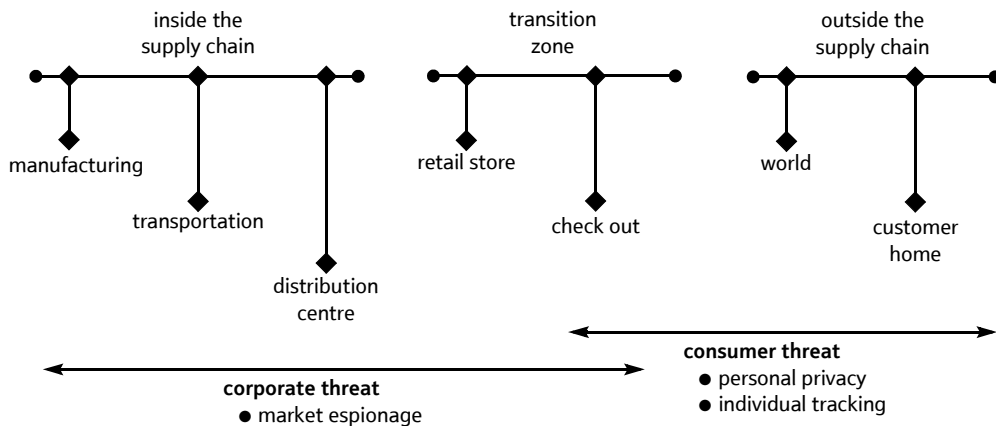


Fig 2 RFID tags — an ocean of privacy issues from the manufacturer to the consumer.

registration activities, could create detailed profile information. While from a point of view of privacy this trend of increased interaction can be seen as a great threat, some companies can foresee great business opportunities to reduce costs, and increase services. A large amount of embedded devices will ease the process of collecting personal information. A better collection and aggregation of personal data will allow different businesses to have a deeper view of consumers' behaviour and provide a better service.

Several questions arise about the value of this data. If personal data is valuable, why should not individuals benefit from this value? Can the market set a correct price for personal information by returning value to the consumer? Can we guarantee that information is acquired and disclosed only for legitimate purposes? These questions will soon require an answer.

3. Understanding privacy in pervasive computing

If privacy in pervasive computing is such a hot topic, why has the impact on technology been rather minimal? One reason is surely the fact that only a few research groups around the world have developed comprehensive pervasive or sensor systems. To date, such systems have only been deployed in restricted environments. For example, measuring humidity in a forest is unlikely to cause any great privacy uproar. Another reason is the ambiguity of people's perception of privacy. The definition of private is normally found in the field of legal studies, and technologists have a hard time to define a model that considers not only technical, but also social and economic implications. Only recently the research community has studied conceptual models of privacy to assist system designers and service providers in the deployment of (pervasive) computing.

Definitions and discussions about privacy have a long history with the expectations of individuals continually evolving in different cultures. As early as 1890 the paper 'The Right to Privacy' [7] defined privacy as 'the right to be let alone'. Today, privacy is more often about selecting what information we would like to disclose.

Over the course of the 20th Century the privacy focus has shifted with technological developments and social threats. The exploitation of detailed public records during the World War II, by Nazi Germany, allowed them to identify the Jewish population in many cities. Many European countries have developed laws to prevent such misuse of centrally stored information within their own country. During the 1960s, and particularly the 1970s, the introduction of information technology, and the use of mainframe computers and

databases, prompted the demand for new national laws on the collection of personal data. Westin at this time defined information privacy as '... the claims of individuals, groups or institutions to determine for themselves when, how and to what extent information about them is communicated to the others' [8].

This section first takes a look at work to model users' perception of privacy, along with the factors that can be controlled in order to make the release of information to be acceptable. Section 3.2 and 3.3 follow by discussing the development of international guidelines for privacy, and how they fit with the emerging world of pervasive computing.

3.1 Privacy models

Privacy has always been a very broad term, encompassing both fundamental human rights [9] and less definable personal factors. Definitions of privacy can vary widely according to the context and environment in which they are used. In the last quarter-century, definitions have often related to data protection, and these focus on the privacy of data items that contain personal secrets. It is often more easy to define privacy as a border between the society (government and private sector) and one's personal affairs. Marx [10] refines the concept of borders by introducing natural, social, spatial or temporal, and ephemeral borders. Natural borders are those governed by the natural senses, and physical boundaries such as clothes, walls, sealed envelopes and direct telephone calls. Social borders govern the expectation that information is shared within a social group, such as friends, family, work or healthcare. Spatial or temporal borders separate aspects on one's life, such as moving to university or starting employment. Finally ephemeral borders are based on the assumption that information is transitory and not captured or preserved longer than expected.

With the development and availability of personal information technology, privacy models have considered the interactions between the users and the digitised world. With the Internet, users can be exposed to systems that provide different degrees of privacy and security. Users should be able to verify what privacy protection is implemented and how their data will be used. In multimedia communications environments, Adams [11] has identified four key factors that affect the user perception of privacy — information sensitivity, information receiver/manipulator, information usage, and the context of disclosure.

Perceived infringements of privacy can lead to users rejecting the technology and thereby decreasing its commercial value. A new challenging tussle is emerging between the subjects that share information and the manipulators that exploit it for their own value. This

conflict arises because the intent of the manipulator cannot be clearly identified. Sensitive information disclosed to a trusted party may not affect our privacy, while low-sensitivity information can create resentment if disclosed to the wrong people. The user's control and feedback from the computing environments are important variables that affect the perception of privacy.

The introduction of pervasive computing raises the level of the challenge to protect privacy. Computing devices embedded in the fabric of everyday life will require systems that are able to evolve with the needs of the society and able to interact with the users to meet their requirements for privacy. Lederer [12] has proposed a cohesive model of privacy in pervasive computing by synthesising Adam's user perceptual model with Lessig's societal model [13]. Lessig's model illustrates privacy as a balance of four different forces — law, market, norms and technology. The discussion in section 3.2 and 3.3 about information practices covers some aspects of law and norms. Section 4 of this document goes on to discuss how technology can play a role in the protection of privacy. The market can be also used to control privacy. Companies must protect their brand through reputable dealings, including the treatment of private data. Also, the data subject can be returned value through the release of the personal data. This can be through better services, or simply as monetary rewards. For example, it is common practice in the USA to offer rebates on goods when personal information is disclosed. Lederer extends his model by introducing the metaphor of faces. In the real world people choose which face to present in different situations, and this concept can be extended to pervasive computing services. A face, in this case, is a meaningful representation of a user's privacy preferences in that context.

Protecting privacy in this 'brave' new world depends on the same factors identified by Lessig and Adams. We must receive value for releasing our information, have trust in practices, have protection and control through technology, and have legal recourse should our rights be infringed. We need to control our personal data, and to receive feedback on how such data is communicated and used in order to build trust.

3.2 Fair information practices

One of the influential pieces of early privacy legislation was the US Privacy Act of 1974 [14], which set down a number of fair information practices. Even earlier than this, many European countries had begun to implement laws to protect information privacy. Within the Organisation for Economic Co-operation and Development (OECD), there was concern that the development of disparate legislative approaches to privacy in member countries (including Europe and the USA)

would hinder trans-border flow of information, and thus '... cause serious disruption in important sectors of the economy, such as banking or insurance' [15]. Hence, in 1980, the OECD encapsulated eight principles among its privacy guidelines to member states. These OECD guidelines have formed the basis for much discussion and development of guidelines and legislation around the world over the past quarter-century. Even today the formation of an Asia-Pacific privacy standard across the APEC economies is starting from a set of principles very close to the original OECD guidelines [16].

The eight OECD principles are reproduced below. Whereas they appear as guidelines 7 to 14 in the OECD document, we have renumbered them here as principles 1 to 8.

1. Collection limitation principle

There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.

2. Data quality principle

Personal data should be relevant to the purposes for which they are to be used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.

3. Purpose specification principle

The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfilment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.

4. Use limitation principle

Personal data should not be disclosed, made available or otherwise used for purposes other than those specified in accordance with Paragraph 9 except:

- a) with the consent of the data subject; or
- b) by the authority of law.

5. Security safeguards principle

Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.

6. Openness principle

There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller.

7. Individual participation principle

An individual should have the right:

- a) to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him;
- b) to have communicated to him, data relating to him within a reasonable time; at a charge, if any, that is not excessive; in a reasonable manner; and in a form that is readily intelligible to him;
- c) to be given reasons if a request made under subparagraphs (a) and (b) is denied, and to be able to challenge such denial; and
- d) to challenge data relating to him and, if the challenge is successful to have the data erased, rectified, completed or amended.

8. Accountability principle

A data controller should be accountable for complying with measures which give effect to the principles stated above.

Reproduced from 'Guidelines on the Protection of Privacy and Transborder Flows of Personal Data' OECD 1980 [15]

In 1998, the OECD reviewed the continued relevance of the 1980 OECD guidelines in consideration of the 'development and diffusion of digital computer and network technologies'. Their declaration [17] stated that they reaffirmed their commitment to the 1980 OECD privacy guidelines. Despite this, the OECD guidelines are not without their critics. Some criticise the goal of the OECD principles as furthering economic trade instead of preserving the rights of individuals [18], while others question the relevance of the OECD guidelines to the modern technological world [19]. Much of the discussion around the OECD guidelines concerns the number of discretionary clauses and lack of requirements for legal enforcement. Instead, in this paper we are primarily concerned with the principles

themselves, and how well they fit with a pervasive computing world.

Clarke [18] points out that the OECD guidelines may be contradictory in that they state that the principles are '... valid for the processing of data in general, irrespective of the technology employed', while they are also limited to data on which 'automatic processing' is performed. This leaves some room for debate about what constitutes automatic processing. However, by the former statement, we can consider the OECD principles to apply to personal data gathered by pervasive sensors and other devices. In the following section we consider how well the OECD principles fit into the world of pervasive computing, both in terms of the wording of the principles themselves, and problems of implementation.

3.3 *The OECD principles and pervasive computing*

In this section we consider some of the problems with the application of the OECD principles to the world of pervasive computing technology. We find that there remain areas of privacy concern that are not covered by the wording of the OECD principles, alongside outstanding technical problems with their implementation in this new world.

3.3.1 Personal data and identity

Perhaps the main question is what is 'personal data'? The OECD guidelines define personal data to mean 'any information relating to an identified or identifiable individual (data subject)'. This means that much information, where a person's identity may not be immediately discerned, may fall outside the scope of the OECD principles. This is of great concern where pervasive devices are collecting huge volumes of data, which may only be collated to ascertain personal information at a later stage. Furthermore, the OECD principles can be interpreted to mean that data collected anonymously would be free from restrictions on use. This raises concerns, not only about the later identification, but also about tracking or behaviour analysis, along with invasions of privacy through directed marketing and customised services. For example, we do not require identity to track the path of a person leaving school at the end of lessons. For others, the offering of services based upon the identification of physical characteristics (height, weight), the clothes we wear and the objects we carry may be seen as intrusive — for example the offer of a new dietary product.

There is also a problem with the use of 'identity'. The OECD appears to only consider the identity of a physical individual. However, in an information world we

may present ourselves through the use of multiple identities. For example, in our use of the Internet, a single individual might typically use multiple log-in names, e-mail and IP addresses. Many of these cannot immediately be linked to an individual's physical identity, but this is possible through aggregation with other data. Perhaps more fundamentally, the user may regard such aliases as part of their identity, or to hold value that cannot be simply thrown away continuously to protect their privacy. For example, an alias might have been used to build a transaction history (e.g. on eBay), or to establish a presence in an Internet chat community. The use of identity to mean the physical individual also then opens the possibility of surveillance of assets and physical goods, whether owned by a company or an individual. Although a burglar surveying one's house for various goods electronically must be considered a security risk, should it be considered an invasion of privacy? Similarly corporate espionage might not relate to an individual and falls outside the scope of the OECD considerations. Of course, there is also a grey area between individual identity and assets. For example, the identification of an asset such as a car, can lead easily to the identification of the potential drivers.

We have seen that the OECD principles have not clearly considered the full extent of the privacy issues that arise in the new age of information technology, networks, and particularly pervasive computing. In the remaining part of this section we consider the problems of implementing the principles in a pervasive computing world. We first look at the problems surrounding the interface of the physical and information world, and the involvement of the data subject. We then look at the management of how the data is stored and used for the remainder of its life cycle.

3.3.2 Data collection

The collection of data is mainly governed by OECD principles 1 to 3. To summarise, the collection of data should be limited, obtained with the knowledge or consent of the subject, and relevant for a purpose that has been disclosed to the subject.

In the age in which the OECD principles were developed, and indeed in many uses in today's information world, the subject is present when such data is recorded. In this manner, the subject may read such notices, fine print, and online privacy statements before, or during the process of giving up their personal data.

Perhaps the most fundamental change in the world of pervasive computing is this lack of a two-way interface between the subject and the information world. A camera may record your movements, but how do we notify the subject and obtain their consent? The

current conventions of 'CCTV cameras in operation' and roadside speed camera signs are not scalable with the spread of pervasive devices and their uses — or will become meaningless to the extent of displaying 'pervasive computing devices in this area' signs. The OECD guidelines already only specify obtaining consent 'where appropriate'. With the spread of pervasive computing, the number of uses will grow for which it may seem impractical to ask consent. More and more of our privacy will fall beyond our ability to control.

In section 4, this paper looks at some techniques for controlling our privacy. Broadly these technologies take two approaches. The first seeks to minimise the collection of personal data through anonymity techniques. The second, and more immature area, looks for ways whereby the data subject can provide consent and retain control without their immediate presence.

3.3.3 Data usage, storage and access

In this section we concern ourselves with the OECD principles 2, 4, 6, 7 and 8. Broadly, the principles listed deal with the processes to control the legitimate usage of the data, the maintenance of correct data, and the access of the data subject to their personal data.

The OECD model relies on the 'data controller' being accountable for their adherence to the other privacy principles (as stated in principle 8). The OECD guidelines specify that 'adequate sanctions and remedies' must be in place to ensure the good behaviour of the data controller. In Europe [20] and other implementations of the OECD principles, this means the existence of supervisory bodies and supporting laws to enforce compliance. Each data controller must register the categories of data collected, and the purposes to which they are put, with such supervisory bodies. These bodies then have powers to audit the compliance of the data controllers with their specified intentions.

The data subject is given powers to request, from a data controller, what personal data is maintained, and to challenge the purpose for which it is held. However, in the OECD principles, and the EU implementation, such communications may be charged at a reasonable cost (to protect the data controller involved). This process relies on the fact that notice has previously been given to the data subject that their personal data is being collected. Without such notice, it is hard for the data subject to identify which data controllers may hold their personal data. As has already been stated, such notice will become harder to give in a pervasive computing world, and additionally it will become impossible without assistance for the data subject to

maintain knowledge of all data controllers that have some of their data.

Section 4.2 examines some techniques that allow the data subject to maintain a role in the maintenance and usage of their personal data after it has been released. Proposed technologies seek to maintain the data subject's control remotely through the use of cryptography and policies, while giving data-tracking and auditing capabilities back to the data subject.

4. Technical approaches to privacy

Before diving into specific techniques to enhance and protect privacy, it helps to have a technical overview on the security characteristics of sensor networks. These networks may consist of hundreds or thousands of low-power, low-cost wireless nodes, each with limited hardware capabilities. From a security point of view one of the main risks is the reliance on wireless network communications, so that adversaries, even if they are physically distant, can easily eavesdrop on the radio transmission. Unfortunately traditional cryptographic solutions can only partially solve the problem. The fact that the devices are extremely limited in computation and communication resources means that they will require very lightweight security protocols. Another problem is the fact that every node represents a potential point of attack. An adversary in control of a few nodes inside the network can then launch attacks against the whole sensor network. Technologists can build tamper-resistant devices, but this raises the cost of such devices. In this paper we acknowledge this problem space, but concentrate on the overall privacy of the data subject. Low-cost transmission security

techniques are evolving, but perhaps as important for privacy is the control of the information flow, as opposed to protecting that flow from eavesdroppers. The first OECD principle states that '... there should be limits to the collection of personal data'. If personal data is not collected, then it cannot be misused, and expensive solutions to control such usage become unnecessary. Given the difficulties with establishing explicit consent (due to lack of two-way communication, or limited device computational power), techniques to restrict the release of information, such as anonymity and pseudonymity, are often considered a better approach.

4.1 Anonymity, accountability and pseudonyms

Anonymity techniques ensure that a user may use a resource or a service without disclosing their identity (see Fig 3). In the communication domain we can define anonymity as the inability to link a communication to any particular sender or receiver [21]. The assumption is that if data cannot be related to the individual, it poses no threats in terms of privacy, and therefore there is no need to restrict its collection. However, it must be realised that associated information about an identified individual may be used to attack the anonymity of other data. Also, while anonymity might protect some definitions of privacy, as described in section 3, anonymous data may still be used for malicious purposes.

Accountability may be considered to be an opposing objective to anonymity. In order to have accountability, we need to identify who has taken various actions. For some sensor networks and applications, accountability




Fig 3 Protecting privacy through pseudonyms — decreasing correlation leads to increasing anonymity (reproduced from Pfitzmann and Köhntopp[21]).

is important. For example, we may wish to know which doctors have accessed the drugs storage in a hospital. Pseudonyms can be used to preserve anonymity in normal usage, but allow accountability in certain cases. A pseudonym normally means an ID that can be used to uniquely link an action to an identity and authenticate the party (or asset) involved. The strength of anonymity decreases, as more is known about a particular pseudonym — for example as it is used for more and more transactions.

The ultimate strength in terms of anonymity is achieved by changing the pseudonym for every transaction. In this manner, correlation of different transactions (and other data) cannot be used to infer identity. Another property of continually changing the pseudonym is that the item or person cannot be traced though the repeated use of the same pseudonym, which may or may not be desirable. For example, we do not wish a mugger to be able to tell that an anonymous person regularly takes a certain route between their shop and the bank at the end of the day, but we do wish to analyse road traffic flows of anonymous cars.

Pseudonym techniques are currently being proposed for use in RFID tags [22]. Each time a tag is queried by a scanner, the output must be indistinguishable from truly random values, and unable to be linked by unauthorised parties to the actual ID of the tag. This kind of technique is simple and cheap to implement on the tag, but authorised readers must have access to the list or chain of pseudonyms that will be used, and be able to map backwards from the pseudonym to the actual ID.

While pseudonym techniques are useful for anonymity, they do not offer perfect protection for personal information for two reasons. Firstly, as has been explained, anonymous information may be correlated with other sources of information, and even anonymous information may hold some value against the interests of the data subject. Secondly, ubiquitous computing environments include sensors that sense our real identity or other physical properties. For example, it is much harder to continually change our physical appearance to fool a digital camera, than it is to change an electronic ID in an RFID tag. This is compounded by observations performed by people — such as knowing when a person leaves a house in the morning. Although pseudonyms might have uses such as to hide number plates on cars from parties other than the police or highways authorities, we must realise that our current social interactions and presence in the real world cannot be completely hidden behind pseudonyms. Unless we decide to live in a completely digitised world and interact only through computers and networks, some information about our real identity will always be disclosed.

The final problem with anonymity is that we risk removing valuable services. For example, we can share pseudonym information about car registration numbers with the police, but for every additional party we wish to be able to offer services, we must explicitly arrange for them to be able to break our anonymity. A common approach to this problem is to use trusted intermediary (or group of intermediaries) to protect our information, but still allow provision of multiple services. Past approaches have taken the approach of mixing information from different subjects, or aggregating subject information. A few examples of such systems for use over the Internet are Mix-nets [23] for anonymous e-mail and Crowds [24] for Web browsing.

Interest has been awakened in solutions to protect privacy during the provision of location-based services. This is particularly important in the light of European directive 2002/58/EC [25]. This directive brings the European privacy directive 1995/46/EC [20] and its telecommunications counterpart 1997/66/EC [26] up to date for 'new advanced digital technologies' and the 'introduction of new electronic communication services'. Along with new guidelines for the use of e-mail, services, directories, anonymous communication and transport data (such as routing records), the directive in Article 9 also specifically focuses on location-based services. It states that: 'Where location data ... relating to users ... can be processed, such data may only be processed when they are made anonymous, or with the consent of the users'. It goes on to state that, prior to consent, the users must be informed of the 'type of location data' that is being used, along with the 'purposes and duration of the processing', and any third parties involved in the delivery of value added service. Furthermore, users must have the means to temporarily withdraw their consent per communication of location information. Given the arduous tasks of implementing such notification and consent procedures, it is likely that location-based services will instead choose to process anonymous data.

Two interesting approaches to providing anonymity for location-based services are Mix Zones [27] and k-anonymity [28]. In Mix Zones, the user's location data is transmitted to a trusted intermediary. The user has the ability to specify location regions or zones, within which they are willing to share their location information. Critically, while within one of these zones, instead of presenting the user's ID to the location-based service, a pseudonym is generated and used instead. A new pseudonym is created every time the user changes zone. Although potentially the tracing of individuals is possible, given high enough populations and errors in location accuracy, the application's confidence in tracking a single individual is quickly diminished. In the approach of Gruteser and Grunwald [28], the user

identification is also removed, but, in addition, the user location information is blended together such that a user's location is indistinguishable from the location of $k-1$ other data subjects.

We can see that anonymity solutions are key technologies to delivering services in the world of pervasive computing. They rely on the principle of minimal release of information, such as through the mixing or aggregation of data. If data is not released, then it is impossible to abuse. Pseudonyms can also be used for selective release of information through the careful release of data that allows the pseudonym to be matched to an identity. Similarly encryption techniques can be used to distribute information in a controlled manner. In section 4.2 we look at techniques for controlling access to information (OECD principles on collection limitation and purpose specification), and methods for tracing and auditing its use (OECD principle on openness, participation and accountability).

4.2 Privacy management

In his controversial book 'The Transparency of the Society' [29], Brin introduces a world in which privacy is non-existent. The Internet allows people to gather information across great distances and to correlate information at unimaginable speed. Brin warns that governments and major businesses are exploiting technology for ubiquitous spying and suggests that the only possible remedy is to have a world where information is free and everybody has the capability to spy on everybody else. In this manner we can at least see who is spying on us. While this approach might work in anarchist socialism [30], where the people have the power to back up the information and to enforce correct behaviour, it is hard to imagine it working in today's society. Even if we do not share the same pessimistic vision of Brin, we can agree that a total protection of our privacy will hardly be achievable, or desirable. What we can expect is the introduction of more technology to enhance our privacy.

Many Web sites and Internet applications now carry associated privacy policies. Similarly there are privacy notices in the real world (e.g. CCTV) and small print on subscription forms. Aside from the odd tick-box, how many people actually manage to read such notices and assert control over their privacy? In a Federal Trade Commission Workshop on Consumer Data [31], the Excite@Home privacy officer stated that only 100 out of 20 million visitors accessed the privacy policy the day after they featured in an Internet privacy segment of a popular TV show [32]. To address this problem, work has started within the W3C on the Platform for Privacy Preferences or P3P [33]. This project enables Web sites to express their privacy policies in machine-readable XML formats. Before a page is accessed from the user's

browser, an agent first checks the privacy policy for that page, and alerts the user to any discrepancies between the page and the user policy. Along with the criticism of the legal foundation for P3P, the technical criticism of P3P has been curiously bipolar. On one hand, there are concerns that the policy expression is not complicated enough to accurately capture a company's (or user's) privacy policy [34], and that it does not allow negotiation over privacy settings or operate over more than just HTTP [35]. On the other hand, there are suggestions that user policies across multiple applications or Web clients is unmanageable, and that users will be unable to understand and express their own complex policies and have to fall back to default settings [36]. Most critics are agreed that P3P is a step in the right direction — just that it has some way to go before maturity.

So is P3P a solution for pervasive computing? The answer is likely to be 'no' for a number of reasons. If the complexity of policies is hindering the development of a privacy solution for Web access, then these problems will be massively compounded in the pervasive computing domain. Following the privacy model of Adams (see section 3.1), we might wish ultimately to express policies in terms of data, destination, purpose and contextual information. For example, I might wish to allow Tesco to use my location information only when I am in the Ipswich-based stores, and only for the purpose of notifying special offers. We can only imagine how complicated it might become to express such policies for every possible interaction in a highly pervasive computing world of the future. Although our ability to express computer-understandable policies will grow, we cannot get rid of the requirement for the user to understand what is happening with their data. However, policies are likely to form another key component in the support for personal privacy — but perhaps in limited domains where they can be easily understood and codified.

The greater restriction of P3P that limits its applicability to pervasive computing is the requirement for the user (or at least a user agent) to be present in the information transfer. In pervasive computing, the information may flow from a user-controlled device, but might equally flow from another sensory device in the environment. In this latter case, there is no clear binding to an identified subject. For example, although cameras and microphones might pick up privacy sensitive information, there is no easy way to identify at that point, either the person involved, or how to examine their privacy preferences. Langheinrich has proposed one solution to this problem in his Privacy Awareness System [37]. In his system, sensors beacon their identity to the environment, so that users can then contact a privacy proxy for that device and negotiate their privacy

settings. An alternative solution that does not require standardisation of local wireless communication protocols would be for a user device to employ a positioning service, and to examine sensors near that location for their privacy policies. Apart from removing the requirement for a standard announcement protocol and wireless communications medium, this solution also has the advantage that other locations (perhaps in advance of our travels) can be examined. Both approaches have applicability for general pervasive device control beyond just privacy settings — for example, the detection and setting of a central heating system.

A similar binding between the data subject's preferences and the data read by the sensing device is achieved with RFID tags by using 'soft blocking' [38]. When the RFID reader scans tags within range, one or more of these tags may be a soft blocker tag. In the simplest scenario, the presence of a soft blocker tag indicates that information about other tags (that are flagged as private) may not be transmitted. In a more complicated scenario, different soft blocker IDs may indicate different treatments of other tags, and the privacy policy associated with that soft blocker tag is found by a database look-up.

Another problem with the approach of P3P is that the information is delivered to a single application, and therefore controlled by a single policy. In a highly pervasive computer world, information from one sensor may be delivered simultaneously, along with other sensor information, to numerous applications. For this reason Casassa-Mont et al [39] suggest attaching a 'sticky' policy to the information. Information fields can be encoded using identifier-based encryption (a technique to allow the string-based encryption of data, and the subsequent generation of the decryption key [40]) and delivered to multiple applications. The application then refers to a trust authority and attests to fulfilling the conditions of the use policy before gaining the keys for access to the information. The trust authority, in this manner, also is able to perform tracing and auditing of the data flow and usage. One advantage of this end-to-end approach is that the information can be communicated freely over any untrusted communications medium from the users to the applications. The use of sticky policies suffers from the problem of how to attach such sticky policies to the data in the first instance. Although the techniques discussed above, such as from Langheinrich [37], might allow the user to interact with the sensing device, the construction of a suitable policy might be extremely complex. This problem is exacerbated since the policy is constructed knowing only the data, and any user context information. At the point of policy construction no information is available on the destination of the data,

or indeed on the purpose for which it will be used. Thus, the sticky policy approach might at first appear a poor match for the implementation of either Adams' privacy model, or the OECD principles. However, it is easy to imagine that such a solution might be extended to include negotiation over destination and usage at the point that the application contacts the trust authority. In fact, Casassa-Mont et al suggest that the policy might specify that the trust authority must obtain explicit user authorisation before releasing the data. A similar mechanism could be added to initiate further policy negotiation with the user, or a user policy agent.

The solution of Casassa-Mont et al shows how the tracing and auditing of privacy sensitive data may be achieved. Such a solution can be used to realise Tygar's suggestion that strong audit mechanisms are available to 'watch the watchers' [41]. Everyone using personal data should be subject to auditing, and alternatives to the current European (and OECD) model, where data collection and purpose are registered with a central authority, should be examined. Such distributed alternatives aim to maintain the link between the data subject and their personal data, and to allow the data subject to participate in the auditing function.

Opponents of such strong auditing can make the case that the complexity and expense of achieving a secure solution will be prohibitive and will cripple the deployment of new services. Another advocate of sticky policies is the work by Karjoth et al on a 'Platform for Enterprise Privacy Practices' (E-P3P) [42]. In contrast to the work by Casassa-Mont et al the E-P3P approach uses sticky policies to consistently manage a data subject's privacy within an enterprise. The policy is constructed within the enterprise in accordance with the external policy to which the data subject has agreed concerning release of the information. Access to the data is then restricted to specific users or roles within the enterprise, and for specific purposes. Such work allows the organisation to have clear visibility and control over personal data, and to easily demonstrate its accountability to auditors or data subjects.

5. Research challenges

In the absence of direct interaction on the part of the data subject, the use of policies is essential to protecting privacy in pervasive computing environments. Although it can be seen such work is starting to develop, more needs to be done to make such solutions realistically achievable across a range of applications.

Along with further work on binding policies to data (beginning with binding data subjects and sensors), we can imagine further complexities that have not, to the authors' knowledge, been solved. One example is the

multiple 'ownership' of data. For example, a camera will inevitably include multiple data subjects in a single frame, and may also claim ownership of some of the information such as the background or timestamp. Allowing these multiple parties to each control their privacy will be a huge problem. Simply going for the highest common privacy protection may not be an adequate solution since this will block some users from obtaining legitimate services. Another potential problem occurs during the decomposition or aggregation of data. The splitting, or aggregation of privacy policies in such cases will be hard since it is not clear how much information is lost in such operations, and whether the previous policies are still valid.

6. Conclusions

A comprehensive review covering all angles of the privacy problem and potential solutions for pervasive computing has not been possible in this paper. Although privacy is not a new problem, the examination of privacy in the area of pervasive computing is immature, and we have tried to provide the reader with a flavour of how current legislation and technology fit with this emerging field, along with some of the new research being conducted. What emerges is that many of the components that we require to protect privacy in this new age are either in place, or beginning to develop. However, a comprehensive approach that sets forth a set of privacy principles for pervasive computing, and a technical framework to aid those principles is missing. Perhaps, at least in the case of technology, the absence of an overall framework is correct. For example, it is hard to imagine the development of a policy-based control mechanism or auditing function that might be universally applied.

The cost and complexity of a universal policy-based system, along with the problems of standardisation, may mean that such a system is never implemented. Instead, perhaps, the problem of privacy is better addressed in pieces. Through restricting the range of applications, and hence the data, purposes, and other contextual information, we can begin to make policy-based systems manageable. For example, it is easy to imagine that policies might be expressed to control location data to a range of services operating over a common privacy-enabled middleware platform.

Acknowledgements

Thanks to our many reviewers, including Vivekanand Korgaonkar, Arnaud Jacquet, Ben Strulo, Alan Smith, Clazien Wezeman and Gabriele Corliano. Thanks also to Marek Rejman-Greene for some interesting discussion and debate, and to Bob Briscoe for his contribution to kick-start the work.

References

- 1 Kahn R, Katz H and Pister K: 'Emerging challenges: mobile networking for 'Smart Dust'', *J Comm Networks*, pp 188—196 (September 2000).
- 2 Weiser M: 'The computer for the 21st century', *Scientific American*, 256, No 3, pp 94—104 (1991).
- 3 Krikorian R: 'The Net Comes Home', *New Scientist* (February 2003).
- 4 Overby C S: 'The X Internet and consumer privacy', *Forrester Report* (December 2003).
- 5 Culler D et al: 'TinyOS: an operating system for sensor networks', to appear in Rabaey J (Ed): 'Ambient Intelligence', Springer (2004).
- 6 Sarma S E et al: 'Radio-frequency identification: security risks and challenges', *RSA CryptoBytes*, 6, (2003).
- 7 Warren S and Brandeis L: 'The Right to Privacy', *Harvard Law Review*, IV, No 5 (December 1890).
- 8 Westin A: 'Privacy and Freedom', Atheneum, New York (1967).
- 9 Laurant C: 'Privacy and human rights: an international survey of privacy laws and developments', *Electronic Privacy Information Center*, Washington, DC, USA (2003).
- 10 Marx G: 'Murky conceptual waters: the private and the public', *Ethics and Information Technology* (July 2001).
- 11 Adams A: 'Multimedia information changes the whole privacy ballgame', *Proceedings of Computers, Freedom, and Privacy* (2000).
- 12 Lederer S: 'Everyday privacy in ubiquitous computing environments', *Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing, UbiComp* (2002).
- 13 Lessig L: 'The architecture of privacy', *Taiwan Net Conference* (1998).
- 14 US Privacy Act of 1974 — <http://www.usdoj.gov/foia/privstat.htm>
- 15 OECD — Recommendation Concerning Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data, (September 1980).
- 16 Greenleaf G: 'Australia's APEC privacy initiative: the pros and cons of OECD lite', *Privacy Law and Policy Reporter* (2003).
- 17 OECD — Ministerial Declaration on the Protection of Privacy on Global Networks (October 1998).
- 18 Clarke R: 'Beyond the OECD guidelines: privacy protection for the 21st century', (January 2000) — <http://www.anu.edu.au/people/Roger.Clarke/DV/PP21C.html>
- 19 Justice M K: 'Privacy protection, a new beginning: OECD principles 20 years on', *Privacy Law and Policy Reporter* (1999).
- 20 EU Directive 1995/46/EC of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data (October 1995).
- 21 Pfizmann A and Köhntopp M: 'Anonymity, unobservability, and pseudonymity — a proposal for terminology', *Workshop on Design Issues in Anonymity and Unobservability*, Berkeley, California (2002).
- 22 Juels A: 'Privacy and Authentication in Low-Cost RFID Tags', *RSA Laboratories* (2003).
- 23 Chaum D: 'Untraceable electronic mail, return addresses and digital pseudonyms', *Communications of the ACM*, 24, No 2, pp 84—90 (1981).
- 24 Reiter M K and Rubin A D: 'Crowds: Anonymity for Web Transactions', *ACM Transactions on Information and System Security* (1998).

Maintaining privacy in pervasive computing

- 25 EU Directive 2002/58/EC of the European Parliament and of the Council Concerning the Processing of Personal Data and the Protection of Privacy in the Electronic Communications Sector, (July 2002).
- 26 EU Directive 1997/66/EC of the European Parliament and of the Council Concerning the Processing of Personal Data and the Protection of Privacy in the Telecommunications Sector, December 1997.
- 27 Beresford A and Stajano F: 'Mix zones: user privacy in location-aware services', IEEE International Workshop on Pervasive Computing and Communication Security (PerSec) (2004).
- 28 Gruteser M and Grunwald D: 'Anonymous usage of location-based services through spatial and temporal cloaking', ACM/USENIX International Conference on Mobile Systems, Applications and Services (MobiSys) (2003).
- 29 Brin D: 'The Transparent Society', Addison-Wesley (1998).
- 30 Guerin D: 'Anarchism: From Theory to Practice', Monthly Review Press (1970).
- 31 Federal Trade Commission: 'Workshop on the information marketplace: merging and exchanging consumer data', Washington DC (March 2001).
- 32 Fred H C: 'Principles for protecting privacy', The Cato Journal (March 2002).
- 33 Cranor L et al: 'The Platform for Privacy Preferences 1.0 (P3P 1.0) specification', W3C Recommendation (April 2002) — <http://www.w3.org/TR/2002/REC-P3P-20020416>
- 34 Pedersen A: 'P3P — problems, progress, potential', Privacy Laws & Business International Newsletter (February 2003).
- 35 Thidadeau R: 'A critique of P3P: privacy on the Web', (August 2000) — <http://dollar.ecom.cmu.edu/p3pcritique/>
- 36 Birchman J A: 'Is P3P "The devil"?', Law and the Internet Seminars, University of Miami School of Law (May 1998) — <http://www.law.miami.edu/~froomkin/sem97/birchman.html>
- 37 Langheinrich M: 'A privacy awareness system for ubiquitous computing environments', 4th International Conference on Ubiquitous Computing (UbiComp) (2002).
- 38 Juels A and Brainard J: 'Soft blocking: flexible blocker tags on the cheap', Manuscript (2003) — <http://www.rsasecurity.com/rsalabs/staff/bios/ajuels/publications/softblocker/softblocker.pdf>
- 39 Casassa-Mont M, Pearson S and Bramhill P: 'Towards accountable management of identity and privacy: sticky policies and enforceable tracing services', IEEE 14th International Workshop on Database and Expert Systems Applications (DEXA'03) (September 2003).

- 40 Boneh D and Franklin M: 'Identity-based encryption from the Weil pairing', Crypto (2001).
- 41 Tygar D: 'Security with privacy', ISAT 2002 study (December 2002).
- 42 Karjoth G, Schunter M and Waidner M: 'Platform for enterprise privacy practices: privacy-enabled management of customer data', 2nd Workshop on Privacy Enhancing Technologies (April 2002).



Andrea Soppera attained an MSc in Telecommunication Engineering from the Polytechnic of Turin, awarded in July 2000. He also holds a diploma in Network Engineering from the Eurecom Institute of Sophia Antipolis. He first worked for Techno-Concept — a US-based start-up in the domain of software radio communication. As a system designer, he was responsible for the delivery of a concept prototype for software radio GSM networks. He joined BT's Network Research Centre in 2001. His research has covered the analysis of denial of service issues and security in distributed systems. He has also been involved in the design and development of publish-subscribe messaging technology for large-scale distributed systems. He is currently working on privacy and security for sensor networks and pervasive computing systems.



Trevor Burbridge is a member of the Networks Research Centre, researching solutions to the issues of scalability, management, and distributed control in the field of group communications. He is currently working towards overcoming the privacy barriers to the deployment of new services enabled by the deployment of pervasive computing. In recent years he has been active in the development of scalable publish-subscribe event notification systems, and in the past has focused on large-scale enterprise integration and the distribution of business intelligence and data. He attained an MEng(Hons) in Computer Software and Systems Engineering from the University of York in 1994, and pursued postgraduate research in computer agent co-ordination before joining BT in 1997.