Territorial Privacy in Ubiquitous Computing

Bastian Könings and Florian Schaub
Institute of Media Informatics
Ulm University, Germany
{ bastian.koenings | florian.schaub }@uni-ulm.de

***Abstract*—Smartphones define a trend towards increasing combination and integration of sensing capabilities with almost ubiquitous inter-connectivity. Resulting location-based services and context-aware applications will benefit users by adapting better to the user application needs. However, there is a lack of effective means for controlling privacy in such systems which will likely increase further with future ubiquitous computing systems. Territorial privacy is a concept that moves away from the information-centric view in traditional systems to a context-centric approach. In this paper, we define and model territorial privacy in the context of ubiquitous computing. We further discuss potential observers and disturbers in our model and provide an overview on how territorial privacy can be controlled in different environments, ranging from personal to public.**

I. Introduction

Nowadays, personal gadgets and other devices communicate wirelessly with each other and also with services in the cloud. This inter-connectivity enables new and exciting applications. Location-based services only mark the beginning of a new range of future applications centered around context-based information services and seamless inter-connectivity. At the same time, the user's ability to control and manage the information exchange seems to be diminishing.

A prime example for this issue are smartphones, such as Android or iOS devices. These devices pack a multitude of sensors. Ranging from obvious ones like the microphone and camera to localization systems based on multiple location sensors (GPS, WiFi, cell of origin) to integrated sensors (proximity sensor, light sensor, accelerometer, gyroscope, digital compass). All these sensors combined provide detailed information about the location and environment of device and user, as well as the handling context of the device. Applications can utilize this information to provide enhanced context sensitivity and user adaptation. At the same time privacy issues arise from the combination of detailed information about the user context and almost ubiquitous Internet connectivity. Potentially, an application could pass on any sensor output to cloud-based services and their providers. What information actually leaves the device is usually invisible to the user. For example, iOS, the iPhone operating system, indicates with a little arrow when an application is using the localization system to obtain the current location context. Another symbol indicates when an application uses data communication services. However, if, to what extend, and to whom a particular application communicates the location context outside the device remains opaque. Current mobile operating systems provide some privacy controls to address these issues but fall short of their goal. iOS provides a control interface to grant and deny applications access to location context and also warns the first time an application wants to access the device location before granting permanent access.[1] The Android operating system provides a fine-grained system of security permissions to control which application has access to which system resource, including sensors. However, requests for such permissions are statically declared in an application's manifest file.[2] When installing an application, the user is presented with an overview of the application's required system resources. At that time, the user has to make an all-or-nothing decision [1], without knowing for what purpose an application requires a certain sensor and what is done with the acquired information.

But how should privacy controls be designed in such systems to be effective? A fine-grained access control matrix that governs application specific access to individual sensors maybe desirable but lacks usability and scalability, especially on smartphones with continuously increasing numbers of applications and sensors. Furthermore, such approaches provide inadequate support for context-based dynamic access control. A paradigm shift in terms of privacy controls is required to ensure that users can effectively manage their privacy and that the privacy settings match the user's intentions closely in a specific context and situation. Thus, we advocate to model privacy controls in mobile environments based on the user's context and location in contrast to the information-specific approaches prevalent in traditional computing.

Currently, the described privacy issues exist mainly in the smartphone and mobile computing environment. The combination and integration of sensing and communication capabilities will likely increase in the future and so will privacy implications. In envisioned ambient assisted living and ubiquitous computing scenarios people will be surrounded by hundreds of smart sensors that will communicate with each other autonomously. While mostly filed under "visions of the future", especially ambient assisted living systems are already being field tested.[3] Increasing surveillance of the public with advanced surveillance and sensing systems will also affect personal privacy and civil liberties in the future.

In this paper, we discuss the rather old concept of territorial privacy and how it can be potentially leveraged to address the outlined challenges for privacy control. We will focus our

[1] http://support.apple.com/kb/HT1975
[2] http://developer.android.com/guide/topics/security/security.html
[3] e.g. in the MonAMI project: http://www.monami.info/

8th International Conference on Wireless On-Demand Network Systems and Services (WONS 2011), Bardonecchia, Italy, January 2011, IEEE
Author version. Definitive version at http://dx.doi.org/10.1109/WONS.2011.5720177

discussion on future systems with sensing and communication abilities. We first discuss the concept of territorial privacy and provide a definition (Sec. II), before proposing a model for territorial privacy in ubiquitous computing systems (Sec. III and IV). We discuss approaches for controlling territorial privacy in Section V. Section VI concludes the paper.

II. Territorial Privacy

Current privacy research mostly focus on mechanisms for protecting sensitive personal information, like user profiles of web applications, bank data, or location information. While addressing information privacy concerns might be a sufficient approach for most common ICT domains, it remains only one aspect of privacy in the context of emerging ubiquitous computing systems. The pervasive nature of such systems equipped with multiple sensors, actuators, processors and wireless communication capabilities, constitute a new facet of privacy needs. Privacy issues will more than ever emerge from the physical environment of a user. For example, by gathered information from sensors in the user's proximity or by actuators disturbing the user with intervening actions. This fact leads to a different expectation of privacy. We call this aspect of privacy *territorial privacy*, which refers to a more traditional privacy understanding such as *"the right to be let alone"* [2] or being in *"a state in which one is not observed or disturbed by others"* [3].

In contrast to an information centered approach where privacy is controlled by protecting particular information, the concept of territorial privacy aims to provide a more user centered approach. Privacy decisions of users are often based on their physical or spatial context. Traditionally, achieving privacy was as easy as going into a room and closing the door, thus avoiding undesired observations and disturbances by setting up physical boundaries. Whereas in a traditional scenario the physical boundaries of a room would also mark the boundaries of a user's private territory, this situation will drastically change in future ubiquitous environments. Invisible embedded sensors, actuators and in particular wireless communications could widen the boundaries of a private territory far beyond its physical boundaries. As a consequence the ability to perceive and control who is observing or disturbing a user in her private territory will decrease or even cease to exist. Thus, perceiving and controlling those new *virtual* territorial boundaries with the ease of traditional privacy controls outline the main goals of the territorial privacy concept.

A. Defining Territorial Privacy

Defining the term privacy remains one of the most intractable problems in privacy research. Therefore, our intention is not the provision of a widely applicable and comprehensive definition of this term, we rather aim at providing a specific definition of territorial privacy as it is used in our context and in the context of ubiquitous computing systems. We base our definition on the traditional notion of privacy as *"a state in which one is not observed or disturbed by others"* [3].

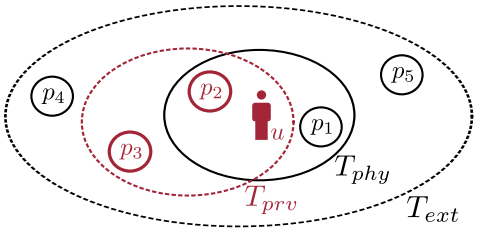

Fig. 1. A user and different participants p_1-p_5 in a physical territory T_{phy}, extended territory T_{ext} and private territory T_{prv}.

We extend this definition by several means. First, territorial privacy is realized by protecting a private territory T_{prv} of a specific user u. In a territory, different entities could be present, that we call participants P of that territory. A participant can either be an observer or a disturber. Observers are entities receiving realtime information gathered in the user's proximity. Disturbers are entities which could actively intervene in a user's environment. Furthermore, participants can be part of the user's physical territory or participate virtually in the user's extended territory T_{ext}. The boundaries of the private territory are defined by a subset of only desired participants P_d, thus $T_{prv} \subseteq T_{ext}$. We will provide a more detailed description of these terms in the following section.

With respect to these assumptions we define territorial privacy as *a state in which a user u is able to perceive his extended territory T_{ext} and enforce his private territory T_{prv} by allowing only desired participants P_d.*

III. A Model for Territorial Privacy

Respecting the former definition we will introduce a basic model for territorial privacy based on the concepts of *territories* and *participants*. An overview of our conceptual model is depicted in Fig. 1.

A. Territories

We distinguish between three territories: The user's *physical territory* T_{phy}, the user's *extended territory* T_{ext} and the user's *private territory* T_{prv}. The physical territory refers to environments of a user, which are demarcated by physical boundaries like walls or doors. Thus, a physical territory might refer to a house, room, or a car, and it includes all physically present entities. The extended territory is the virtual expansion of the physical territory to encompass remote entities connected via communication technologies. Therefore, it applies $T_{phy} \subseteq T_{ext}$. Finally, the private territory is a subset of the extended territory, but not necessarily a superset of the physical territory, thus it applies $T_{prv} \subseteq T_{ext}$ but not necessarily $T_{prv} \supseteq T_{phy}$. This means that a user may have the ability to exclude physically present entities from his private territory. For instance, a user might want to disable a surveillance camera in the same room.

B. Participants

Participants are all entities that are either physically or virtually present in a user's extended territory and are able

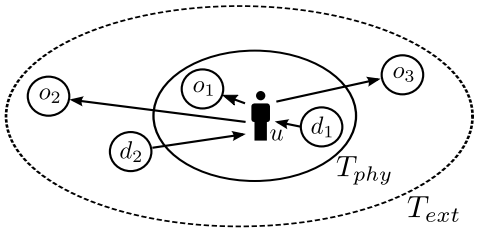

Fig. 2. A user's physical territory T_{phy}, extended territory T_{ext} and participating observers o_1-o_3 and participating disturbers d_1, d_2.

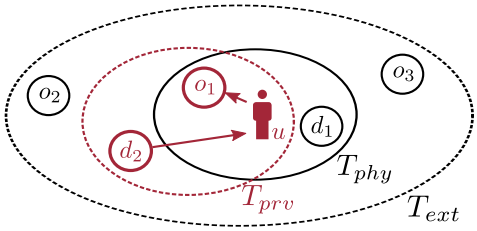

Fig. 3. A user's physical territory T_{phy}, extended territory T_{ext} and private territory T_{prv} with desired observer o_1 and desired disturber d_2.

to observe or disturb the user in a certain way. A participating entity may either be a human, a computer system, or a sensing device. We distinguish between participating observers O and disturbers D with $O \cup D = P$. Observers and disturbers are connected by directed edges from and to the user, respectively. Fig. 2 shows a user's extended territory with participating observers o_1 to o_3 and disturbers d_1, d_2. Disturber d_1 and observer o_1 are physical participants, while o_2, o_3 and d_2 are virtually participating in the user's territory. A further discussion and classification of possible observers and disturbers will be given in the following Section.

Knowing the extended territory T_{ext} of a user, the goal of territorial privacy is now to protect the private territory T_{prv} by excluding undesired observers and disturbers. In Fig. 3 the private territory consists of desired observer o_1 and desired disturber d_2. All other participants are excluded from the user's private territory. A more detailed discussion and formalization of this demarcation process is given in form of an observation model in [4].

IV. Classification of Participating Entities

In ubiquitous environments a user might be observed or disturbed by different entities, like embedded sensors, autonomous devices, but also other persons. In general, all entities can either participate physically or virtually in a user's territory. A physically participating entity is located inside the user's physical territory. Thus, the entity is in physical range of the user, e.g., in the same room, and may therefore be physically perceived by the user, either visually or acoustically. A virtually participating entity is located outside the user's physical range, but connected by communication technologies to a physical participant.

The following subsections will provide a classification of participating entities in terms of observers and disturbers.

A. Observers

We refer to a participating entity of a user's territory as an *observer*, if this entity receives any realtime information about the user gathered in his current context. We distinguish between *active* and *passive* observers. Active observers gather realtime information about a user. Passive observers are entities, which receive this gathered information from active observers. We further classify active observers into *humans*, *ambient sensors*, *body sensors*, *personal device sensors*, and *personal device detectors*.

1) Humans: A person, which is in range of a user's physical territory can observe the user either visually or acoustically. For instance, a person in the same room as the user can see what the user is doing or hear what the user is saying.

2) Ambient Sensors: Ambient sensors that are installed or embedded in the physical environment, can gather information in many ways. The most obvious sensors of this kind are cameras and microphones. But also other sensors, e.g., sensors for motion, temperature or light, are considered to be active observers. Some might argue that many of those sensors are not related to privacy as they do not gather privacy sensitive information. However, even very simple sensors like brightness or temperature sensors, are gathering information, which in a larger context might lead to a privacy invasion of a user. For example, the historical analysis of light and temperature values in an office room, might allow to infer the presence of a person. Thus, a user should always be aware of all kinds of sensors that observe him in his current territory.

3) Body Sensors: Body sensors are placed in the proximity of the user's body, e.g., wearable sensors in clothes measuring the temperature or pulse of a user. Also implanted sensors, which measure blood pressure or the blood sugar level fit into this category.

4) Personal Device Sensors: Personal device sensors are installed in personal devices like PDAs, smartphones, or watches. The primary role of those sensors is estimating the state of a device. For example, accelerometers, digital compasses, or proximity sensors measuring the movement, direction or position of a device, or GPS units measuring the device's location. Even if those sensors do not directly sense information of a user, the gathered information can be assigned to the device owner, if the ownership can be unambiguously determined. Once the association of a device to its owner is known, gathered information from device sensors will provide information about this specific user, like his movements or location.

5) Personal Device Detectors: One of the most privacy threatening sensor classes are personal device detectors. Those are sensors which detect the presence or location of a personal device by wireless communication technologies. For example, a smartphone is typically equipped with wireless communication modules for GSM, WiFi and Bluetooth. Newer devices might also be equipped with near field communication (NFC). These communication technologies allow the detection or even precise localization of a device, even if the device is not actively communicating. For example, a GSM base

station can receive the signal of a mobile device from several kilometers away. This can be used by the network provider for localization, which can only be avoided if the device is completely turned off. In a similar way, a WiFi access point can detect nearby WiFi devices in a range of about hundred meters. Based on the communication technology, a device can be identified by an unique ID, e.g., the IMEI number of GSM cellphones, or the MAC address of WiFi modules. If this identifier can be matched to a real user, these device detections and localizations imply observations of this user.

The main issue of device detectors is their wireless nature. For a user it is very hard to perceive the presence of these observers as they are not limited by physical boundaries. Further, the prevalence of smartphones and other personal devices, like music players, or even sport shoes [5] with continuously active wireless communication modules, makes it even harder to avoid this form of observation.

B. Disturbers

We call an entity which is actively intervening in a user's physical territory a *disturber*. We distinguish between *physical* and *remote* disturbers. Remote disturbers are remotely controlling a physical disturber, which is located in physical range of a user. A disturbance can have different dimensions, which are *visual*, *acoustical*, or *motion*. We further classify physical disturbers in *humans*, *ambient output devices*, *personal devices*, *autonomous devices*, and *remote controllable devices*.

1) Humans: Another person may disturb a user in his private territory, by physically passing the borders of the territory, e.g., entering a room by walking through a door. Further, a physically present person may acoustically disturb a user by making noise or talking.

2) Ambient Output Devices: Ambient output devices refer to devices that are installed or embedded in the user's physical environment and can provide visual or acoustic output. For example, a switched on screen, or blinking status LEDs may visually disturb a user. A acoustic disturbance may be caused by the sound of loud speakers.

3) Personal Devices: A personal device could have similar output modalities, like a display, status LEDs or speakers. But also tangible outputs, e.g., vibrations, may be available that could be a disturbance to the user in his current territory.

4) Autonomous Devices: Autonomous devices represent a special class of disturbers as they can disturb a user in several ways. For example, a vacuum cleaning robot may disturb the user in all three dimensions, by moving around, making disturbing sounds and flashing lights. The main issue with autonomous devices is that they are harder to control than other devices.

5) Remote Controllable Devices: The last class of disturbers refers to devices which can be remotely controlled by other entities. For example, in a smart home the lights, heat, or curtains might be remote controllable, which may lead to a disturbance of the user if activated in his presence.

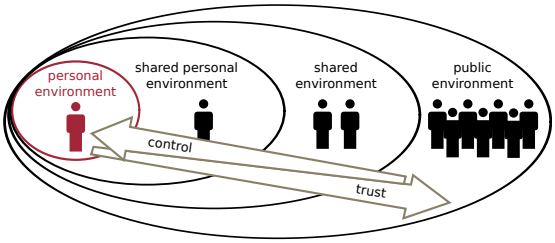

Fig. 4. Environments from personal to public with decreasing user control and increasing required trust.

V. Approaches for Control

The user's ability to control the boundary of the private territory directly relates to the user's ability to exclude undesired observers and disturbers from the user's extended territory spanning physical and virtual environments. In [4] we presented a graph-based process for excluding undesired observers from the private territory. When realizing such control processes in practical systems it becomes a key factor in which type of environment the user is currently active. The user's home is easier controllable than public and shared environments. We distinguish four different environment types:

- The *personal environment* constitutes the user's most private sanctuary. The user has full control about the environment and entities within. An example for a personal environment is a personal bed room or home office.
- The *shared personal environment* extends the personal environment and is shared with others. For example a family home. Entities are still controllable but the mediation between interests of different users in the same environment needs to be taken into account.
- A *shared environment* is accessible to a restricted number of users, whereby an individual has only limited control over the environment. An example is the work office, where a user shares work space with co-workers. Infrastructure services maybe under control of the building management rather than the users.
- The *public environment* is open to anyone but out of control of individuals. Governments and public service providers exercise control in such environments.

The personal and public environments are similar in definition to those of Daskala and Maghiros [6]. However, we explicitly distinguish two types of shared environments compared to one group environment [6] to reflect the user's level of control over these environments.

Furthermore, these environments should not be seen as discrete values, they simply serve as categories which subsume many environments with related configurations on a continuum from personal to public environment. Fig. 4 shows how these environments relate to each other. In general, the more public an environment, the harder it is to control the territorial privacy of individual users, resulting in larger territorial privacy boundaries, or privacy bubbles [6].

A. Personal and Shared Personal Environments

In personal and shared personal environments system components are under the users control, e.g., smart home and home automation systems. Therefore, a territorial privacy management component can be instantiated that takes user preferences in terms of territorial privacy into account to adapt system behavior accordingly. System entities in the personal environment are considered trustworthy. However, unobtrusive feedback can provide reassurance to the user that his or her territorial privacy boundary is being respected. In shared personal environments preferences from multiple users have to be taken into account to define the system state.

Some concepts relating to control of territorial privacy in personal (shared) environments have been proposed.

ATRACO [7] uses privacy policy ontologies to govern access to a user's private territory and activity sphere. Territory access requests for observers and disturbers are handled by a territorial privacy controller which matches the requests with the privacy policies in the user's profile. Gong et al. [8] propose that a user's personal device broadcasts the user's privacy preferences as policies into the environment. The system then interprets the policies and adapts data gathering activities accordingly. Geo-fencing [9] is a proposed approach to control the extend of the virtual territory by limiting WiFi coverage of the home network to physical boundaries.

B. Shared and Public Environments

In shared and public environments the individual's control capabilities are reduced. Therefore, the amount of trust required to be placed in third parties increases. The individual user must trust the environment that privacy is respected. Therefore, territorial privacy mechanisms in shared and public environments need to be designed from a different perspective than those in personal environments. The environment and the perception of the environment play an important role in establishing trust in the privacy awareness of deployed systems. For example, a public square overladen with surveillance cameras does not suggest a privacy-friendly environment. Companies, public service providers, and government agencies can actively shape the perception of shared and public environments by supporting privacy awareness and providing feedback to users. One possible approach are privacy beacons as suggested by Langheinrich [10]. Surveillance cameras as well as other observers could send out privacy beacons to inform users that they are being recorded, for what purpose, who long the data is retained and if privacy enhancing mechanisms are in place, e.g., face blurring or blocking in CCTV footage. This approach could also be combined with privacy policies broadcast by the user's device [8], observing systems should then provide feedback if those policies are respected and implemented.

VI. Conclusions

In a world, in which physical and virtual environments start to converge due to the ubiquitous presence of smart and interconnected sensor technology, controlling privacy with an information-centric view becomes increasingly difficult. The concept of territorial privacy offers a more usable and intuitive view on controlling individual privacy by using location and context information as a basis for privacy decisions. The presented model for territorial privacy is user centric and encompasses desired and undesired observers and disturbers. Our classification of potential observers and disturbers facilitates the development of suitable control mechanisms for territorial privacy in smart environments. The implementation of these control mechanisms depends on the user's current environment.

In a (shared) personal environment, like the home, comprehensive privacy systems can be realized that adapt the behavior of system components to the user's privacy preferences. Dynamic adaptation to changing privacy requirements and taking multiple individuals into account pose interesting challenges for the future. In shared and public environments an emphasis needs to be placed on informing the user about potential privacy implications, while services under the user's full control should still respect privacy preferences. Significantly higher trust in infrastructure providers is required in such environments. Individual users should be made aware of the actual privacy level they can reasonably expect in a given situation.

Currently, privacy systems often focus either on personal or public environments, while some initial approaches are also applicable in both scenarios. However, disturbers are not explicitly addressed by current approaches. We are currently developing a comprehensive approach for territorial privacy that enables users to control and monitor their privacy in the presence of observers and disturbers while seamlessly moving between different environments.

References

[1] A. Shabtai, Y. Fledel, U. Kanonov, Y. Elovici, S. Dolev, and C. Glezer, "Google android: A comprehensive security assessment," *IEEE Security & Privacy Magazine*, vol. 8, no. 2, pp. 35–44, 2010.

[2] S. D. Warren and L. D. Brandeis, "Right to privacy," *Harvard Law Review*, vol. 4, pp. 193–220, 1890.

[3] The Oxford English Dictionary, *"Privacy" Definition*, 2nd ed. Oxford University Press, USA, 2005.

[4] B. Könings, F. Schaub, F. Kargl, and M. Weber, "Towards territorial privacy in smart environments," in *Intelligent Information Privacy Management Symposium (Privacy 2010), AAAI Spring Symposium Series*, Stanford University, USA, 2010.

[5] T. S. Saponas, J. Lester, C. Hartung, S. Agarwal, and T. Kohno, "Devices that tell on you: privacy trends in consumer ubiquitous computing," in *Proc. of 16th USENIX Security Symposium*. Boston, MA: USENIX Association, 2007, pp. 1–16.

[6] B. Daskala and L. Maghiros, "Digital territories," in *Proc. of the 2nd IET Intl. Conference on Intelligent Environments*, 2006, pp. 221–226.

[7] B. Könings, B. Wiedersheim, and M. Weber, "Privacy management and control in ATRACO," in *Ambient Intelligence*, ser. LNCS. Springer Berlin / Heidelberg, 2010, vol. 6439, pp. 51–60.

[8] N.-W. Gong, M. Laibowitz, and J. Paradiso, "Dynamic privacy management in pervasive sensor networks," in *Ambient Intelligence*, ser. LNCS. Springer Berlin / Heidelberg, 2010, vol. 6439, pp. 96–106.

[9] H. Tokuda, M. Beigl, A. Friday, A. Brush, Y. Tobe, A. Sheth, S. Seshan, and D. Wetherall, "Geo-fencing: Confining Wi-Fi coverage to physical boundaries," in *Pervasive Computing*, ser. LNCS. Springer Berlin / Heidelberg, 2009, vol. 5538, pp. 274–290.

[10] M. Langheinrich, "A privacy awareness system for ubiquitous computing environments," in *Proc. of the 4th Intl. Conference on Ubiquitous Computing*. London, UK: Springer, 2002, pp. 237–245.