# A Privacy Preserving Overlay for Active Spaces

Jalal Al-Muhtadi, Raquel Hill, Roy Campbell

Department of Computer Science,
University of Illinois, Urbana-Champaign

(almuhtad@cs.uiuc.edu, rlhill@uiuc.edu, rhc@uiuc.edu)

## Abstract

Based on our experience in building and deploying ubiquitous computing environments, we identify some design guidelines for privacy. We also propose a novel mechanism that meets these guidelines and provides a practical and reasonable tradeoff between user privacy and environment automation and personalization.

## 1. Introduction

Ubiquitous computing is poised to have a profound effect on how humans interact with machines, physical spaces, services, everyday devices, and other humans. Ubiquitous computing envisions a world that features an overabundance of embedded processors, wearable computers, smart consumer devices, sensors, actuators, and digital communication, which are all tightly coupled to form an "active space." Active spaces enhance education, distance learning, remote collaboration, resource sharing, remote discussions, group activities, and scientific studies. Intrinsic to the notion of such environments and applications is the accurate capturing and processing of context information sensed through a variety of different mechanisms and sensors. Contextual information may include things such as users' locations, users' identities, group activities, environmental contexts (temperature, light and sound), group activities, location, orientation and status of devices and services, etc. Context awareness enables significant functionality to ubiquitous computing applications, users, resources and the ways they interact. It allows ubiquitous computing environments to tailor themselves according to users' preferences and expectations, and to reconfigure the available resources in a way that enriches users' experiences. However, the physical outreach of active spaces, and the plethora of sensors and devices implanted in ubiquitous environments could transform the ubiquitous computing environment into an all-encompassing surveillance system. Novel mechanisms are needed to give users control over their private information and how and when this information is disclosed.

In our research in ubiquitous computing we developed a distributed middleware system, called Gaia [1-3], which provides the necessary core services to construct, manage, and program general-purpose ubiquitous computing environments. We refer to these environments as active spaces. We use active spaces for regular graduate seminars [4], remote collaboration with off-campus researchers, group meetings, PhD defenses, and brainstorming sessions. During these usage scenarios, the active space utilizes context and fine-grain location information [5-7] to authenticate principals, activate roles, record attendance (if needed), and so on, while minimizing distractions and allowing the space to learn, adapt, and facilitate users' interactions with their surroundings. As we move towards deploying this technology in a new "smart" building, and with the advent of highly accurate location tracking systems based on UWB [8], we believe that providing good privacy guarantees is key to real-world deployment of this technology.

In this paper, based on our experiences with deploying active spaces, we address users' need for identity and location privacy. We briefly discuss design guidelines for preserving users' privacy. We use these guidelines to propose some privacy preserving solutions. Finally, we present some major challenges to ensuring privacy and some future directions for privacy research.

## 2. Design Guidelines

We deployed active spaces in different usage scenarios, which allowed us to assess privacy needs and concerns of the active space users. Furthermore, we identified some design guidelines that help to establish a practical tradeoff between users' privacy on one hand, and context awareness and space automation on the other. These guidelines address issues of total anonymity, decoupling identity and location, and customizable privacy. We briefly mention these guidelines in this section.

### 2.1 Total Anonymity

A key concept in ubiquitous computing systems is the ability to capture relevant contextual information and act upon it, which allows the environment to customize itself and adapt according to the current situation. We define total anonymity as the state of being unknown for an infinite amount of time. With respect to a ubiquitous computing environment, this means that neither a user's identity nor location can be inferred from context information. Total anonymity in a ubiquitous computing setting is not feasible because context information may reveal information about a user's identity, location, etc. Many scenarios, like classrooms and seminars, may require some level of authentication, attendance recording, or authorization (e.g. only the presenter has access to the slideshow application). Furthermore, access control mechanisms and security audits become useless if total anonymity is permitted. For these reasons, we must use mechanisms that strike the desired balance between preserving privacy and enabling some kind of value-added services for users based on context information.

### 2.2 Decoupling Identity from Location Privacy

Cooper et al. [9] identify three kinds of privacy: content, identity, and location. *Content privacy* is concerned with keeping data or content private. *Identity privacy* is concerned with hiding the identity of the user. *Location privacy* is concerned with hiding the location of the user. Content privacy relates to confidentiality and, in many cases, it can be achieved through encryption. We propose that decoupling one's identity from one's location information provides the person with a better level of privacy. For example, in some of the scenarios that we explored, there is a need to perform authorization checks before an individual can invoke a particular service. In traditional systems, some component within the infrastructure authenticates users before they access privileged services. However, because the infrastructure is expected to be context and location aware, the identity, location, and the state of the user are all revealed. Nevertheless, if information decoupling is executed properly, so that the information required by certain components in the system is well-defined, it is possible to allow these components to obtain only the amount of information that is required for successful operation. For example, the authentication service may only need to verify that an entity belongs to a specific role within a given time interval. In this case, the authentication service does not need to know the exact identity or location of that entity. We exploit this feature to provide a configurable trade-off between location/identity privacy on one hand and security and value-added services on the other.

### 2.3 Customizable Level of Privacy

We believe that it is necessary to propose a system of relative privacy that allows a user to choose the level of desired privacy. The level of privacy is derived from the users' interactions with the ubiquitous environment, and the amount of personalization or value-added services that a user desires. Furthermore, users should be given the ability to fine-tune the privacy level. This gives users the ability to tweak their privacy levels according to their preferences or their current situation. In practice, users may choose to give up some information in return for value-added services, but users should control when they share this information to access these services.

## 3. Proposed Privacy Mechanisms

While context awareness will enable significant functionality to ubiquitous computing, the ubiquitous environment may pose serious threats to the basic privacy rights of users unless suitable privacy protection mechanisms are introduced. It is essential to provide novel protocols for obfuscating the identity and/or location of the users, resources, and applications in the system. Onion Routing [10], Crowds [11], and Mist [12] are examples of protocols that introduce some level of anonymity.

We identify two different models for preserving privacy, an *infrastructure-based model* and an *ad-hoc model*. The *infrastructure-based model* for preserving privacy assumes that a trustworthy infrastructure exists, and users can set policies to identify who, what, when and under what context their location information is disclosed. To meet this requirement, it is possible to use the dynamic roles and the context aware security policies supported by Gaia [13] to allow users, administrators, and service

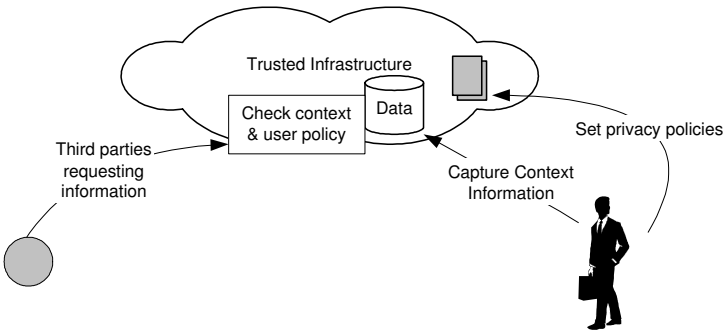


Figure 1: Trusted Infrastructure Model

providers to define access control policies on how and when their location information can be disclosed. This model is illustrated in Figure 1.

The assumption that a trusted infrastructure will protect private information is inadequate in some settings. Furthermore, the infrastructure may be prone to insider attacks or to policy changes that could affect the privacy of users. The other model for privacy is an *ad hoc model*; where the location of a user, the role(s) he assumes, and his identity is decoupled to provide a higher degree of anonymity. The decoupling is achieved by (1) only revealing the least amount of information needed to access the required resource, and (2) introducing a level of indirection, so that no single component in the system can capture all the information, unless enough number of components collude. We describe this in more details shortly.

In order to enable this ad hoc model for decoupling identity privacy from location privacy, we utilize Mist [12, 14] an overlay network for communication that decouples location privacy and identity privacy in ubiquitous computing environments. We select Mist over other anonymous communication protocols because it features location and identity decoupling, provides two-way anonymous communication, and targets ubiquitous computing environments. We give a brief overview on how Mist works. Mist consists of a privacy-preserving hierarchy of *Mist Routers* that form an overlay network. This overlay network allows users to communicate privately. The Mist Routers route communication packets using a hop-by-hop, handle-based routing protocol with limited public-key cryptography, thus making communication untraceable by eavesdroppers and intermediate routers. At the leaf level of this hierarchy, special Mist Routers called "*Portals*" are introduced. These devices are installed in the different Active Spaces. Portals are devices capable of detecting the presence of people and objects through the use of various location sensors; however, they are incapable of positively identifying the users. To successfully, decouple location and identity, these portals should not be given the user-badge ID assignments, for example.

The actual authentication or role name activation of a user takes place above the portal level in the hierarchy; therefore the actual physical location of the user cannot be deduced. More specifically, it takes place at a special Mist Router referred to as a "*Lighthouse*." Only a Lighthouse is able to authenticate the user or activate some of his roles. However, the Lighthouse does not know the actual physical location of the user (thanks to the hop-by-hop routing protocol). To illustrate, in Figure 2, Alice, who is in active space 3, is detected by the Portal in that space. The Portal only detects Alice's badge ID (or other information embedded into other devices that Alice is carrying or wearing) however, this information alone is insufficient to indicate that this is actually Alice. The CS Building Mist Router is designated as Alice's Lighthouse. A secure channel between Alice's devices and her Lighthouse is established, and traverses the Portal, node 1, node 2, and finally node 3. Encryption is employed to prevent private information from leaking. Instead of having a traceable source and destination addresses, packets over this secure link are routed through the use of handles that are valid only over a single hop. The intermediate nodes translate an incoming handle to an outgoing one. Thus, intermediate Mist Routers can only route to the next hop correctly, but do not know the actual destination or source. The true location of Alice can be found only if all intermediate Mist Routers collude. Note that in the example, Alice's Lighthouse can only infer that Alice is located somewhere within the CS building.

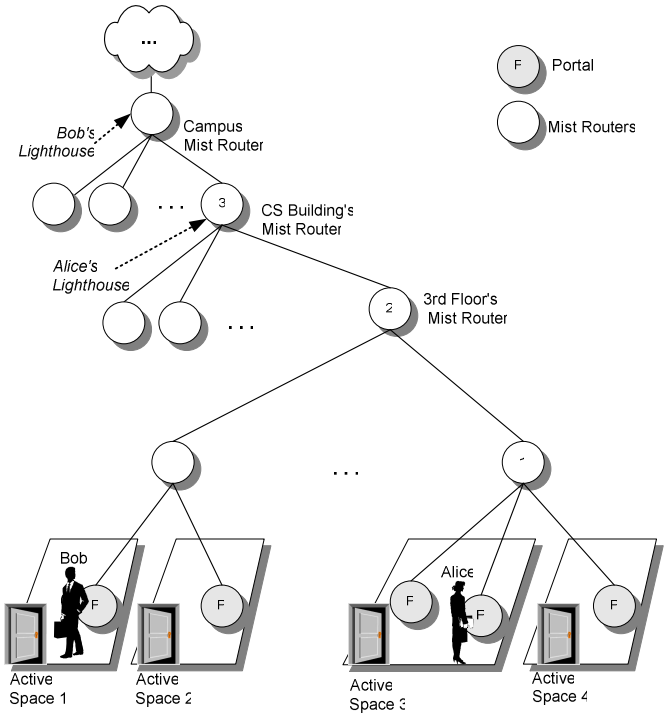


Figure 2: The Mist Communication Protocol

Our authentication and authorization protocol enables users to access privileged services that require some level of authentication while preserving their location and identity privacy. The protocol extends Kerberos to support user devices and utilizes the privacy decoupling that Mist provides. In the extended version of Kerberos, no single component is entrusted with all the private information of users.

An active space security service exists for every *Active Domain*. An Active Domain is a collection of active spaces, and the interconnecting networks, which are managed by a single administrative authority. These domains resemble Kerberos "Realms." Like Kerberos, the security service consists of three components. The first component is the AS (Authentication Service), which provides a single sign-on point for the Active Domain, using any devices and gadgets the user currently possesses. The TGS (Ticket Granting Service) issues "tickets" that can be used by the user to access available services in that space. These tickets are signed, as a protection against tampering, using the private key of the TGS. To decouple sensitive information further, and limit linkability, the system incorporates different TGS components for different services. For example, there is a TGS for the "printing service" and a TGS for the "lighting service" for instance. Finally, a database is maintained that contains necessary information for the authentication of users within the Active Domain.

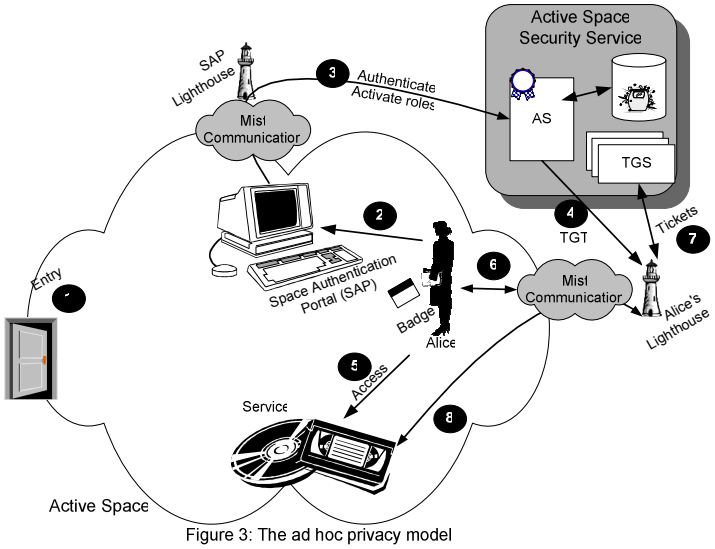

Figure 3: The ad hoc privacy model

The protocol is illustrated in Figure 3. We give a brief overview of the protocol. Within every active space, we assume the existence of one or more "*space authentication portals*" (SAPs). These are special types of Portals that can be located at the entrance of an active space, or other convenient places. The SAP will feature a collection of wireless and wired base stations and device readers that enable users to authenticate with the active space using any authentication devices they are carrying or wearing. Users enter an active space (step 1 in Figure 3). To gain access to privileged services, users can authenticate through the space authentication portals (SAP) (Step 2). To achieve privacy, the SAP itself does not have sufficient information to authenticate users. However, it has a Lighthouse through which it can communicate with the authentication service (step 3). Mist communication is used here to prevent the authentication service from determining the exact physical location of the authenticated user. Through its Lighthouse, the SAP contacts the authentication service with a set of authentication requests, each representing a different authentication device. Upon successful authentication, the AS, like Kerberos, issues a ticket granting ticket (TGT) for that user (step 4). Recall that in Mist, every user has a Lighthouse that stores his relevant information. The TGT issued for a user can be stored on his personal device (if available) or can be stored in his Lighthouse. The TGT in our system is a cryptographic data structure that contains one or more roles that the user can utilize to access a certain service (step 5). Upon accessing a protected service, the user requests a ticket to access this service, by going through the Mist to his Lighthouse. This TGT is stored at the user's Lighthouse. In step 7, using the TGT stored at the Lighthouse for the target user, the Lighthouse can communicate with the TGS and request tickets to access the required service. The TGS issues the necessary cryptographic tickets. These tickets do not contain any references to the real name or identity of the owner; they just incorporate an unforgeable pseudonym or contain a role name for the user that allows him to access the service without revealing his exact identity. Using the information in these tickets, the service can make a decision whether to authorize the user or not (step 8).

To illustrate the decoupling of information, Table 1 shows what information is revealed to which component in a scenario where a user is trying to access a printer in a particular room by getting a ticket from the printing TGS.

**Table 1: Information decoupling in a printing scenario**

| *Component* | *Information revealed* |
|---|---|
| Authentication Service | The identity of the user and his role names. No information on current location, context, or service interaction is available. |
| Printing Service TGS | Relevant role name and a "hint" on possible user location (i.e. the user must be using one of the printers managed by this TGS). No information on the exact location or identity. |
| Printer | Whether the user is authorized to print or not and his location. No information about his identity or roles. |

## 4. Challenges

Privacy in ubiquitous computing environments poses many challenges that need to be addressed. We identify some of the challenges that we are currently investigating and trying to address in our active space infrastructure.

### 4.1 Modeling Levels of Privacy

Providing a user with the ability to specify and maintain a certain level of privacy poses many challenges. To address this design concern, we must define a model of privacy for the environment. Among other things, this model should illustrate the various levels of privacy, define how the levels relate, specify policies per privacy level, identify the actions or functions that enforce the policies, define the relationships between user and administrative policies and identify user actions that may change the level of privacy.

One major challenge to defining the relationships between user and administrator polices, is that of balancing user needs for privacy with administrative needs for linkability. Linkability is the capacity to correlate actions to a specific user. Linkability is especially important when the ubiquitous infrastructure is under attack or when other unsafe conditions exist.

### 4.2 Safety Policies

Many issues may affect user privacy. An issue of particular interest to our research is understanding how safety policies affect user privacy. We feel that in many instances the need for safety may override that of privacy. For example, when a ubiquitous infrastructure experiences attack, administrators may need to know the exact identity of all users in order to identify the attacker and limit the damage to the infrastructure. Furthermore, in a medical scenario, a doctor who is not authorized to view a patient's medical records under normal circumstances, may need to access these records when a medical emergency occurs (e.g. the patient's attending physician is not available during an emergency). We must understand how to identify the scenarios in which safety policies should preempt privacy policies. We must also define protocols that securely escalate the safety policy over the privacy policy.

### 4.3 Information Flow Analysis

Another interesting challenge is that of information flow analysis. This involves understanding the relationship among data that is collected by administrators or other users of a ubiquitous computing environment and the information that can be inferred from this data, particularly when different data can be gathered from multiple sources. Thus, given data X and Y, can we infer a user's identity, location, or role?

## 5. Conclusion

In this paper we identified some design guidelines and outlined some challenges for privacy in ubiquitous computing environments. In addition, we briefly described our deployed infrastructure that enables some level of location and identity privacy while allowing users to benefit from the context-awareness and automation that a ubiquitous computing environment provides.

## 6. References


[1] M. Román, C. K. Hess, R. Cerqueira, A. Ranganathan, R. H. Campbell, and K. Nahrstedt, "Gaia: A Middleware Infrastructure to Enable Active Spaces," *IEEE Pervasive Computing (accepted)*, 2002.

[2] M. Roman and R. Campbell, "GAIA: Enabling Active Spaces," presented at 9th ACM SIGOPS European Workshop,, Kolding, Denmark, 2000.

[3] M. Roman, C. K. Hess, R. Cerqueira, R. H. Campbell, and K. Narhstedt, "Gaia: A Middleware Infrastructure to Enable Active Spaces," *IEEE Pervasive Computing Magazine*, vol. 1, pp. 74-83, 2002.

[4] M. Roman, J. Al-Muhtadi, B. Ziebart, R. Campbell, and M. D. Mickunas, "System Support for Rapid Ubiquitous Computing Application Development and Evaluation," presented at System Support for Ubiquitous Computing Workshop (UbiSys '03) in conjunction with UbiComp '03, Seattle, WA.

[5] J. Al-Muhtadi, A. Ranganathan, R. Campbell, and M. D. Mickunas, "Cerberus: A Context-Aware Security Scheme for Smart Spaces," presented at the First IEEE Annual Conference on Pervasive Computing and Communications (PerCom 2003), Fort Worth, Texas, 2003.

[6] A. Ranganathan and R. H. Campbell, "A Middleware for Context-Aware Agents in Ubiquitous Computing Environments," presented at Middleware 2003 (submitted) http://choices.cs.uiuc.edu/~ranganat/Pubs/MiddlewareForContext.pdf, 2003.

[7] A. Ranganathan, J. Al-Muhtadi, S. Chetan, R. Campbell, and M. D. Mickunas, "MiddleWhere: A Middleware for Location Awareness in Ubiquitous Computing Applications," presented at 5th International Middleware Conference (Middleware 2004) (accepted), 2004.

[8] UbiSense, "Local position system and sentient computing." http://www.ubisense.net/.

[9] D. A. Cooper and K. P. Birman, "Preserving Privacy in a Network of Mobile Computers.," presented at IEEE Symposium on Research in Security and Privacy, 1995.

[10] M. Reed, P. Syverson, and D. Goldschlag, "Anonymous Connections and Onion Routing," *IEEE Journal on Selected Areas in Communication, Special Issue on Copyright and Privacy Protection*, 1998.

[11] M. Reiter and A. D. Rubin, "Crowds: Anonymity for Web Transactions," *ACM Transactions on Information and System Security (TISSEC)*, vol. 1, 1998.

[12] J. Al-Muhtadi, R. Campbell, A. Kapadia, D. Mickunas, and S. Yi, "Routing Through the Mist: Privacy Preserving Communication in Ubiquitous Computing Environments," presented at International Conference of Distributed Computing Systems (ICDCS 2002), Vienna, Austria, 2002.

[13] G. Sampemane, P. Naldurg, and R. Campbell, "Access Control for Active Spaces," presented at the Annual Computer Security Applications Conference (ACSAC), Las Vegas, NV, 2002.

[14] J. Al-Muhtadi, R. Campbell, A. Kapadia, D. Mickunas, and S. Yi, "Routing through the Mist: Design and Implementation," UIUCDCS-R-2002-2267, March 2002.