

# The development of a privacy-enhancing infrastructure: Some interesting findings

Patrik Osbakk  
Computing Laboratory  
University of Kent  
Canterbury, Kent  
CT2 7NF, UK  
+44 1227 823824  
pjo2@kent.ac.uk

Nick Ryan  
Computing Laboratory  
University of Kent  
Canterbury, Kent  
CT2 7NF, UK  
+44 1227 827699  
N.S.Ryan@kent.ac.uk

## ABSTRACT

Providing privacy protection for ubiquitous environments is a complex task that has only recently become a hot topic. In this paper we describe the current state of our privacy-enhancing infrastructure and address some issues that have arisen during its evolution. In particular we contrast users' online and offline privacy concerns, and their perceived and actual performance in configuring access control mechanisms. We also present a brief assessment of cryptographic performance on small devices.

## Keywords

Privacy, Context, RBAC, PIV, CCS, P3P

## 1. INTRODUCTION

Privacy concerns have been expressed since the birth of ubiquitous computing [1]. Yet it is only in the last couple of years that research into privacy protection for ubiquitous computing has really taken off. Though work in this area is still very much in its infancy, the variety of existing research [2][3][4] shows that privacy is not a simple issue that can easily be addressed. Privacy issues surround all aspects of ubiquitous computing and must be taken into account throughout the design. This paper will describe our privacy-enhancing infrastructure for context-awareness and present some of the issues that have arisen during its development.

## 2. PRIVACY POSITION

The foundation of our infrastructure is the beliefs held and the assumption made about privacy with respect to ubiquitous and context-aware computing. Throughout this work we think of privacy in terms of information flow. The adopted definition is "Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others" [5]. As such, the ownership of information is with the subject. It is thus important to note that privacy does not equal isolation, nor is this desirable. Rather, the issue is one of control.

Control can only extend to the point of information release. Once released the subject cannot control how information is used. The process of releasing information is therefore thought to be like a business transaction where the disclosure of information follows some form of agreement that governs how the information will be used. Such agreement can be explicitly stated, though often it will be determined implicitly by cultural and social norms, existing relationship between the parties, etc. An important factor in determining whether or not to release information is therefore the trust placed in the recipient regarding their intention to honour such an agreement, whether

explicit or not. Legislation requiring adherence to explicit agreements may offer extended protection beyond the point of release, but enforcement would be difficult so it might at best act as a deterrent against improper use.

The difficulty of achieving perfect privacy protection should not discourage the provision of some protection. The offline world, with which we interact everyday, provides far from perfect privacy given the extensive use of technologies like credit cards, loyalty cards, mobile phones, CCTV, etc. Ubiquitous computing systems should at least aim to uphold similar levels of privacy.

## 3. CONTEXT-AWARE SYSTEMS

Our work on privacy has focused on context-aware systems. This field is especially interesting in terms of privacy because much of the information used is personal and, often, sensitive. Brown and Jones go so far as to state that "Context-aware applications, above all others in the pervasive field, can be regarded as anti-privacy" [6].

The most well known context-aware applications are those that use location, but context can be much more [7]. The definition we use states that: *context is information related to an entity, where the information may include other entities*. An entity can be anything from "people, places, and things" [8] to activities and concepts. What is important is the existence of relationships linking entities, or entities and values, in a potentially extensive network.

## 4. INFRASTRUCTURE

In our privacy-enhancing infrastructure, each entity has at least one associated Context Manager (CM). Multiple synchronised CMs can, if necessary, be employed in distributed environments. A CM acts as a point through which all context information owned by the associated entity flows. It is the CM that is responsible for providing privacy protection, storage, and handling of context information.

Applications (agents) can be either context consumers (clients) or context producers (services). Depending on their type, agents communicate with the CM to either retrieve or store context information. The collection and use of context is separated, as others [9] have found desirable. This approach simplifies the development of applications as they need only consider the usage of context and the underlying mechanism for gathering context.

Communication is an important aspect of the infrastructure. To support heterogeneous devices with unknown and variable connectivity, a component-based approach has been used with plug-ins to support different communication media. The infrastructure does not rely on external security as, even when

network security is present, it may be deemed to be inadequate or flawed [10]. Instead cipher plug-ins are used to secure the communication. These, rather than a fixed cipher, ensure that differences in requirements and legislation may be supported, and allows use of custom cryptographic processors if available.

## 5. APPLICATIONS

A number of simple applications have been developed to use the infrastructure, including an iButton Context Capture application for gathering information, a Context-Aware Desk Display that shows information about the occupier, and a Web Presence application for publishing context information.

The iButton Context Capture application uses iButtons [11] to gather context information. iButtons are small, uniquely identifiable, devices. They may contain memory or have abilities like temperature sensing. iButtons may be identified and their memory content or sensor values read by a PDA with an adapter. The identity can be used to look up logically linked context information.



Figure 1: Context-Aware Desk Display

The Context-Aware Desk Display consists of an enclosed TINI [12] microcontroller with an attached LCD and keypad. The display has been designed with the ordinary name tag sometimes found on desks in mind. As well as displaying the name of the person occupying the desk it also displays additional context information like the person's email address, if they are in today or when they are expected to be in next. The information is collected from the person's CM and the application can easily be extended to display more information.

It has long been the vision of projects such as Cooltown [8] that every entity should have a web presence. The Web Presence Application provides this for entities in our infrastructure. Although the CM itself can be seen to provide a form of web presence, it is not directly accessible by standard web browsers. Thus the Web Presence Application provides the link between a CM and web browser.

## 6. EXTERNAL INTEGRATION

Being able to integrate this privacy-enhancing infrastructure with other external infrastructures is something we previously identified as desirable [13]. The idea is to utilise existing sensor networks and application, e.g. those in a context-aware office or in a smart home, whenever possible instead of having to develop and deploy duplicates. Such integration can be achieved by the use of a proxy. The proxy is responsible for translation and communication of information between the infrastructures. The proxy can be developed as part of either infrastructure or as separate application if desired.

We have developed a single two-way proxy to integrate this infrastructure with that of another [14], also developed at Kent. The proxy is a plug-in resource to the privacy-enhancing infrastructure thus not requiring any changes to be made to either infrastructure. Since resources are able, given the proper

access, to detect context events it is capable of constantly keeping the external infrastructure up to date. Information is gathered from the external infrastructure by regularly querying it for information and looking for changes. This could be made more efficient if desired by for instance using two single-way proxies located at either infrastructure.

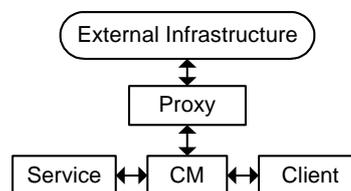


Figure 2: Integration with external infrastructure

Independent of how and where the proxy is deployed it is subject to the same access control mechanism as a service or client. Hence integration also has the advantage of allowing the privacy protection mechanisms to be applied to otherwise unprotected infrastructures. Although information would still flow freely within the unprotected infrastructure, control is gained over access from the outside. The integration has allowed for greater flexibility in specifying access from the outside.

## 7. PRIVACY PROTECTION

So far, we have said little about the privacy protection mechanism utilised in the infrastructure beyond acknowledging its existence. An important aspect of privacy is the existence of an access control mechanism providing control over the information flow. Though other areas such as trust, authentication, etc. are also important this section will focus on the access control mechanisms evaluated as part of the development of this infrastructure.

Our early work [15] used a Classification and Clearance Scheme (CCS) to control access. In this, each context element was assigned a classification level. The more sensitive, the higher it was classified, with a level of 0 indicating public information. Sites, services and other participants were each assigned an individual clearance value. The clearance given indicated the level of trust, and thus the maximum sensitivity of the information they could access. By default, unclassified information could not be accessed and recipients without clearance were limited to public information. Although working well with a small number of context elements and consumers the mechanism is limited by scalability problems.

The mechanism employed in the current infrastructure uses Role Based Access Control (RBAC) [16] instead as this allows for better scalability. The idea behind RBAC is to use roles to group permissions together and thus to make administration easier. Our implementation is based on the RBAC<sub>0</sub> [16], but with an important difference. For simplicity we have opted to use an automatic role activation mechanism. Agents will always be granted the best possible access given their current set of roles. In our RBAC mechanism the permissions that make up roles are implemented as lists of access controls. Each access control grants access to one context item. Any combination of read, write, and history access can be specified, including none of them. In the case of history access a limit can be set on how far back access is granted. For further customisation we have introduced the possibility of agents being assigned a personal permission that overrides the access granted by general roles.

## 8. FINDINGS

A number of interesting issues have arisen during the course of this work, some of which are discussed here.

Firstly, we began by assuming that people want to enjoy the same level of privacy online as they do offline. This assumption set the target level of protection that we have worked towards. However a survey of potential users appears to indicate that this may not always be true. The responses showed 97% being either very concerned or concerned about their privacy when online. In contrast, only 71% expressed similar levels of concern about their offline privacy. The sample size is currently small (31 responses) and shows a bias towards computer and Internet literate people, yet there appears to be a clear difference here: people seem to be more concerned about their privacy online than offline. Why this is the case is still unclear. Perhaps it is because people fear the unknowns that are encountered with new technology? Another hypothesis is that the greater concerns exist due to a lack of trust in online technology, its providers, and beliefs about the way personal information is handled online. It is not hard to imagine people at the receiving end of malicious online activities, e.g. spam and viruses, to lose faith in any online technology. Irrespective of the origin of such a heightened concern, the implications are not good. We may have missed the opportunity to form an early relationship of trust between the technology, its users and providers. This can potentially have a negative impact on the wider acceptability and desirability of ubiquitous technology. It may also mean that new technology will be more thoroughly inspected and possibly judged harder than traditional technologies. Although the sample size is limited, the clear distinction between the subjects' online and offline concerns suggests that this may indeed be a wider phenomenon.

Secondly, our survey also provides some interesting findings about the CCS and the RBAC mechanisms described above. We asked who should be allowed to access information regarding the respondent's location, activity, and contact details. We then asked them to express these preferences using both the CCS and RBAC mechanisms. The results showed that, on average, the setup of the CCS mechanism was 79% accurate whereas that for the RBAC mechanism was 87%. Thus the RBAC was more accurately setup to represent the subjects' preferences. This makes sense since the preferences that can be expressed using the CCS mechanism are limited to a single list of classifications. However, when asked how accurately they felt they had been able express their preferences 65% stated that they had been able to express their preferences either very accurately or accurately with the CCS mechanism whilst for the RBAC mechanism the figure was only 48%. Hence the subjects felt that they had been able to express their preferences more accurately using the CCS mechanism, when in fact the opposite was true. The subjects found it easier to express their preferences using the CCS mechanism. These results provide an indication that it may be difficult for people to see the real effect, at any particular state, of an access control mechanism. It also suggests that people may not always be aware of how accurately they have managed to setup their privacy protection and that ease of use may be misleading. Another worrying finding was that in the event of the access control mechanism being inaccurately setup it most often resulted in too high access being granted. In the case of CCS 70% of the inaccuracies caused too high access, the corresponding figure for RBAC was 68%. Thus in the cases where people mistakenly feel they have setup a protection mechanism accurately the consequences are likely to be a false sense of protection. This in

turn may cause people to mistrust the technology if they later come to realise that their expectations have not been met.

Finally, it is not uncommon for the performance of cryptography operations to be questioned in ubiquitous computing. So far we have described how the infrastructure uses cipher plug-ins to secure communications. The main plug-in used here employs both an asymmetric cipher, RSA [17], and symmetric cipher, AES [18]. The RSAAES plug-in, as we will refer to it, aims to strike a balance between security, manageable key distribution, and performance. It has been developed in Personal Java, just as the infrastructure, and uses the Bouncy Castle crypto package [19]. The experiences gained from the development and use the RSAAES plug-in in our infrastructure suggest that, although resources are limited, reasonable levels of performance may be achieved on the targeted devices. Benchmarks have shown, not unexpectedly, that the most time consuming operation to be the RSA key generation. The table below show the average time to generate a keypair with the respective key sizes on two different generations of handhelds as well as on a PC.

	512-bits	1024-bits	2048-bits
iPAQ 3660	1085	6147	44320
iPAQ 4150	636	4249	32615
Evo N1015	66	1235	12035

Figure 3: RSA key generation time in ms (average of 5 runs)

Given that the key size recommended by RSA Laboratories [20] for general use, at the time of writing, is 1024-bits then we end up in the region of 4-7 seconds. Keypairs can, however, be reused provided that the private key has not compromised, so the achieved performance is more than adequate. Of more importance is the time it takes to encrypt and decrypt a message since this will always delay communication. The table below shows the average time, over 5 runs, to encrypt or decrypt 15000 bytes of randomly generated data. The real message sizes used in the infrastructure vary from request to request with common examples about 1000 bytes. However message size increases with the complexity of a request, so to show that even complex messages can be handled, a larger message size was used in the benchmark. Throughout the benchmark the AES key size were fixed at its maximum length, 256-bits.

	512-bits	1024-bits	2048-bits
iPAQ 3660	573 / 921	592 / 1056	788 / 1851
iPAQ 4150	255 / 346	223 / 621	277 / 979
Evo N1015	20 / 30	18 / 54	20 / 140

Figure 4: RSAAES encryption/decryption time in ms

Progress in manufacturing faster devices makes significant differences. The iPAQ 3660 takes roughly twice the time to perform the same operation as the newer iPAQ 4150. Given the often short validity of context information, one could reduce the key size as appropriate. Any gain in performance would have to take into consideration, though, the shorter life span of the keypair. It should also be noted, as stated previously, that so far in our infrastructure the majority of the requests are significantly smaller. This reduces the encryption/decryption time, for example on the iPAQ 4150 the average time for the respective operations are 105/295 ms with a 1024-bit key and 1000 bytes of data. In comparison to other processes yet to be optimised, the impact on performance incurred by the use of the RSAAES plug-in is largely insignificant. In an ideal world the complete encryption/decryption process on both sides should take less than a second, including the fulfilment of the request itself. Although this is as yet hard to achieve between two handhelds, the overall performance is still reasonable.

