



HAL
open science

Public Privacy and Brick Houses Made of Glass

Stephen Marsh, Ada Diaconescu, David Evans, Tracy Ann Kosa, Peter R. Lewis, Sheikh Mahbub Habib

► **To cite this version:**

Stephen Marsh, Ada Diaconescu, David Evans, Tracy Ann Kosa, Peter R. Lewis, et al.. Public Privacy and Brick Houses Made of Glass. 12th IFIP International Conference on Trust Management (TM), Jul 2018, Toronto, ON, Canada. pp.137-148, 10.1007/978-3-319-95276-5_10 . hal-01855989

HAL Id: hal-01855989

<https://inria.hal.science/hal-01855989v1>

Submitted on 9 Aug 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Public Privacy and Brick Houses made of Glass

Stephen Marsh¹, Ada Diaconescu², David Evans³, Tracy Ann Kosa⁴,
Peter R. Lewis³, and Sheikh Mahbub Habib⁵

¹ University of Ontario Institute of Technology, Oshawa, Canada
`stephen.marsh@uoit.ca`

² Telecom ParisTech, Paris, France `ada.diaconescu@telecom-paristech.fr`

³ Aston University, Birmingham, UK `d.j.evans@aston.ac.uk`,
`p.lewis@aston.ac.uk`

⁴ Stanford University, USA

⁵ Technische Universitaet Darmstadt, Germany `sheikh@tk.tu-darmstadt.de`

Abstract. In this work in progress paper, we present a description of a new view of privacy in public, examining how it is possible to ascertain the privacy levels of individuals in context and in groups, and different ways of visualising these Public Privacy levels. We examine how awareness of one’s Public Privacy may have an impact on behaviour and privacy protection in general, and propose future work to examine the concept in more detail.

1 Introduction: Privacy and the Privacy Paradox

Most people shut their curtains, and don’t live in glass houses – at least, not without the aforementioned curtains. We might like to watch Big Brother, but that doesn’t mean we all want to participate. And yet, for various reasons – time, efficiency, lack of choice, lack of awareness, and so on – we intentionally sometimes and unintentionally oftentimes leak or broadcast information to whomever might be listening. Certainly, it’s possible to adjust our privacy settings on social networks and so forth, but in reality, the tools we use and the environments, such as smart cities, we exist in, thrive on sharing data around and, crucially, *about* us.

A perhaps more pernicious problem is that the people around us may well be either unaware or uncaring about the effects that their own behaviour with respect to privacy and sharing can have on the world around them. These effects can range from inadvertently exposing private information that can be used against others to deliberately sharing pictures or movies that include others in them, with similar potentially bad results.

Privacy, then, is a social problem, with social ramifications, even when there are people who care about it.

The effect is that whilst we might believe that we live in brick houses with the curtains closed, those around us may well just be turning those houses into glass, and no amount of effort on our part can put curtains around the whole building.

Recent work on Privacy Awareness [5] acknowledges that, whilst information is ‘out there’, it is no bad thing to help people become more aware of both where it is and how it might be linked to other extant information to provide fuller pictures of people (profiles).

Other work in the realm of Device Comfort [11, 3, 2] focuses on the ability of the very technology that causes the problems to ameliorate these problems when the technology’s sensing (and reasoning) capabilities are combined with human social norms as expressive tools. That is, trust, comfort, regret and so forth are useful tools to assist people to better *understand* the way in which their actions, in context, can effect, reveal, or damage their security and privacy postures.

In this paper we propose *Public Privacy* as a toolset to address some of the problems associated with the social aspect of privacy, within the context of Privacy Awareness. In the Public Privacy paradigm, we assign privacy states [6] to group spaces, which we intend will have the result of making more public the privacy sensitivity and actions of the members of that space. Additionally, we conjecture that it then can provide suggestions, tools, and societal pressure in different contexts to raise more awareness and potentially initiate or support behaviour changes.

2 On Privacy, Problems, What has come before, and What it means

This work falls generally in the arena of Privacy Awareness, although with a significant twist. It is also firmly situated in boundary regulation (see for example [7]), but again from the point of view of physical space. In this, the work is directly related to Altman’s [1] dialectical conception of privacy in which boundaries are used and negotiated. In our work, the boundaries may already be established but the individual in the shared space can choose to participate or indeed exercise forms of control, thus, the idea of Public Privacy support in technical tools allows for boundary and personal privacy negotiations in a sociotechnical setting. It is similar to Petronio’s communication privacy management [14] theory in that it allows this negotiation but the toolset is paramount here. Boundary regulation is also a feature in [7] but once again, the emphasis is on online networks, where physical social boundaries may not be present.

As well, [13] discusses privacy in a sociotechnical situation, in particular unpacking how privacy is determined and considered in sociotechnical systems, but uses Altman’s theory to examine how privacy could be understood in digital systems.

Here, we are bringing the model back to shared social spaces in which technology is an adjunct: people do stuff with it, but it’s not the technology that enables people to work together, in the sense, for example, of [15] and other similar CSCW-related work. That is, we don’t care so much about the stuff, we care about providing awareness so that people can negotiate boundaries together so that they can work and play together in comfort and safety. To achieve this effectively, it might also be necessary or beneficial to design the awareness-raising

and mediation system such that it explicitly considers the awareness of users, and its role in promoting and being sensitive to it, in decision-making. Thus, we draw on and further develop work on computational self-awareness, particularly, public self-awareness [8] and social self-awareness [4].

Further, the work in this paper differs from related work in that, whilst it certainly is about ‘digital’ privacy, its concern in things like boundary regulation is in the physical world. As well, whilst in some cases examining the content of messages is used to determine if privacy is about to be violated or harmed, such invasive methods are not necessary here: we are focused on behaviour in society where privacy is respected, not the online aspect of how much information is ‘out there’ in social networking sites, for instance. The end results may be similar: increased awareness, better management of information and boundaries, and so forth, but the methods differ.

3 Public Privacy

In the Public Privacy paradigm, we focus on privacy as two⁶ things in a group space (a meeting room or a coffee shop, for example):

- A goal to be worked towards (at some level⁷ or state)
- A state, in a space, in context, that can be expressed to everyone in that space.

In the former, everyone in the space has their own ‘desired’ level of privacy and sets up their own devices or behaviours toward that end, whilst when in a space that desire may conflict with that of the space itself, and so the latter expresses how the two (individual and public) may differ, with attendant possibilities we will discuss further below.

Following Kosa [6] we place privacy as an abstract notion within a set of states, or levels, which are shown in table 3.

One option is to represent these states as numbers, since we need to be able to compute the state of a space from that of the people in it (amongst other things). The most basic aspect of public privacy is that any space within which there are people (and devices) has an inherent privacy state, which we can compute and display in a public form (and forum). For simplicity, we have added numbers to the states here. The higher the number, the less ‘private’ the state.

3.1 Formalising the Public Privacy State of a Space

The Public Privacy State of a shared space S is a function of:

⁶ It can definitely be both at the same time.

⁷ Kosa [6] talks about privacy states, which we think is a reasonable way to express the level of privacy at a suitable concrete level, and has the associated benefits of understandability and computational tractability of sorts. In this work we interchange level and state because different states obviously have different achievable privacy, or levels of privacy.

Sn	State	Physical Self	Digital Self	Example
1	Private	Existence is unknown	Existence is unknown	A child hiding
2	Unidentified	Existence is known	No identity data	A shadow
3	Masked	Existence is visible	Limited identity information	An organ donor
4	De-identified	Existence is unconnected	Non-specific identity information is known	Unpublished identity information is available about a patient in a study
5	Pseudonymous	Existence is connected but accuracy is unreliable	Identity data could apply to multiple persons	Reference to common characteristic, e.g., female person such as Jane Doe
6	Confidential	Existence is connected but limited distribution	Limited identity data available to defined person in a certain role	A doctor with access to her patient's records
7	Identified	Existence is connected with unlimited distribution	Data is available with few or no controls	Social networking sites
8	Public	Existence is completely transparent	Digital self is live-cast, online and cross-referenced	Babies or small children (limited control)

Table 1. Privacy States (from [6]) with Sn (State Number) added

- the Context, Con of the space;
- The Privacy States of the individuals in that space (including that of their devices present), which we represent as $P_S = f(\rho_{S,1} \dots \rho_{S,i})$, where i is the number of individuals in the space, and $\rho_{S,n}$ is the Privacy State of individual n);
- Privacy ‘requirements’ of the individuals in the space, which we represent as $A(S, \text{Con}) = f(\lambda_1(S, \text{Con}_1) \dots \lambda_i(S, \text{Con}_i))$ with again $\lambda_n(S, \text{Con}_n)$ that of individual n .

Each of these can be calculated, as we discuss in this section. Note also that the Public Privacy state of a space is highly dynamic based on the individuals joining and leaving the space, as well as its context (which may change from time to time). For example, a meeting room may be in a context of *Business Meeting* or *Legal Meeting* or *Interview, Presentation*, and so forth, whilst a coffee shop would usually be in the context of *Public Space*, but this also might change in different circumstances.

In this instance, we consider that the Context of a space is obtainable in some fashion, and that the Context is the same for every individual in the space (see section 4 for a further discussion on this).

The Context of a Space. Spaces are used to *do* things. What this means is that it should be possible to determine from moment to moment what the space is purposed for. This context may change over time for a given space, or the space may indeed be designed in a way that ensures its context (but even that may be subverted: the best parties usually end up in the kitchen).

For our purposes, it is sufficient to split the context in a relatively rough manner. The contexts available to a space are:

1. Private-Meeting: A closed session where attendees are potentially invited.
2. Public-Meeting: A session where anyone can potentially attend, but which is dedicated to a specific task. For example, a workshop or conference session.
3. Public-Static: A space, such as coffee shop, restaurant, where people can enter and leave freely. Can also include spaces like supermarkets, malls, etc.
4. Public-Transient: Slightly different from Public-Static, the Public-Transient space is one which is used as a place through which people move. For instance, railway stations, airports, etc.

Without getting into a discussion about the language of architecture, consider that in general some spaces are designed to facilitate certain contexts. Lecture theatres are designed for Private-Meeting or Public-Meeting contexts, for example. However, it's entirely possible to convert such spaces into, say, Public-Static spaces at need. Other spaces, such as railway stations, are much more difficult to turn into certain contexts.

Contexts advise behaviours. Each context has specific expectations associated with it. These are expectations of behaviour (how many people are talking at once, how many listening, who pays attention to whom, and so on) which naturally impact privacy expectations for the space. It's unusual to for instance tweet about what is happening in a private meeting, but more acceptable from a public meeting such as a workshop or a press briefing.

Regardless of space, to an extent, the context is what is important to us here since it sets its own expectations. We would suggest that the privacy expectations of the contexts listed above decrease as we move down the list.

It is a matter for future work to determine different other contexts and sub-contexts for spaces.

Private Privacy: The Privacy State of Individuals. In the functions above, $P_S = f(\rho_{S,1} \dots \rho_{S,i})$, where i is the number of individuals in the space. How do we determine ρ for any individual, since this is perhaps the most important aspect of the problem?

Ironically, more information is of great utility in this computation, but let us gloss over that for a moment. What is needed is something akin to that calculated by the Device Comfort methodology [11]. However, in and of itself, that may prove troublesome, since by its nature the concept can be seen as privacy-invasive on an individual level.

In general, then, there are two different aspects here:

- What is currently being done by the individual;

- The history of the individual.

In the first, we can determine something about what is being done at a specific time by observation, which includes the status of their devices, if any: what apps are being run, what OS is installed, and what patches, for instance, all of which have an impact on the potential for the device to be used in privacy invasive ways either intentionally or not.

For the purpose of the current work, we can place a value on the status and action being taken. More privacy problematic behaviours (such as texting or using social networks, using cameras and so forth) have a higher ‘risk’ score.

For the latter, we use Kosa’s states [6]. To determine the state of the individual based on past history, privacy awareness tools can be used (see e.g. [5]).

At this point, we can combine the two aspects to make sense of the person in the moment. What usually takes precedence in a situation where there is a conflict (wildly differing states for each of the two aspects) is perhaps a matter of preference, although we would argue, much as with trust, that historical behaviour is a relatively good predictor of what immediate future behaviour might be. This being the case, the history aspect is usually the one which is chosen. That said, conflict should be acknowledged and so here we can put certainty values on the result.

How Do You Feel About Privacy? The second aspect of the computation is what we call a privacy ‘requirement’ for an individual. In this, we agree with [12] with regard to context. In order to compute the context and requirements within it, we need but ask. [3, 2], used a simple method to determine both how people felt about sharing information, to whom, and also what for, with the context being indicated by purpose of information sharing. Another method for eliciting what we might call privacy stance is in using labels such as those identified in [16]. Note that there are also certainly tools for expressing privacy preferences, including for instance the Platform for Privacy Preferences, and these can also be used readily in the calculations here, but the approach taken is one that aims explicitly for simplicity (we just don’t need to make it harder than it has to be [10]).

3.2 Expressing the Privacy State of a Space

Ironically, the privacy state of a public space is both a public and a private notion. As a public measure, it is available to everyone in that space to show how everyone’s tools and behaviours come together. As a private measure it is affected by the context of the individual as well as the space. In both it becomes a warning, a guide, a queryable service, a political statement, and a suggestion, as we will see.

Publicly... A public expression of the privacy level of a space is just that: something of which everyone in the space is or can be aware. We can express

it on a shared display in the space (in our development we are using Apple TV apps on shared screens as well as other shared devices - the key here is that the displays are social and shared). The value can be expressed in different ways. Currently we are working with colours, icons and numbers, potentially in combinations. We are planning experiments to determine how these may be received and understood, and as such they are in flux.

Privately... We envisage that Public Privacy in private (sic) can be a tool for the individual in various ways.

To begin with, the expression of the privacy state of the space can be shown on a personal screen (when, for instance a public shared screen is unavailable or unsuitable for a particular use, for reasons of different ability or requirements) and thus available to all.

Moreover, Public Privacy, when discussed in the context of the individual, becomes even more interesting because it can be changed from that individual point of view. To illustrate what is meant here, consider a brief scenario, where there is a group of strangers in a shared space (say an ice cream parlour), and a new individual (our subject) enters. The privacy state of the space, as calculated, is shown to the subject. For the sake of argument, let's say the state is one that reflects a low level of privacy, a low level of privacy awareness, and a social (in this space) acceptance of sharing as 'okay.' Our subject is, for one reason or another, a person who values their privacy in special ways - pictures aren't okay, for example - and the calculated level of privacy is, in ordinary circumstances, outside of their comfort zone (to put it another way, it's way too low to make them happy). However, in context, that makes sense. Plus, no-one in the place is taking pictures (they're all enjoying the ice cream instead) and so the privacy level in that context is reasonable for our subject because their own context is specific to pictures. That given, two things may happen: firstly, the contextual privacy view of public privacy state is that it is reasonable, but also the public view of privacy within other individual devices is updated to reflect that context. In other words, the 'requirement' that pictures of other people in the space not be taken can be added to individual contexts whilst the public view may not change (unless cameras start to be pulled out).

3.3 And, how do we use this knowledge?

Public Privacy is a Privacy Awareness tool, but it is one with something of a twist: we are not concerned here with online privacy *per se*. We are concerned with potential actions in a contextual physical space, populated by others, with their own goals, devices and behaviours, that may have an impact on an individual's privacy. This could be because personal or shared information might be further shared outside the group, inadvertently or purposely, images could be shared, videos might be taken, and so on. Each of these things has an impact on our privacy state as individuals. Public Privacy makes the potential for these things known and to some extent we hope understood.

This has different possibilities in various contexts, some of which we illustrate here. Before entering into these discussions, a few words about trust and associated topics are in order. As we note in e.g., [9], trusting behaviour is the answer to two questions: how much does one have, and how much does one need (and the obvious happy result is that if you have more than you need, you can say you ‘trust’ the other). In many circumstances, we can use artificial tools, often legal or social ones, to increase one or decrease the other, with the effect that some form of relationship might be entered into. We will use this observation in some of our scenarios below./footnoteWe are sure other scenarios exist!

- In a situation where for example business discussions around a private topic might be ongoing, the members of the discussion have different options, including Non-Disclosure Agreements (NDAs). These tools have an effect on trust that may or may not be beneficial - the act of asking people to sign a legal document may decrease the amount of trust needed, but it doesn’t always match the circumstances, and it doesn’t always encourage others to feel comfortable in a given situation. However, if everyone was made aware of the Privacy state of the room in which the meeting was taking place, based in this instance on past history and potentially trust in the members of the discussion, a level of comfort may be achieved which allows discussion to take place without the legal niceties.
- Related to this, there may be circumstances where the members of the discussion simply don’t want to have their names on such documents, and the ability to discuss privacy in such a space, aided by the Public Privacy model at the time, is a powerful social aid.
- Entering into a space with a specific Privacy State allows a participant in that space to better protect themselves with more technical privacy and security tools as needed, since forewarned may well be forearmed in this case. The Privacy State is additionally valuable for tools that use Device Comfort [11] since it provides information that such tools can use to automatically adjust their security posture.
- The members of a space which has a certain Privacy State can observe changes to that state as others enter and behave accordingly (for more discussion of this, see below in section 5).

4 Current Status, Future Work

Public Privacy is a work in progress. We are currently implementing the system using both ‘public’ and ‘personal’ tools, include Apple TV apps to be able to display the Public Privacy state on shared screens, iOS and Android apps for displays on personal screens, and a cloud-based approach which allows a multi-platform approach with web browsers.

More importantly, we are using the scenarios above, amongst others, to design experiments that will illustrate the usefulness of the idea in physical spaces and for different cultures.

There are some other limitations to the model thus far. For instance, the context of an individual in a space might not in fact be that of the space itself. It is entirely possible to be a small group of people within a coffee shop having an at least semi-private meeting. In this case, the contexts of the individuals in the meeting form something of a *sub* space for that group of people. In other words, some spaces can encompass multiple contexts at the same time. We leave the calculation and implementation of this for future work.

A further avenue to explore in operationalising Public Privacy, could be to represent privacy states not as scalar values but as logical propositions. An individual's *logical privacy labels* could represent different expectations and preferences over their behaviour and the behaviour of others. Then, upon entering a social situation, the individuals' propositions could be combined, and 'solved' for contradictions, which are then used to prompt or otherwise raise awareness with the respective users. For example, the person walking into the ice cream parlour can realise that their own propositions (expectations and preferences) are entirely compatible with the prevailing ones (no contradictions), or can be prompted that there are existing contradictions, that the group may wish to address. We therefore arrive at a form of computer-mediated privacy agreement. In general, one advantage of this approach would be to enable us to capture, reason about, and attempt to reconcile a more qualitative, multi-dimensional trust, and hence facilitate decisions based on a richer set of human expectations.

Having operationalized the Privacy State of a space we introduce the idea of Alliance States as a method for aligning and codifying individual and group's intentions in space.

This is extended from the idea of a designed alliance, a social ritual where a group of people designs the emotional atmosphere it wants in a space and context, and, its conflict protocol, how the group wants to be when things get difficult between them. We propose that we codify a 'digital-alliance' between systems.

Often, what comes up in the emotional-atmosphere part of the Designed Alliance are concepts like - trust, safety, and confidentiality. Building on Kosas Privacy States, we create Alliance States, that codify levels of trust, safety and confidentiality. Alliance States could have some quantitative assessment of their quality - i.e. a system could have trust as 2/10, i.e. low, safety 8/10 i.e. high, confidentiality 1/10 i.e. low or unknown.

A 'public' space would quantify its Alliance State, and a 'personal' device entering a public space could read the Alliance State of the space. Systems could negotiate Alliance States, before sharing information. The comparison between a Public and Private alliance states and only interact when certain criteria are met between the Alliance States - and, if they are different, then the systems work to adjust their states so that there is a match before they interact.

5 Discussion: How societies, or at least outlooks, might change

One goal of Public Privacy is to change minds. How might this work? It is our view that much social change happens because society makes it clear that the change is needed. This can happen in various ways - laws, for example, are in many cases expressions of societal expectations, and social pressures are also expressed in interpersonal encounters. Consider, for instance, smoking or driving under the influence. In many cases, since expectations change, so have the ways in which people are influenced in these circumstances (people in many countries don't smoke in different spaces not just because the law says they can't (the law says nothing about smoking in your home), but also because often if they try they are told (asked!) not to). And so, slowly, opinions and potentially behaviours change.

Right now, we have a privacy problem. It is seen, at least in part, as a socially acceptable thing to take pictures, and share them and other information, regardless of other peoples' opinion, wants or needs. The potential effects are not always beneficial, as has been pointed out elsewhere, including in initiatives such as the Privacy Awareness Weeks of several territories.

In part, we see Public Privacy as a tool to raise social awareness within individuals of the consequences of their actions. If one enters a space and sees the privacy state change, and behaviours change with it, this is the first step to this awareness. A second step comes when people take the public awareness tool and use it to *ask* others for change - instead of 'please don't smoke' the refrain becomes 'I'd rather you didn't tweet about this' or 'please don't share that picture', or even 'you really shouldn't be sharing all this stuff.' Once privacy instead of profligate sharing becomes an acceptable stance, we believe⁸ change can happen.

6 Conclusions

Public Privacy is a tool, an awareness raising mechanism, and a method of social change (and as such a political stance). It uses a technical methodology to estimate the Privacy State of shared spaces - that is, spaces where more than one person and their devices are present. This shared Public Privacy State can be used in a variety of different ways:

- To help people in the space better protect their own information (or private selves) either automatically or with human input;
- To raise awareness of issues to help change behaviours or at least educate, and not least
- To help people get things done in contextual shared spaces.

The Public Privacy tool is currently being developed based on the technical discussion in this paper. Experiments for both correct behaviour and acceptance are planned.

⁸ Granted, perhaps naïvely!

References

1. Irwin Altman. *Environment and Social Behaviour*. Brooks/Cole Publishing, 1976.
2. Saghar Behrooz. Trust-based framework for information sharing in health care. Master's thesis, University of Ontario Institute of Technology, 2016.
3. Saghar Behrooz and Stephen Marsh. A trust-based framework for information sharing between mobile health care applications. In Sheikh Mahbub Habib, Julita Vassileva, Sjouke Mauw, and Max Mühlhäuser, editors, *Trust Management X: 10th IFIP WG 11.11 International Conference, IFIPTM 2016*, pages 79–95. Springer AICT, 2016.
4. K. Bellman, J. Botev, H. Hildmann, P. R. Lewis, S. Marsh, J. Pitt, I. Scholtes, and S. Tomforde. Socially-sensitive systems design: Exploring social potential. *IEEE Technology and Society Magazine*, 36(3):72–80, Sept 2017.
5. Simone Fischer-Hübner, Julio Angulo, Farzaneh Karegar, and Tobias Pulls. Transparency, privacy and trust technology for tracking and controlling my data disclosures: Does this work? In Sheikh Mahbub Habib, Julita Vassileva, Sjouke Mauw, and Max Mühlhäuser, editors, *Trust Management X: 10th IFIP WG 11.11 International Conference, IFIPTM 2016*, pages 3–14. Springer AICT, 2016.
6. Tracy Ann Kosa. *Towards Measuring Privacy*. PhD thesis, University of Ontario Institute of Technology, 2015.
7. Airi Lampinen, Vilma Lehtinen, Asko Lehmuskallio, and Sakari Tamminen. We're in it together: Interpersonal management of disclosure in social network services. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '11*, pages 3217–3226, New York, NY, USA, 2011. ACM.
8. Peter R. Lewis, Marco Platzner, Bernhard Rinner, Jim Torresen, and Xin Yao, editors. *Self-Aware Computing Systems: An Engineering Approach*. Springer, 2016.
9. Stephen Marsh. *Formalising Trust as a Computational Concept*. PhD thesis, Department of Computing Science, University of Stirling, 1994. Available online via <http://www.stephenmarsh.ca/Files/pubs/Trust-thesis.pdf>.
10. Stephen Marsh, Anirban Basu, and Natasha Dwyer. Rendering unto Cæsar the things that are Cæsar's: Complex trust and human understanding. In *Proceedings IFIPTM 2012: Trust Management VI, Surat India.*, pages 191–200. Springer AICT, 2012.
11. Stephen Marsh, Pamela Briggs, Khalil El-Khatib, Babak Esfandiari, and John A. Stewart. Defining and investigating device comfort. *Journal of Information Processing*, 19:231–252, 2011.
12. Helen Nissenbam. *Privacy in Context*. Stanford Law Books, 2009.
13. Leysia Palen and Paul Dourish. Unpacking "privacy" for a networked world. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '03*, pages 129–136, New York, NY, USA, 2003. ACM.
14. S. Petronio and W.T. Durham. Communication privacy management theory. In L.A. Baxter and D.O. Braithwaite, editors, *Engaging theories in interpersonal communication: Multiple perspectives*, pages 309–322. Sage, Thousand Oaks, CA, 2008.
15. Mohamad H. Salimian, Derek Reilly, Stephen Brooks, and Bonnie MacKay. Physical-digital privacy interfaces for mixed reality collaboration: An exploratory study. In *Proceedings of the 2016 ACM International Conference on Interactive Surfaces and Spaces, ISS '16*, pages 261–270, New York, NY, USA, 2016. ACM.
16. Pamela J. Wisniewski, Bart P. Knijnenburg, and Heather Richter Lipford. Making privacy personal. *Int. J. Hum.-Comput. Stud.*, 98(C):95–108, February 2017.