Challenges in Retaining Privacy in Smart Spaces*

Jimmy C. Chau and T.D.C. Little
Department of Electrical and Computer Engineering
Boston University, Boston, Massachusetts
{*jchau,tdcl*}*@bu.edu*

June 1, 2013

MCL Technical Report No. 06-01-2013

Abstract–Advances in mobile computing, sensors, controls, ubiquitous networking, and other indoor automation infrastructure enable buildings to operate more intelligently, providing improved energy efficiency, safety, convenience, and quality of life. However, many features of these "smart spaces" require sensing, aggregation, analysis, and storage of potentially sensitive information about room occupants. The privacy of the information manipulated by smart spaces quickly becomes a key barrier in realizing the full value of ambient systems and is the focus of this paper. We approach this challenge by first surveying current privacy definitions and mechanisms (access control, k-anonymity, and differential privacy) under the assumption of ambient sensors and networking found in smart spaces. We then identify how existing approaches are not suitable for smart spaces under major smart space privacy scenarios and propose adaptations with strong potential for addressing these scenarios.

*In *Proc. 4th Intl. Conf. on Ambient Systems, Networks and Technologies*, June 25–28, 2013, Halifax, Nova Scotia, Canada, and *Procedia Computer Science*, Volume 19, 2013. This work was supported primarily by the Engineering Research Centers Program of the National Science Foundation under NSF Cooperative Agreement No. EEC-0812056.

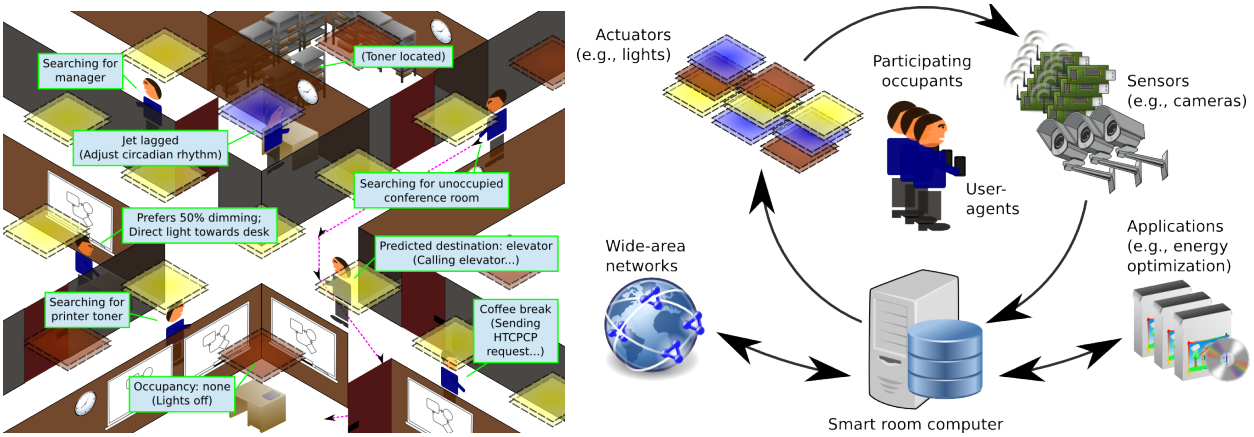

(a) (b)

Figure 1: (a) Highly automated smart spaces provide services to improve productivity, energy-efficiency, and health in a variety of situations. (b) Illustrates the interactions between different parts of a smart space. By gathering information about the smart space from sensors and by providing access to this information to software applications, the smart space computer controls the actuators in the smart space system. Occupants can interact with the smart space infrastructure through the sensors and actuators or through mobile devices that act as user agents. Through the smart space computer, the smart space can also access resources on larger networks, such as the cloud.

1 Introduction

Ambient systems offer many potential benefits to humans in terms of efficiency, safety, health, convenience, and productivity. Our recent work in this area focuses on benefits achieved by the adoption of intelligent and interactive ("smart") lighting systems developed by the NSF Smart Lighting Engineering Research Center [1]. A typical use case here is illustrated in figure 1a. Here we envision individuals in an office building in which the infrastructure ("smart spaces") provides services to identify the location of assets (people, objects, and the like), resources (conference rooms, elevators, & lighting), and rules for their use and consumption.

In the most basic case, we seek to turn lights off when no user is present. In more exotic cases, we monitor user activities and predict behaviors to anticipate where services are required (e.g., elevator arrival or light consumption to support tasks), or we quantify the occupancy and location of certain individuals to maximize the use of conference rooms. In these and many other related use cases, the smart space—sensing, computing, and networking infrastructure—needs to detect, process, and disseminate a great deal of information about individual users. This information can be highly personal and dangerous if exposed inappropriately or maliciously. We focus on the challenge of mitigating the effects of the sensing and aggregation of this information in smart spaces without jeopardizing the benefits that they promise.

We investigate three popular methods of protecting privacy and illustrate their limitations. In the process, major barriers to protecting privacy in smart spaces are highlighted. We also introduce an extension to differential privacy to make it more applicable to emerging smart spaces.

1.1 Capabilities of smart spaces

The scope of smart spaces is vast, growing daily as new applications are developed for mobile computing platforms, such smart phones, in the context of home automation, or for vehicular networks. While these applications are sometimes exotic, they are also practical as concepts seeking to improve human life. Examples include indoor positioning, thermostat control, and our main focus: lighting technologies. Specific innovations include color-controllable lighting units; optical sensors, such as cameras; and visible light communications, which can reuse lighting infrastructure to provide wireless networking capabilities [2, 3, 4].

These technologies can be combined to create highly-automated rooms that anticipate and respond to their occupants' needs. For example, a smart space can find and track occupants using sensors as they move throughout a building to automatically turn lights on near occupants to light their way and off as the occupant passes to conserve energy. Temperature, air flow, sound volume, and other room settings can also be automatically tuned to suit the occupants' personal preference or to match their current activities. Additional examples are illustrated in figure 1a. Through such automation, smart spaces have not only the potential to provide energy savings, but also to improve the productivity of their occupants.

In addition to automation, the sensing, networking, and computational capabilities of smart spaces can provide information services, such as providing directions for navigation, managing resources to facilitate sharing, tracking shoppers to increase sales, interacting with utility companies through smart meters, or monitoring to allow nurses to remotely serve patients (telemedicine). Research into circadian rhythm control by regulating light levels [5] also demonstrates the potential to improve health when a smart space can identify and track individuals. This latter point is critical: realizing many of the benefits of ambient systems requires empowering the computing infrastructure with sensitive information about individuals.

1.2 Organization of the paper

The remainder of the paper is organized as follows: section 2 explains the problems and the desired outcomes for smart space privacy; section 3.1 explores using access control methods to protect privacy in smart spaces; section 3.2 explores applying k-anonymity to smart spaces; and section 3.3 explores applying differential privacy to smart spaces. Section 4 concludes the paper.

2 Overview of privacy problems

While smart spaces have the potential to improve quality of life for their occupants, unless everyone who observes the smart space's data is fully trusted to handle all of the data, smart spaces can also cause their occupants to suffer privacy breaches. Although this trust is reasonable in special cases, such as remote in-home medical care, in which only the patient and medical professionals, who are customarily trusted, interact with the smart space, this assumption of trust is not appropriate in most other scenarios.

For example, in shared buildings, such as office buildings or shopping centers, occupants may not be comfortable sharing their smart space data with each other. In these scenarios, malicious occupants can use the smart space to gather data about other occupants that would otherwise be

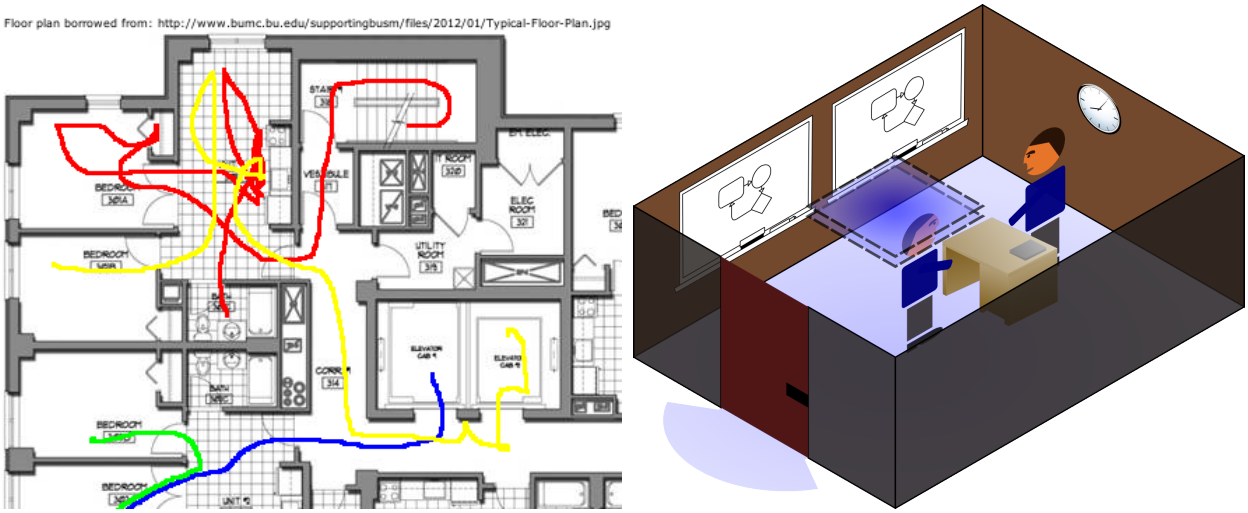

(a) (b)

Figure 2: A smart space may inadvertently leak information in many ways. (a) Indoor positioning or occupancy data (needed for many smart space applications) can disclose the location, activities, and relationships of occupants. (b) Shows how personalized lighting may disclose the presence and identity of occupants within a room: if light matching Alice's personal preference leaks out of a room through a slit under the door, people outside of the room may correctly infer that Alice is inside the room.

inaccessible. One example of this is illustrated in figure 2b, where the capability to adjust a room's settings to meet personal preferences can be used to discover the presence or location of occupants.

In other cases, occupants may not entirely trust the smart space's system administrator. For example, if a shopkeeper analyzes her store's smart space data to maximize profit, a conflict between the shopkeeper's desire for more revenue and the shopper's desire to keep purchasing decisions from advertisers may prevent shoppers from fully trusting the shopkeeper.

Additional privacy problems can arise from the software or devices that comprise the smart space. Like smart phone applications, while smart space software applications from third parties can provide desirable features, they can also contain trojan code: components that access and covertly disclose private information. Similarly, hardware obtained from untrusted parties can leak information. Unfortunately, it is very difficult to screen for such hidden threats [6]. These scenarios, already common in smart phone applications, will be problematic in ambient systems as well.

Prior works have investigated privacy problems in smart spaces or related scenarios. While a few of these papers present a general solution [7], many focus on only one feature of smart spaces, such as protecting location privacy for a particular application [8], and do not address privacy protection broadly for multiple services. We seek a general framework on which to base the development of a privacy protection paradigm that can be applied in a large set of use cases in smart spaces.

Among the works that do aim to provide a general solution for smart spaces, most rely primarily on access control mechanisms to prevent unauthorized access to data [7, 9]. Unfortunately, as we explain in section 3.1, despite access control mechanisms, information can still be leaked to untrusted or partially trusted entities; additional privacy protection is necessary to mitigate these

disclosures.

In contrast to previous works, we seek general solutions to protect privacy even when the smart space must accommodate untrusted entities. This is achieved by exploring privacy mechanisms that are typically used to regulate third-party database access and evaluating them in the context of smart spaces.

2.1 Smart space privacy goals

In order to discuss privacy for smart spaces, we need a working definition of privacy. Unfortunately, no universal consensus exists on the definition of or requirements for privacy; instead, privacy is a nebulous concept that is time-dependent, context-dependent, and subject to different interpretations [10]. These characteristics make privacy difficult to achieve or to measure. Instead of evaluating privacy mechanisms against a universal definition, we compare their utility (whether desired applications can work under the restrictions of the privacy mechanisms) and the integrity of their privacy guarantees (the extent to which their privacy protections hold) against the other mechanisms in various smart space use cases.

Each use case is a desirable smart space service implemented on the architecture illustrated in figure 1b. In this architecture, we assume that only a central computer is fully trusted and that communications to and from this central computer are secure against eavesdropping and tampering. Other components or participants (such as third-party software applications or cloud-based services) can be untrusted or partially trusted. These assumptions mitigate the need to fully vet each part of the smart space while avoiding the need to implement decentralized privacy mechanisms.

3 Approaches to privacy

3.1 Access control

Privacy is often defined by an access control policy that specifies the conditions under which information may or may not be released. These policies consider what information is requested, who is requesting the information, and other contextual information. For example, the access control policy for a telemedicine smart space may specify that access to the patient's medical information is only granted during business hours to the patient's physician. Another, more permissive access control policy may allow all information to be shared or sold to advertising partners.

According to the access control definition of privacy, a system provides or preserves privacy if and only if it adheres to the specified access control policy. However, the policy may not adequately reflect the expected or desired level of privacy and may still allow information to be leaked in undesirable ways. For example, as illustrated in figure 2b, an access control policy that allows a personalized lighting system to access occupancy information may unintentionally disclose this information to untrusted coworkers sharing the smart space; outsiders without direct access to protected occupancy information can still infer the identity of occupants behind closed doors by observing the smart space's response to the protected information (in this case, personalized lighting that escapes under the door).

Table 1: k-anonymization is illustrated: (a) a sample raw dataset and (b) the k-anonymized dataset, where $k = 2$, are shown.

Name	Age	Gender	Score
Alice	17	F	85
Bob	23	M	70
Charlie	16	M	72
Dave	26	M	74
Eve	15	F	61
Frank	22	M	90

(a)

Name	Age	Gender	Score
*	< 23	F	85
*	≥ 23	M	70
*	< 23	M	72
*	≥ 23	M	74
*	< 23	F	61
*	< 23	M	90

(b)

This example highlights two unfortunate limitations of access control. First, the onus to anticipate unintended disclosures lies with the access control policy designer. In this case, since the policy's designer did not anticipate that escaping light can carry occupancy information, the policy was not designed to prevent this disclosure. Anticipating unintended disclosures is especially difficult in smart spaces due to the vast variety of data used, the large number of applications and other entities using this data, and the complex (and sometimes unknown) interactions that occur within smart spaces.

Second, using strict policies to prevent undesirable disclosures can be impractical since such policies may disable otherwise desirable components of the smart space that inevitably share or disclose information. For example, a strict policy would prevent personalized lighting within an office building because personalized lighting would reveal occupancy information to untrusted or adversarial coworkers. Similarly, a strict policy would prevent smart meters from communicating to utility companies since potentially sensitive usage information may be inferred from these communications.

Finally, a strict access control policy would prevent third-party applications, devices, and services from using smart space data, crippling third-party smart space products, since they may contain trojans or deviate from the specified policy in an unverifiable manner. As a consequence, using access control to achieve privacy in smart spaces may hamper the development of smart space products by new, not-yet-trusted developers.

3.2 k-anonymity

An alternative is k-anonymity. Unlike access control, which can only protect privacy by preventing data releases, k-anonymity aims to protect the identity of the person whom is the subject of the released information [11]. In this way, even if desirable applications disclose the k-anonymized sensitive information, smart space occupants remain protected by anonymity. This also allows untrusted or partially trusted third-party applications, devices, and services to be integrated into the smart space without necessarily sacrificing privacy. Similarly, the shopkeeper described in section 2 can analyze her shop's smart space data while allowing her customers to retain their privacy through anonymity.

k-anonymity is achieved by determining which attributes are quasi-identifiers (QIs): attributes (such as name, age, and location) that can be used with other sources of information to identify

people. The values of these attributes are generalized so that they cannot specifically identify any individual; instead, any anonymized record can belong to any of k people.

For example, from the raw dataset shown in table 1a, k-anonymization can be performed by generalizing each QI (name, age, and gender), to form table 1b, so that each tuple of QIs can match at least $k = 2$ people. As a result, an adversary who can match name, age, and gender to identities would still be unable to determine to whom each anonymized entry belongs. For example, the 5th entry in table 1b could either be Alice's or Eve's entry to the adversary.

However, this generalization makes k-anonymized data less precise than data protected with just access control mechanisms. This reduction in precision can degrade the performance of smart space applications. For example, one approach to achieve (probabilistic) k-anonymity is for everyone in a smart space to falsely report being in being in $N - 1$ different regions, where N is the number of regions expected to cover k people [12]. With this anonymized data, an energy-efficiency application that automatically switches lights on and off based on occupancy would not work well because lights would remain on in $N - 1$ out of N regions, which are falsely reported as being occupied.

Still, despite this compromise in utility, k-anonymity is susceptible to failure in several ways [11, 13]. One assumption of k-anonymity is that the anonymizer can determine which attributes in the private dataset also appear in other datasets to determine which attributes should be treated as QIs [11]. Unfortunately, as admitted in reference [11], identifying all QIs is a challenging task that requires knowing about all external sources of information. If a QI is missed when anonymizing data, adversaries can use the attribute to narrow down the anonymizing set. k-anonymity is especially susceptible to this risk in smart spaces; due to the rich abundance of information, many attributes, and hence, potentially unanticipated QIs, exist.

k-anonymity in smart spaces is further complicated because smart spaces generate multiple data releases over time. Unlike privacy mechanisms for traditional databases that only need to protect a single snapshot of the database, smart space privacy mechanisms need to handle perpetually updating information (or, effectively, an ever increasing number of snapshots). Unfortunately, if multiple data releases describe the same or overlapping groups of people, composition or intersection attacks may drastically reduce the size of the anonymizing set [13]. For example, Eve in table 1b completely loses her anonymity on her 23rd birthday because her k-anonymized age transitions from < 23 to ≥ 23; since Eve's entry will be the only one that transitions on Eve's 23rd birthday, an adversary that knows Eve's birthday can identify her entry with certainty.

Intersection attacks may also be possible if separate smart space applications submit separate queries on the same snapshot, yielding multiple data releases. Separate queries and multiple data releases may be generated if different applications and services require data to be k-anonymized in different ways. For example, an indoor navigation service may prefer the anonymization method described in [12] while an energy-efficiency application may prefer to use larger, continuous regions instead of fake locations for anonymization. By overlaying both sets of anonymized location information to find their intersection, the precise location of participants can be determined.

3.3 Differential privacy

In contrast to k-anonymity, differential privacy (DP) composes well: combining multiple differential privacy data releases or combining them with external sources of information does not catastrophically obliterate the privacy guarantees [13]. In addition, like k-anonymity, DP can work in

smart spaces with untrusted or partially trusted components, occupants, and external services since it does not rely on the entities that access information to enforce its policy.

DP aims to ensure that the probability of getting any outcome remains the same whether or not any individual participates; in this sense, participation does not diminish privacy since it does not significantly affect the result. This is achieved by randomizing and adding noise to the functions that use the dataset to mask the effects of individual participants. Formally, as stated in Dwork [14], a randomized function f gives ϵ-differential privacy if for any input datasets D_1 and D_2 differing by no more than one entry and for any set of outputs $S \in \mathrm{Range}(f)$,

$$\Pr[f(D_1) \in S] \leq \exp(\epsilon) \times \Pr[f(D_2) \in S]$$

A smaller ϵ (closer to 0) means that each participant's participation has a less significant effect on the outcome, thus providing greater privacy to each individual.

However, the guarantees provided by DP inherently limit the utility of the resulting data in a smart space. Since the individual should not be able to cause the probability of getting any outcome to change by more than a factor of $\exp(\epsilon)$, which is deliberately kept near unity to protect privacy, personalization, such as adjusting the lighting to match the occupant's personal preferences, is not possible.

For example, imagine that we have a dataset that tracks the presence of people in a room and their lighting preference. We wish to design an ϵ-DP function f, with $\epsilon = 0.1$, that uses this dataset to set the room's lighting levels to match the occupants' preference, or, if no occupants are in the room, to turn off the room's lights. Let's say that Bob, an occupant in the room, represented by an entry in the dataset, prefers 82 lux of neutral white lighting. Ideally, if we ignore privacy, f always sets the lights to 82 lux when Bob is present and to 0 lux when Bob is not present. However, since 0.1-DP requires

$$\Pr\left[f(\text{Bob here}) = 82\right] \leq e^{0.1} * \Pr\left[f(\text{nobody here}) = 82\right]$$

if the lights are off with 90% probability when nobody is in the room, then the lights cannot be 82 lux (or even on) with greater than $1.1 * (100\% - 90\%) = 11\%$ probability when Bob is in the room.

Fortunately, smart space applications that work with large groups of people that don't require personalization and can work with noisy aggregated data remain possible. For example, one potential application is to adjust settings, such as the temperature and lighting levels, in a large auditorium to match the audience's overall preference. DP can also be applied to analyze data to learn general trends that can provide useful insights to help optimize the smart space.

3.3.1 Privacy budgets

Implementations of DP, unfortunately, only allow for a limited number of queries, known as the *privacy budget*. While this limitation is acceptable for database applications (for which DP was designed), in which a snapshot of the database at one point in time can yield useful insights, most smart space applications require frequent and perpetual information updates and will not work properly with stale information. For example, knowing the average preference for lighting level in an auditorium five weeks ago is insufficient for a smart space application to determine the appropriate lighting settings now.

In database applications, the privacy budget is intended to prevent adversaries from repeatedly querying a dataset to gain more certainty about the result. This attack would otherwise be possible since averaging several noisy estimates of one value yields a more accurate estimate of that value. Unless the number of queries is bounded (by the privacy budget), adversaries would be able to almost completely eliminate the noise by averaging, defeating the guarantees of differential privacy.

Dwork et al. [15] present ways to implement differential privacy without limiting the number of updates allowed. Unfortunately, these methods require the queries' answers to monotonically approach a fixed bound with each changing update; this requirement ensures that the query results will eventually stop changing significantly and thus, stop providing any significant update. While the approaches in [15] are suitable for monitoring one-shot events as they happen, the approaches are not as suitable for smart space applications because the smart space applications will eventually stop getting useful updates.

We propose partitioning time-varying smart space data over time, with a separate privacy budget for each time segment. Smart space applications require frequent updates because the needed information changes unpredictably over time. For example, imagine a smart space application that sets the ventilation airflow in an elevator to be proportional to the number of occupants. At time t_0, there are a_0 occupants, so the application learns that there are approximately $a_0 + N_0$ occupants, where N_i is the noise added to the query at time t_i. At a later time t_1, perhaps ten minutes later, the previous occupants have likely left and new occupants have arrived. Let the number of people in the elevator at t_1 be a_1. Assuming that a_0 and a_1 are independent, the estimate of a_1 does not reveal any additional information about a_0 and vice-versa. Thus, multiple differential privacy queries about unpredictably changing information, made in sufficiently different times cannot be used to increase certainty about any particular answer.

This ever-growing privacy budget may be implemented by providing an hourly, daily, or other periodic privacy budget that may only be used to query data from the associated time period. However, more analysis needs to be done to determine how to implement the periodic privacy budget to achieve the desired ϵ for ϵ-differential privacy, especially in cases where the answers to queries at different times cannot be assumed to be completely independent.

3.3.2 Interactions in the smart space

Unfortunately, despite this adaptation and despite the compromises in utility accepted to provide more robust privacy guarantees than just access control or k-anonymity alone, DP may not be sufficient to provide an acceptable level of privacy for smart space occupants. Although DP can add enough noise to mask the maximum effect that any individual's dataset entry can have on query results, occupants in a smart space also have the potential to affect the entries of other people through interactions with those people. In the worst case, one influential individual may be able to affect the attributes of all other individuals. In this case, either that individual will be individually observable (through other dataset entries) despite DP, which results in no privacy for that individual, or enough noise to mask the effect of all entries in the dataset will need to be added, which renders the resulting data entirely useless. Neither option is desirable.

4 Summary, discussion, and conclusion

Future smart spaces, equipped with the means to gather, share, and use information about themselves and their occupants, have the potential to greatly improve quality-of-life by providing convenience, safety, and efficiency. However, these same capabilities also create opportunities for vast and intrusive privacy invasions. In an effort to mitigate these privacy problems while retaining the benefits, we investigate privacy-protection mechanisms in the context of these envisioned smart spaces to build a general smart space privacy framework.

Unlike prior works, we assume a generalized threat model in which devices or applications that consume information are not fully trusted. This broader threat model is more realistic since it accounts for potentially untrustworthy visitors, coworkers, shop keepers, service providers, and other data-consumers that can interact with a smart space. Furthermore, this threat model allows for a more flexible smart space product ecosystem, in which untrusted or partially trusted third parties can contribute to the development of smart spaces without sacrificing privacy.

Unfortunately, within this threat model, existing access-control mechanisms are insufficient to prevent data disclosure or to protect privacy. Other privacy definitions and mechanisms, such as k-anonymity and differential privacy are able to provide more robust privacy guarantees by sacrificing precision, and hence utility, but they each have their own weaknesses when applied to smart spaces. k-anonymity can fail if previously unanticipated data is made available, as time passes, or if smart space components interact badly. Differential privacy provides more robust privacy guarantees and it can be adapted for certain smart space applications by partitioning the data to deal with privacy budget limitations. However, it can still fail if occupants interact with each other (as they are likely to do since they share the smart space).

Although none of the explored approaches preserve privacy in general-purpose smart spaces, the possibility of finding a suitable privacy paradigm has not been ruled out; it may still be possible, with a sufficiently inspired definition of privacy, to provide a satisfactory degree of privacy to the smart space participants without rendering the smart space useless.

In the meantime, more specific use cases, such as remote home care, or single-application smart spaces can still be secured with simple privacy mechanisms such as access control and k-anonymity. However, future smart space privacy research must both account for the adversary's ability to infer information from a smart space's response to private information and consider the important role that untrusted or partially trusted parties play in the smart-space ecosystem.

References

[1] Smart Lighting Engineering Research Center, [Online]. Available: `http://smartlighting.rpi.edu`.

[2] J. Chau and T. Little, *Improved Design for an Optical Communication System*, New York: Smart Lighting Engineering Research Center, 2011.

[3] J. Chau, K. Matarese and T. Little, *IP-Enabled LED Lighting Supporting Indoor Mobile and Wireless Communications*, MobiSys, 2010.

[4] S. Afshari, S. Mishra, J. Wen and R. Karlicek, "An adaptive smart lighting system", in *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, New York, NY, USA, 2012.

[5] M. Figueiro, N. Lesniak and M. Rea, "Implications of controlled short-wavelength light exposure for sleep in older adults", *BMC Research Notes*, vol. 4, no. 1, p. 334, 2011.

[6] X. Wang, M. Tehranipoor and J. Plusquellic, "Detecting malicious inclusions in secure hardware: Challenges and solutions", in *Hardware-Oriented Security and Trust, 2008. HOST 2008. IEEE International Workshop on*, 2008.

[7] J. Suomalainen, P. Hyttinen and P. Tarvainen, "Secure information sharing between heterogeneous embedded devices", in *Proceedings of the Fourth European Conference on Software Architecture: Companion Volume*, New York, NY, USA, 2010.

[8] A. Beresford and F. Stajano, "Location privacy in pervasive computing", *Pervasive Computing, IEEE*, vol. 2, no. 1, pp. 46-55, Jan–Mar 2003.

[9] R. Campbell, J. Al-Muhtadi, P. Naldurg, G. Sampemane and M. D. Mickunas, "Towards security and privacy for pervasive computing", in *Proceedings of the 2002 Mext-NSF-JSPS international conference on Software security: theories and systems*, Berlin, Heidelberg, 2003.

[10] M. Friedewald, D. Wright, S. Gutwirth and E. Mordini, "Privacy, data protection and emerging sciences and technologies: towards a common framework", *Innovation: The European Journal of Social Science Research*, vol. 23, no. 1, pp. 61–67, 2010.

[11] L. Sweeney, "k-anonymity: A Model for Protecting Privacy", *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, vol. 10 (5), pp. 557–570, 2002.

[12] S. I. Ahamed, M. M. Haque and C. S. Hasan, "A novel location privacy framework without trusted third party based on location anonymity prediction", *SIGAPP Appl. Comput. Rev.*, vol. 12, no. 1, pp. 24–34, Apr 2012.

[13] S. R. Ganta, S. P. Kasiviswanathan and A. Smith, "Composition Attacks and Auxiliary Information in Data Privacy", *CoRR*, vol. abs/0803.0032, no. 0803.0032v2, Mar 2008.

[14] C. Dwork, "Differential Privacy", in *Automata, Languages and Programming*, vol. 4052, M. Bugliesi, B. Preneel, V. Sassone and I. Wegener, Eds., Springer Berlin / Heidelberg, 2006, pp. 1–12.

[15] C. Dwork, M. Naor, T. Pitassi and G. N. Rothblum, "Differential privacy under continual observation", in *Proceedings of the 42nd ACM Symposium on Theory of Computing*, New York, NY, USA, 2010.