Computer Supported Cooperative Work
DOI 10.1007/s10606-012-9177-z

Grounding Privacy in Mediated Communication

Natalia A. Romero[1], Panos Markopoulos[2] & Saul Greenberg[3]
[1]*Industrial Design Engineering Faculty, Delft University of Technology, Landbergstraat 15, 2628 CE Delft, Netherlands (E-mail: n.a.romero@tudelft.nl);* [2]*Industrial Design Department, Eindhoven University of Technology, Den Dolech 2, 5600 MB Eindhoven, Netherlands (E-mail: p.markopoulos@tue.nl);* [3]*Department of Computer Science, University of Calgary, Calgary, AB, Canada, T2N 1 N4 (E-mail: saul.greenberg@ucalgary.ca)*

Abstract. This paper addresses the need of interpersonal privacy coordination mechanisms in the context of mediated communication, emphasizing the dialectic and dynamic nature of privacy. We contribute the Privacy Grounding Model—built upon the Common Ground theory—that describes how connected individuals create and adapt privacy borders dynamically and in a collaborative process. We present the theoretical foundations of the model. We also show the applicability of the model, where we give evidence from a field study that illustrates how it can describe privacy coordination mechanisms amongst users of an instant messaging application and a desktop awareness system. The model describes efficient and effective factors that communicators consider in their decisions to use mechanisms for coordination. The Privacy Grounding Model aims to help designers reflect on how their system supports, or fails to support, people's need for lightweight and distinctive privacy coordination mechanisms, and in particular how communicators within the system create and use privacy border representations for grounding their needs to interact with each other.

Key words: interpersonal privacy, common ground theory, mediated communication, HCI

1. Introduction

The broad adoption of the Internet, mobile phones and social networking applications has made individuals in developed societies almost continuously connected and accessible to others. This 'always available' shift in culture often places individuals under considerable social and organizational pressure. For example, tensions have been observed between expected availability and desire for communication, resulting in unmet expectations and/or undesired communication among the communicating parties involved (e.g. Woodruff and Aoki 2003; Romero et al. 2007a).

By means of sustained and almost continuous connectivity, users of communication media acquire and maintain awareness of others. For the purposes of this paper, *awareness* is defined as an understanding of the whereabouts of others, and/or their activities, feelings, experiences, and/or—more generally—their current status. Conversely, individuals provide to their social network, either through explicit action or through automated technical

Published online: 24 August 2012

means, rich and frequent information about themselves, ranging from microblogging information (status, posts, etc.); real-time audio and video of oneself (e.g., as in media spaces); and contextual information gathered from sensors (such as one's location).

On the one hand, individuals can translate this awareness information into fine-tuned opportunities and expectations of availability, which can enable opportunistic and timely social interaction that many truly enjoy or value for their utility. On the other hand, regularly sharing this information makes individuals accountable for their availability: it can compromise their prerogative to choose whether and when to engage in communication. As a result, plausible deniability is severely diminished (Nardi et al. 2000) as technology hinders an individual's ability to conceal related choices and social behaviour that had been possible over traditional media (Erickson et al. 1999). As observed in the ASTRA study (Markopoulos et al. 2004; Romero and Markopoulos 2005; Romero et al. 2007a), individuals may feel coerced to accept undesired interruptions, resulting in an undesired state of interaction that is associated to affective costs of obligation, or feelings of uncertainty and disappointment towards unresponsive interaction.

The trade-off between privacy and social interaction was described previously by Altman (1975) in the context of unmediated social interactions. This view was later echoed by Palen and Dourish (2003) who argued that individuals are challenged by modern technology to constantly regulate and protect themselves (and others) from undesired interaction. This challenge is slowly gaining broader acceptance by researchers in human computer interaction. Yet interaction design rarely reflects the fact that privacy preferences are always under negotiation and, as argued by Petronio (2002), that privacy borders are set cooperatively by the interacting individuals.

Our research contributes a theory that articulates how connected individuals create and adapt privacy borders both dynamically and in a collaborative process. Our goal is to provide a theoretical foundation for related discussions that will allow the classification and description of interpersonal privacy mechanisms of individuals in a technology mediated communication setting, thus informing the design (or critique) of the relevant systems.

The remainder of this article is structured as follows. First, we provide an overview of related efforts that examine privacy regulation in modern communication media. We then build on the theories of Altman (1975), Petronio (2002) and Boyle and Greenberg (2005) that describe the dynamic and dialectic nature of interpersonal privacy coordination, and how individuals continuously and interactively coordinate their interpersonal privacy borders to communicate their intentions to interact (or not interact) with others. Next, we introduce the Privacy Grounding Model (PGM) as an adaptation of the Common Ground theory (Clark 1996), where it describes the coordination process of interpersonal privacy borders in mediated communication. Subsequently, we report a field study, where we illustrate that PGM is able to describe the dynamic and dialectic process of interpersonal privacy coordination as done 'in the wild'. We conclude by discussing our findings, design implications, and future work.

Our overarching argument is that designing privacy mechanisms in mediated communication is about designing systems that let communicators dynamically negotiate the process of initiating and reacting to communication, conjointly developing a desire state of privacy. We contribute the Privacy Grounding Model (PGM) that describes this coordination process of interpersonal privacy borders in mediated communication.

2. Towards interactive solutions for interpersonal privacy management

Traditionally, privacy in relation to computing and communication systems has been conflated with issues of access to personal information. However, Dourish (1993) argued that users of mediated communication systems in practice are primarily concerned with privacy issues regarding their interaction with others, such as protecting personal space from undesired interaction. Other related research addressed the relevance of supporting information disclosure to manage interpersonal accessibility beyond the control of personal information access. For example, Bellotti and Sellen (1993) defined *principles of feedback and control*. Nardi et al. (2000) introduced the concept of *outeracion*. Palen and Dourish (2003) presented a framework of *genres of disclosure*. All these works inform the design of mechanisms for communicators to feedback their information disclosure to others and to control the disclosure process. Although these authors recognized the importance for communicators to jointly understand their purposes or intentions of disclosing, Belloti and Sellen (1993) and Palen and Dourish (2003) associated this information to existing social knowledge and culture, which are thus handled outside the communication technology design. While Nardi et al. (2000) identified technology support to use parallel tracks for disclosing such intentions, their work is limited to Instant Messaging technologies.

Pursuing the interactive approach, Woodruff and Aoki (2003) evaluated Push-to-Talk technology (a half duplex audio communication appliance that offers separate modes for sending and receiving messages, making speech overlap impossible). They reported that Push-to-Talk is a very low effort and low commitment approach to establishing communication, and that it is effortless and flexible enough to allow users to implement a set of social interaction mechanisms to coordinate their availability. These include the maintenance of plausible deniability, delays, or omissions in responding to communication acts by others. Such mechanisms created a sense of reduced interactional commitment between users, thus supporting a lightweight interactive negotiation process. However, Push-to-Talk was successful only within peer-to-peer communities, as tensions with other communities could not be resolved by only lightweight mechanism. For example, adolescents successfully used interaction mechanisms within Push-to-Talk to mediate communication between themselves. Yet those adolescents did not like to use Push-to-Talk to communicate with their parents, as they found that parents misinterpreted and misused those interaction mechanisms.

An alternative approach to manage accessibility is automated interruption management. In contrast to interactive approach, the system (rather than the user) takes the role of an interruption mediator brokering or filtering interruption attempts. Examples include the Personal Reachability Management Systems (Reichenbach et al. 1997) and the Personal-Level Routing (Roussopoulos et al. 1999). Both used a user's personal privacy rules to offer automated availability management aimed at minimizing recipients' effort when dealing with undesired communicative attempts. A drawback of such automated solutions is that a special effort and premature commitment is required from users, as they have to set privacy rules explicitly and a priori.

Empirical evidence on the use of customized software (Mackay 1999), an experimental groupware calendar system (Palen 1999), and more recently the mySpace portal (Patil and Lai 2005) showed that users are unlikely to configure systems to manage their privacy in this way. Most people perceived the task of creating initial privacy settings as a burden, and in most cases just adopted the default settings. Patil and Lai (2005) further investigated the effect of explicitly disclosing the information being shared by the system, where they expected that people would increase their privacy configuration. Surprisingly, they observed no such effect. Similar results have been reported in more recent studies of privacy within social networks. For example, Lewis et al. (2008) examined a population of 1,710 students: 67 % had fully searchable public profiles, which was the default system setting. Lederer et al. (2004) evaluated the Faces prototype for managing personal privacy in ubiquitous computing settings: while users did configure a priori privacy settings, their preferences often did not match their expectations when confronted with realistic scenarios. The result is that users suffer a privacy breakdown in these conditions. These findings are perhaps unsurprising, as a discrepancy between attitudes and behaviour in the domain of privacy has not only been discussed before (van de Garde-Perik et al. 2008), but also portrays the difficulty of deciding upon privacy preferences outside a concretely specified context.

Work from McFarlane (2002) and Patil and Kobsa (2004) reported that people should be given mechanisms that allow them to both assess and announce interruption moments when opportune. In particular, Patil and Kobsa (2004) reported peoples' need to resort to different strategies to keep their interaction at desired levels by dynamically controlling how they appeared to others. For example, participants' strong desire to control their availability while working was manifested in different strategies: leaving descriptive status messages even if away for 5 min, turning-off the automatic idle indicator to not give the wrong impression of being away, and using different pictures when logged from home or work as a way to indicate one's location. These results were confirmed and extended in an experiment by Romero et al. (2007b): when the system shielded recipients from interruptions, they tend to develop strategies to manually overcome such automatic protection in order to collaborate with the interrupter. In the automatic protection condition the system processed automatically the availability of the recipient by assessing their performance in an individual quiz, blocking interrupters if progress was slow. Recipients in some cases would cheat (by just entering dummy answers to the quiz) to make the system believe that their performance was good enough to allow interruptions. In this way, interruptions were not blocked and recipients could at least help their partner.

A common theme in the afore-mentioned related work is that any mechanism that supports privacy—whether automatic or manual, or for regulating information or interaction—must be lightweight. At the same time, it must be effective if it is to align with established as well as emerging social practices in the coordination of interpersonal privacy. Yet this is not a panacea. In complex scenarios where social tensions could be more visible between participants, lightweight mechanisms seem insufficient. For example, Wiberg and Whittaker (2005) used their Negotiator system to show how manual availability management might create social tension for the interacting parties and thus incur cognitive costs when conflicts occur. To our knowledge, no interactive solutions balance the need for lightweight mechanisms *vs.* more effective mechanisms required to ground conflicts when necessary.

3. Coordination of privacy as border regulation

The previous findings of privacy largely result from evaluations of systems in use. We now turn to theories of privacy, which provides a more general and perhaps more comprehensive perspective on individuals' behaviours regarding privacy coordination. In particular, we build upon Altman's (1975) theory, as it has exerted substantial influence in how researchers in the field of human computer interaction and mediated communication view privacy issues surrounding networked applications, e.g., as used by Dourish (1993).

Altman (1975) defined privacy as a **border regulation process**, where individuals use mechanisms to open or close their borders to regulate if and how interaction takes place within their environment. Altman's definition acknowledges that individuals modify and continuously re-assess their borders in response to stimuli by their environment and their own needs for social interaction.

Figure 1 shows how this boundary regulation process can result in desired and undesired states of interaction with the environment. The figure uses a continuous line to represent an individual's border that is closed to the environment (preventing any interaction) and a dashed line when the border is opened (allowing interaction to take place). When a person attempts to initiate an interaction with someone in the environment, this is shown as an arrow pointing from the person (U) to the environment (E) and vice versa. The diagrams show the possible cases of desired and undesired outcomes for an attempt to interact.

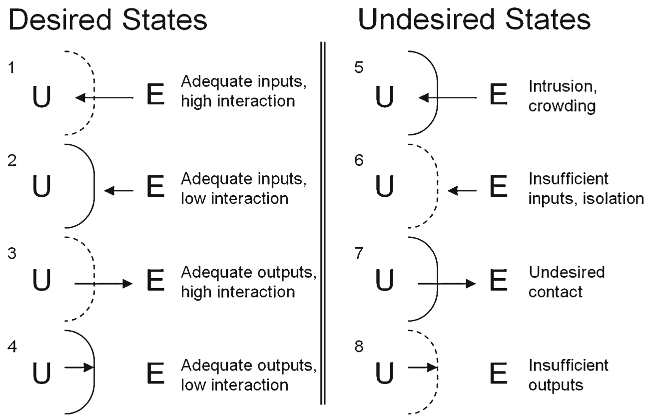

Figure 1. Altmans's model of privacy borders regulation, describing four desired states of interaction (*left*) and four undesired states of interaction (*right*).

For example, the desired state represented in (1) illustrates situations such as when a person (U) opens the door to let someone (E) get in, while in (4) a person is glancing away to avoid unwanted conversation with someone. From the undesired scenarios, (6) illustrates situations such as when a person invites someone for communication, who does not engage (what could cause feeling of isolation). Alternatively, in (7) a person is hoping to avoid someone, but had to spend an evening together (what could violate the wish to be alone).

Altman's perspective on privacy acknowledges but leaves implicit interpersonal interactions that are crucial in mediated settings. Privacy borders are *not* set unilaterally by an individual but are regulated by both communicating parties as a result of pursuing their individual needs and responding to the needs of others. Such a broader perspective is advocated by Petronio (2002). She extended Altman's theory by introducing the element of collaboration in privacy regulation: both sender and receiver are mutually responsible over the information exchanged. Petronio's framework encompassed the perspective of both communicating parties. It describes how the parties develop and use rules to agree on whether to disclose certain information, and on which basis each party takes responsibility over the information disclosed. Petronio also acknowledged the importance of understanding other signs (outside the exchange of information), i.e., to interpret cues of when is the right time to disclose something, to whom, how, and so on. She recognized that this 'implicit' information seems to be crucial for the disclosure decision.

Boyle and Greenberg (2005) developed a privacy vocabulary that builds upon the theoretical work of Altman and others, where they precisely define and describe the nature of the privacy concerns in the context of interpersonal coordination needs (see also Boyle et al. 2009). At a high level, they describe the process of privacy regulation by means of **solitude**, **autonomy**, and **confidentiality** controls. These concepts refer to people's needs to control desired level of interaction (availability), to identify one's behaviours within the community (whether collaboratively or individualistic), and to control access to personal information (fidelity, accuracy) respectively.

In conclusion, these prior works collectively argue for the collaborative and dynamic nature of privacy border management, especially for managing interpersonal privacy in mediated communication. However, there is as yet no account of these mechanisms, nor is there a theoretically motivated understanding of how best to design interactive controls for such systems.

4. Privacy grounding model

We now introduce our own Privacy Grounding Model (PGM). PGM is a descriptive model that provides a generic characterization of the social practices detailing interpersonal privacy coordination activities in mediated settings (Romero and Markopoulos 2009). Reaffirming the previously mentioned prior work, PGM describes how privacy needs and preferences derive and evolve in a fluid way during the course of the interaction between people.

We base our model primarily upon the Common Ground theory of language use (Clark 1996). This theory emphasizes how people engage in a dynamic collaborative process to develop the **common ground** necessary for the success of their communication. According to Clark, people develop common ground to communicate meanings to each other efficiently; in turn communication helps develop this common ground further. Although the Common Ground theory was originally conceived to describe human behaviours in face-to-face communication, it has been successfully used to make predictions and inform the design of system-mediated communications (Clark and Brennan 1991; Monk 2003; Aoki and Woodruff 2005). We further frame the model using Altman's theory and the Privacy Vocabulary (Boyle and Greenberg 2005; Boyle et al. 2009) to represent the coordination of privacy as a regulation of interpersonal borders.

The model is structured as three layers of abstraction that connect theoretical concepts of coordination (components) with their corresponding behaviours (mechanisms and characterizations), as illustrated in Figure 2. In brief, the first **components** layer includes the concepts of Common Ground theory, collaboration, signalling and grounding and the concept of regulation to describe privacy coordination as a collaborative activity. The **mechanisms** layer describes the ways in which communicators operate each component to coordinate privacy borders. Finally, the **characterizations** layer describes the different elements of the mechanisms. These are discussed in turn below.

4.1. Components

The **Components layer** includes three key concepts of Common Ground theory: **collaboration**, **signalling**, and **grounding**. The fourth **regulation** component is

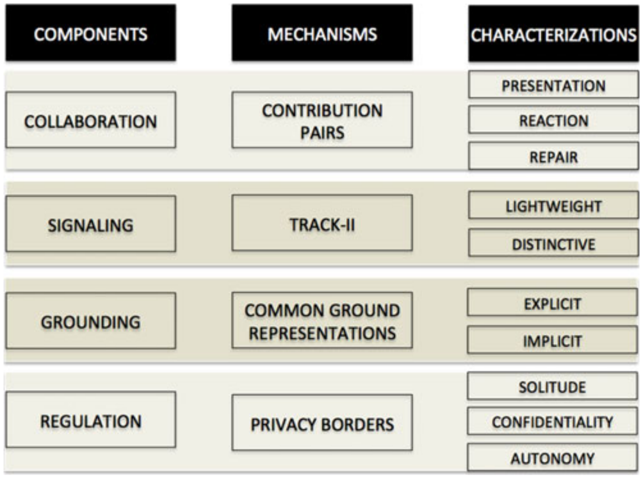

Figure 2. The Privacy Grounding Model (PGM).

based upon Altman's privacy theory. *Signalling* is defined by Clark's theory as the action of communicators to represent and express the meanings of their message. *Grounding* is the joint action of communicators to recognize what the first party meant. Such joint actions accumulate communicators' common ground of their meanings and understandings. *Regulation*, as introduced earlier, is defined as the process where individuals use mechanisms to open or close their borders to regulate if and how interaction takes place within their environment.

PGM describes privacy coordination as a collaborative activity to regulate interpersonal privacy borders. Signalling represents the individuals' action of opening and closing borders to express intentions for interaction with others; and grounding describes the collaborative action that builds a shared understanding of whether a border is opened or closed.

4.2. Mechanisms

The second layer of PGM describes the social mechanisms defined in Common Ground that are associated with collaboration, signalling, and grounding in the frame of regulating privacy borders.

In Common Ground, collaboration models coordination activities as **contribution pairs**. These pairs involve presentations, reactions and repairs by which communicators try to reach closure for each signal used to communicate. An example is the following initiation-reaction pair: 'hi, coffee?'—'sure, give me a moment'. Initiators contribute with presentations of their intentions to interact and recipients contribute with reactions to communicate their understanding of those intentions. These paired contributions are ruled by the principle of *least collaborative effort*: communicators in a joint activity try to express no more than what they perceive as sufficient for advancing the current communication.

Signalling, in Common Ground, is supported by the use of at least two tracks to communicate both the content of communication (**track-I signals**) and the means to coordinate the process of communication (**track-II signals**). *Track-II signals* are therefore the mechanisms used by communicators to represent their privacy intentions when engaging in interaction with others. Examples of track-II signals in a Instant Messaging system are: updating online status, blurring video, changing presence name in the buddy list, or writing explicitly about one's openness to interactions.

Grounding is a key component in Common Ground theory: communicators establish a shared understanding of track-II signals by the development and use of **representations of common ground**. Communicators need such common ground representations if they are to correctly establish a shared understanding of their intentions to communicate. Examples include social rules, shared experiences and shared knowledge of a situation.

Finally, the mechanisms to support the regulation of privacy are the control of **privacy borders** to open and close interaction with others. By means of track-II signals, a privacy border is shared with others, and by means of contribution pairs, common ground representations are jointly developed to agree on the action towards a privacy border.

4.3. Characterizations

At the third layer, *characterizations* typify the social interaction space that reflects the choices made by communicators to develop their coordination practices. Such decisions are made depending on the amount of effort communicators wish to invest and how implicit or explicit they want their coordination intentions to be made.

Collaboration choices are characterized by the use of **presentations** to signal intentions for interaction, and **reactions** or **repairs**, to communicate (dis) agreement or to fix a previously produced presentation, respectively. In Instant Messaging, for example, communicators could use consecutive installments (separate by a small time frame) as separate text messages to repair the intention of the message: "Hi" [enter—1 min] "I'm back" [enter—1 min] "are you there?" In this case, after each signal, the initiator assessed the (least collaborative) effort to be more distinctive in order to achieve an understanding with the recipient of his intentions to interact with her.

Signalling choices are represented by the four characterizations of track-II signals in Common Ground, classified in **lightweight** (brief, background, and simultaneous) and **distinctive**. Track-II signals thus must be in the background and brief: they have to be undemanding in their presentation if they are to allow track-I signals to be more prominent. They must be simultaneous to track-I signals so that the coordination can occur at the same time as the content is presented. They must be distinctive enough from track-I signals so that the recipient can easily understand its coordination purposes. These characteristics describe the effort communicators commit to produce and attend to a track-II signal (in other words how lightweight the signal is), and how implicit or explicit their coordination intentions are made, i.e. how distinctive the signal is.

Grounding choices are characterized by the need of communicators to **explicitly** develop and use common ground representations and therefore incur in additional effort, or whether they can **implicitly** develop and use common ground representations and therefore minimize the effort needed for grounding.

Finally, privacy borders regulation is characterized by the control of solitude, confidentiality, and autonomy. They describe the interpersonal purpose communicators have to engage in the regulation process, characterizing their relation with the individual or community they are interacting with.

4.4. Discussion

The name Privacy Grounding Model suggests that communicators engage in collaborative practices to establish a common understanding of their privacy intentions. This is similar to (and obviously builds upon) how the Common Ground theory explains people's discursive practices. The components of PGM aim to provide descriptors that characterize the interactions people develop to coordinate interpersonal privacy during their mediated communication process.

Signalling and grounding support the collaborative coordination process by providing lightweight mechanisms to establish a common understanding of privacy intentions, even in cases when someone is breaching a misunderstood border. For example, if someone initiates communication even after the recipient has blurred her video to represent unavailability, the recipient could explicitly signal her meaning of unavailability intended by the blurred video, which grounds the recipient's unavailability and how it affects the recipient's responsiveness.

A previously described limitation of mediated communication technologies is that representations of collaboration are not always as efficient as in face-to-face settings. For example, collocated communicators can use verbal silence as a lightweight reaction, where its understanding is easily grounded on the basis of other simultaneous signals including body gestures (nodding, leaning forward, etc.) and physical actions (moving away, opening a book, and so on). In mediated communication, silence is rarely presented along with other signals. The lack of simultaneous channels makes it difficult for the recipient to ground silence in a lightweight manner, and for the initiator to deduce the intentions of recipient's silence. To address such situations, PGM elucidates the needs for signalling and grounding that allow lightweight signals such as a silence to be used as acceptable contributions.

The theory of Common Ground describes how people are continuously grounding their intentions to communicate. Similarly, social interactions require different levels of ambiguity to succeed arguing that, for relationships to run smoothly, they assume a level of ignorance and mutual concealment. Aoki and Woodruff (2005) discuss how communication systems should allow for ambiguity to support the way individuals manage their self-presentation in their social network, and how they react to others. More than hiding the truth or convincing others of a lie, the idea of ambiguity defended by Aoki et al. is to multiply the possible interpretations that could explain a situation so that negative explanations can be avoided if mutually desired. One could think that grounding and ambiguity are the opposite extremes in the spectrum of mechanisms to coordinate interaction. However Aoki and Woodruff observed that the preference for providing ambiguous explanations about certain behaviours (and therefore admitting of multiple interpretations) could better support grounding than no explanation at all. Begole et al. (2004) came to the same conclusion, but from a different angle. They developed the Lilsys availability awareness system, which used sensors to detect and display peoples' availability to each other, where their idea was that such representations would make people accountable when initiating or responding to interactions and thus support social norms. Yet they also observed that the information provided was too precise: there was too little room for ambiguity. Compared to common ground, we note that taking ambiguity into account goes beyond the notion of least collaborative effort.

The remainder of this paper presents a field study conducted to provide evidence of PGM's ability to describe the dynamic and collaborative practices of interpersonal privacy coordination. The field study aims to capture the complexity of real social context and the dynamic of long-term interactions to analyse interpersonal privacy management in mediated communication systems. The context of the study is to analyse the phenomena of interpersonal privacy in mediated communication involving a highly collaborative small online community. The community shares the goal of people exchanging information and interacting with others, but still face challenges when coordinating their needs for interaction. Therefore, with the presented study we aim to observe, collect and analyse interpersonal privacy interactions of highly collaborative online communities in the field.

5. Field study: community bar

The study presents our analysis of the interpersonal privacy practices surrounding the use of a desktop awareness and communication tool called the Community Bar (McEwan and Greenberg 2005; McEwan et al. 2006), or 'CB' for short. (Section 5.2 describes how Community Bar works). The target community consisted of academic workers who were already using the Community Bar regularly for both work-related and social communication. The first author (referred below as the 'ethnographer') acted as a participant-observer. Detailed logs of interaction with this system, including not only exchange of text messages, but also presence video, pictures, and others, spanned 4 weeks and were triangulated with ethnographic data gathered by interviews, naturalistic observations, and diary logs.

We observed (via participant observation) and logged the communication behaviours of the community. We then analysed and described this community's coordination signals using a PGM coding scheme based on the characterizations component of the model.

We note that, unlike most field studies of system use, our purpose was *not* to critique Community Bar's capabilities, nor was it to provide a full account of that community's practices. Rather, the objective of the analysis was to identify confirming and disconfirming evidence regarding key aspects of the PGM: (a) manifestation of collaborative practices in privacy border regulation, (b) track-II signals use to represent different levels of distinctiveness and lightweight privacy borders, (c) pairing presentation and reaction signals to ground privacy borders representations, and (d) using grounding to regulate solitude, autonomy, and confidentiality borders. Ultimately, we wanted to illustrate that PGM provides an adequate description of the dynamic and dialectic process of interpersonal privacy coordination.

5.1. Research methods

We chose field studies *vs.* experimental tests as found in hypothesis testing for a variety of reasons. Field studies allow us to collect richer information compared to those experimental tests most commonly used to evaluate communication systems. We are also concerned that such experimental tests have repeatedly failed to reveal any privacy concerns with respect to awareness systems (e.g., Miller and Stasko 2001; Metaxas et al. 2007). We believe this is because observations done in the safe context of a research study, with a well-defined social context (e.g., fixed information content and recipients), and for short time periods are limited to understand discrete actions regarding privacy. They don't provide and understanding of how privacy needs evolve over time and how it is negotiated. Moreover, participants' behaviours and reports are influenced by the artificial settings (van de Garde-Perik et al. 2008).

To study privacy as a social construct, we have based our research methods on Case Study (Yin 1994) a suitable methodology for answering the 'how' and 'why' questions without requiring control over behavioural events. This methodology is appropriate for covering contextual conditions of contemporary phenomena to validate a theoretical model developed as a characterization of such conditions. Using this methodology, we defined the unit of analysis based on PGM elements, i.e., as sequences of presentations and reactions (or repairs) to intended initiation of interaction, which represented a coordination unit.

Our sampling and data collection technique consist of automatic data logs of all system interactions, diaries and interviews (for more detail refer to next section Data Collection). Qualitative analysis based on coding techniques (Miles and Huberman 1994) to analyse the data logs and to structure the diaries and interviews questions. The coding scheme was developed based on the elements of PGM, with the purpose to describe the instances of privacy interaction in the logs.

5.1.1. *Pilot study*

Prior to our field study, we conducted a 1 week pilot study to verify that the research strategy described above was indeed appropriate in our context of choice, and to gather a first impression of PGM's descriptive capabilities to describe behaviours as gathered in a field study. While also characterised as a field study, the pilot setting was slightly more artificial as. Although the context was real (a community of eight researchers working on a common project), the mediated tool they used—Exodus Instant Messenger—was replaced by one customized for the study. We did this as Exodus had better capabilities for data gathering, and also allowed us to control the available features of the Instant Messenger too. The pilot showed that our research method did allow us to observe evidences of interpersonal privacy coordination, and that the components of PGM were manifested in mediated communication. It also demonstrated that PGM could describe interpersonal privacy border regulation as observed 'in the wild'. Romero (2008) provides complete details of this pilot. We omit discussion of our pilot study results here, as they are very similar to those found in our primary field study, discussed next.

5.1.2. *Data collection*

In our primary field study, we used four methods to gather qualitative and quantitative data regarding users' interaction practices with Community Bar, their understanding of privacy representations in the system, and their reflection on privacy related issues.

First, we implemented a data logger to collect data on people's use of the Community Bar. It registered participants' interactions with the system by collecting both context information (time, place, subject, target) and content information depending on the media item (text, URL, image file, etc.). The data logger captured a total of 13 types of contribution depending on the media items available. From the presence item, the logs collected updates of video snapshots, availability status (both set manual and automatically), pictures, buzz notifications and text messages. Data collection from the chat item included the sender, content, and time of each new message. From the web, sticky notes and photo items, collected data included the sender, content and time of URL's accessed, notes, and photos respectively. Finally the logs also captured the time when users logged on and off the system.

Second, diary logs captured in-situ reflections by participants about their experiences with the system. As mentioned, the diary was implemented as a media item in CB (see 'log item' in Figure 3), autonomously prompting generic and event-related questions to encourage participants' contribution. It randomly selected a general question every hour and sent it to all users. Some questions were event-related: these were provided only to the participants involved in the specific event soon after it took place.

Third, direct naturalistic observation helped the ethnographer acquire insights regarding the community, its structure, and the type of relationships between its members and the embedding of CB usage in these.

Finally, open in-depth interviews were conducted twice. At the beginning of the study, we interviewed eight members of the community, some who actively used CB and others not. At the end of the study, we individually debriefed 11 participants via semi-structured interviews, focusing on their interpersonal interactions and privacy concerns when using CB. Participants were asked to comment on: perceived benefits from CB; the ease and effectiveness of assessing others' reachability in CB; the use of social rules to assess acceptable behaviours; and whether collaboration was considered necessary to achieve a satisfactory use of the system. All interviews were audio-recorded, transcribed and analysed qualitatively.

5.1.3. *Analysis procedure*

A mix of qualitative and quantitative analysis methods was chosen to identify and describe instances of initiation of interaction, and reaction to initiations.

The qualitative analysis included two coders and used the cross-checking technique (Miles and Huberman 1994) to increase the reliability of the analysis. We calculated inter-coder reliability using Miles and Huberman formula (1994, pp.64), which provides a ratio of the total number of agreements to the total number of codes (agreements plus disagreements).

The coding procedure was as follows. First, the two coders identified contribution pairs by describing every signal as a presentation (initiation), reaction or repair, and coupled with another signal to form a contribution pair. Second they described the signals using the four track-II characterizations indicating whether lightweight (brief, background, simultaneous) and distinctiveness were present or not in the signal, using 'yes' (+), 'no' (−) or 'to some extent' (+/−). For the grounding component, all signals were characterized as to whether they explicitly grounded (E) or implicitly grounded (I) interactions.

We used the coding scheme to triangulate the three main data sources: the data logs were coded to identify and characterize interpersonal privacy interactions (unit of analysis); interviews and diaries were designed and structured using the same scheme. This triangulation aims to link coded patterns from the data logs (interactions monitored) with intentions/perceptions of users, allowing us to answer the 'why' question.

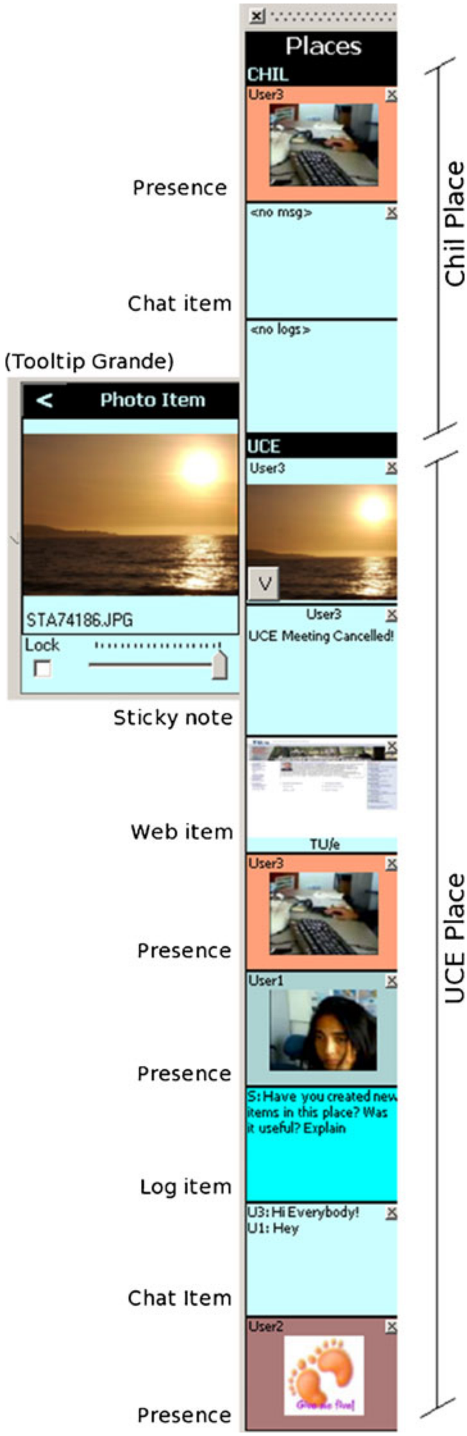

Figure 3. Community Bar—sidebar.

For the analysis, we drew a distinction between media items in CB that are primarily suited for content exchange (track I) and those suited for coordinating the communication (track II).[1] We characterized chat, photo item, sticky notes and web items as **content channels** and presence messages, status, video and picture updates as **coordination channels**.

Finally, basic statistics were applied to our coded units to illustrate interaction traffic and the nature of the interactions observed. In the analysis, the coded units were quantitatively classified to identify patterns of used and relationships between the units.

5.2. Study design

The community under scrutiny had already been using the Community Bar (CB) system for 2 years. Its design was intended to support awareness and casual interaction in small communities (McEwan and Greenberg 2005; McEwan et al. 2006). The group included experienced and novice CB users. Three of the 15 participants of this group had worked on the development of CB, while two others had been involved in the design of the Notification Collage (Greenberg and Rounding 2001), a predecessor of the CB system. The other 10 had no or minimal involvement with CB development.

The 15 participants formed a small cohesive academic group: 11 of them worked at the same research laboratory, 3 were former graduate students now working for other companies or institutions (2 of them in different cities), and 1 member was an external researcher contacting other members of the community only occasionally. The 11 collocated members consisted of one Professor and ten Master and PhD students, out of which five were under the Professor's direct supervision. As mentioned, all members had been using the system prior to the study, ten of them frequently and four intermittently. The ethnographer joined as a temporary member of the community for a period of 4 months; she had used the system for 1 month prior to the actual logging of system interactions reported below.

The community's physical work environment consisted of an open area for students, and few offices around it for the professors and meeting rooms. Four members telecommuted twice a week. Master and PhD students would often stay until late in the evening, or come during weekend to work, or do social activities together (e.g. sports or play computer games).

At the time of this experiment, members of the community used CB on a voluntary basis. Its main benefits were its high fidelity awareness of those who were online (via regularly updated snapshot video and/or status indicators) and its video and textual communication capabilities. People usually used it to engage in brief work or socially related interaction. The telecommuters, like the Professor, had special interest in using CB as it provided the possibility to easily reach other colleagues and students when working from home. Most participants used CB as the main channel to connect with colleagues, but not exclusively. The public nature of the system (everything is broadcasted to everyone) and its restricted audience (limited only to the particular 15 members in the research community) led most participants to use other communication applications in tandem. For example, Email and Instant Messaging systems allowed CB users to engage in more private (one-to-one) conversation as well as to reach a larger number of people.

A webcam regularly posted snapshots of one's presence every minute or so to the community. Its use was recommended but not obligatory. Most participants had web cameras installed at their desks, and the ones who did not could optionally request one for the study period. For the study, four participants did not use a webcam: two because of technical problems or company policies; and two because of personal choice. Participants were asked to use CB as usual, where they could also use other communication media if desired.

The role of the ethnographer during these four months included the following tasks:

- individual interviews with eight group members with the purpose of learning about the community, its members, and their relation with CB (first month);
- the design and execution of the study (note: to minimize influences of the ethnographer in the data logs, all interactions between her and the community took place outside CB during the study); and
- individual interviews with 11 participants to discuss their experience with CB during the study period (fourth month).

An adapted version of Community Bar was installed on participants' desktop and laptop computers. Figure 3 illustrates some of its features. The Community sidebar is an always-visible bar positioned at the side of the screen (it remained visible even if other applications were set to maximized view). The sidebar provides a shared space where users can post and see 'media items'—interactive media that could be posted by community members—as small tiles. Media items could take the form of video, text notes, web pages, availability status, digital photos, etc. The design intent of the sidebar's position on the screen and its media item content was to provide information about the whereabouts and activities of others at a person's periphery of attention. When a person passes the mouse over the tile view of the media item, a larger 'tooltip grande' appears (Figure 3, left side) containing more information and/or interactive controls. The user could also double click that tooltip grande to raise a full-sized window running the highest fidelity and most interactive version of the media item.

Visual notifications were provided to inform users about new items and updates of existing items. A notification turns the background colour of the tile view of the new or updated item brighter until the new information is acknowledged by the viewer (by moving the mouse over the tile).

The media items available at the time of the study were:

- Chat items, representing a textual space for public conversations;
- Presence items, representing presence of people (optionally as video snapshots) connected to the system;
- Sticky notes, used as public static text;
- Web items, to publish web URLs;
- Photo items, to share pictures.

We also added a diary media item, which occasional prompted a participant with a particular question about his or her particular experiences using the system, or where a participant could contribute text detailing their experiences at their leisure.

The adapted version also had a buzz button in the presence item, to attract the attention of others. Clicking the buzz button of the presence item of the person to buzz (the target), generated a visual notification in the target's sidebar; the presence item of the buzzer flashed for a few seconds or until the target user moved the mouse over it. This option was implemented in response to the wishes of CB users, where they wanted an explicit mechanism to coordinate the initiation of an interaction.

5.3. Results

The average number of users logged per day during the first three observation weeks was stable (8.1, 7.4, and 6.6 respectively), but it dropped to 3.6 in the fourth week due to holidays. The peak attendance was 11 users logged on at the same time, which was observed in the second and third weeks during weekdays between 09.00 and 18.00. In the afternoon intervals (from 12.00 to 15.00) between 6 and 10 people were usually online. Several people were typically logged on the system during evening intervals (from 18.00 till midnight).

Table 1 presents the frequency of the interaction signals used by the participants.[2] Interaction signals were classified as either system notifications or user interactions. System notification signals (96 %) represented information that was broadcasted automatically by the system, including video, automatic changes of status, and online/offline indicators. User interaction signals (4 %) represented information generated by explicit user's actions. Most user interaction signals occurred in the content channels (94 %), where the chat message was the prominent signal observed (98 %). The remaining 6 % of user interaction signals corresponded to signals in the coordination channels with presence message (49 %) and buzz (30 %) being the most frequently used.

5.4. Analysis

Using the PGM coding scheme, the two coders identified and described 501 contribution signals related to privacy coordination. Reliability rates between

Table 1. Frequency of interaction signals.

Signals	Frequency (unit)
System notifications (96 %)	**24260**
Video	21975
Automatic Status	1922
Online/Offline	363
User Interactions (4 %)	**1017**
In content channels (94 %)	960
Chat	945
Photo, sticky notes, web	15
In coordination channels (6 %)	57
Presence message	28
Buzz	17
Manual Status	8
Presence picture	4
Total coding lines	**25277**

coders were calculated for each coding category (collaboration, signalling and grounding) and were higher than 80 %, which is very satisfactory; no disagreements remained after discussion. See Romero (2008) for a more detailed report.

The contribution signals included initiations, reactions, omissions, delays, and presence signals that trigger initiation. Delays included reactions to an initiation sent 10 min or longer beforehand. The threshold of 10 min was set to reflect the intermittent nature of chat conversations, where reactions of about up to 10 min are quite usual. In the following subsections we describe how these signals were characterized according to the four PGM components.

5.4.1. *Collaboration*

To describe collaboration practices, we looked at the characterizations of contribution pairs: presentations, reactions and repairs.

Two types of user interaction signals represented more than half (58 %) of all privacy presentations observed. They were text messages in the content channel (39 %); and presence messages, nimbus, alert, invitation, picture and manual status change in the coordination channel (19 %). The second most frequently coded presentation referred to system generated signals in coordination channels (40 %), which described the reference to video links, automatic generated status, and users' connection to the system; and signals generated by users in coordination channels (2 %) described by sticky notes, web links, and photo sharing, which were coded as presentations to potential initiator of conversations. System generated representations were identified as coordination contributions when they related to a user generated signal. Such relations were described as triggers or reactions to user generated contributions. To illustrate, Table 2 is a

Table 2. Contribution pair with a system generated signal (video) triggering user generated signals (chat message) to coordinate the end of a long conversation.

Date/Time	Action	Content	User	Channel	Br	Ba	Si	Di
12:14:28	Msg	You're just upset the team1 had to catch up to the team2 XD	P10	Chat	+	−	+	−
Intermittent social conversation for almost an hour about hockey between P10 and P15								
12:47:42	Msg	He's a presence even if he doesn't score	P10	Chat	+	−	+	−
12:50:26	Video		P10	Presence	+	+	+	−
12:50:59	Msg	okay ;)	P15	Chat	−	−	−	+
12:55:26	Auto	Away	P10	Presence	+	+	+	−
12:58:22	Auto	Away	P15	Presence	+	+	+	−

brief extract from the data logs[3]: it includes the use of video (system generated signal) to signal the end of a conversation (i.e., someone else has just entered the room and as a result P10 has stopped chatting), which in turn triggered a chat text (user generated signal) to ground the delay (i.e., P15 confirms that he understands the conversation is over).

Participants' contributions to initiate communication varied widely. The daily average of initiations per participant was nine. Three participants never initiated any conversation (P3, P4, and P6). The interviews and diaries indicated that participants developed different collaboration strategies depending on the nature of their membership in the community. On the one hand, the 'closer' a person was to the core community,[4] the more common ground already existed, and therefore the less distinctive but more lightweight signals were used for coordinating initiations. On the other hand, explicit collaboration using more distinctive signals (and that made interaction somewhat more heavyweight) were recognized by the core group as a social conduct to maximize awareness and interaction benefits.

> 'What I would usually do is a very quick thing like saying I'm on the phone … usually I kind of say I hear you but not now …and sometimes I say I can't talk now … I usually kind of point to the phone in the video so it gets updated' (Interview Participant 13)

Even in the situations in which individual behaviours conflicted with the general social practices, collaboration was considered important. For example, a participant's decision to not use video was mostly compensated by this person expending extra individual effort to maintain a desirable level of connection with the community. Otherwise such participants felt like they were being treated as second-rate participants and experienced a decrease in their privileges as full members.

> ‘Because not having a webcam does require more effort [for me] to project [my] status’ (Interview Participant 10).

The use of delayed reactions was another example of conflicting expectations between the community and the individuals. Participants reported that delays (18 %) and omissions (17 %) were acceptable coordination mechanisms. In addition, the logs showed that most delays as well as omissions of initiations were not explicitly confronted with messages such as ‘why are you not responding’. In contrast, participants also reported that such mechanisms were less optimal contributions as they left them wondering why is he/she not responding.

Table 3 provides an example interaction reflecting the above insights. Based on the solid common ground between P13 and P15, P15 incurred minimal effort to initiate interaction with P13 (based on P13’s representation of his ‘On’ connection to the system). P15 assumed that the online status of P13 meant that he was available to receive a question without any prior coordination. Similarly, the delayed reaction of P13 was automatically coordinated by the changes of his status, which P13 assumed to be sufficient to coordinate his delay; therefore he reacted contributing incidentally to the content of the conversation instead of explicitly explaining the reasons behind his delay. In addition, P15 repaired the initiation with a buzz, which was likely triggered by P13’s automatic change of status to online. Finally, even when P13 made explicit his interest to contact someone else (P3), he added extra effort to explicitly ground that to P15.

In summary, the need to acknowledge non-collaborative mechanisms confirms that privacy border regulation was not just about allowing or disallowing communication, but also about repairing ‘damage’ after a certain communicative

Table 3. User generated signals (chat and alert) and system generated signals (delays) to coordinate initiation of conversation.

Date/Time	Action	Content	User	Channel	Br	Ba	Si	Di
10:01:54	Conn	On	P13		+	+	+	−
10:02:32	Msg	Hey P13—I have an idea for the social paper	P15	Chat	+	−	+	−
10:07:20	Auto	Away	P13	Presence	+	+	+	−
11:02:34	Auto	Online	P13	Presence	+	+	+	−
11:15:16	Auto	Away	P13	Presence	+	+	+	−
11:18:50	Auto	Online	P13	Presence	+	+	+	−
11:19:27	Buzz	To: P13	P15	Presence	+	+	−	−
11:23:32	Msg	P3, Would you phone me?	P13	Chat	−	−	−	+
11:24:13	Msg	P15…um the deadline for GI is today at 5 pm…	P13	Chat	−	−	−	+

attempt was neglected. The need to repair could be motivated by politeness, but in most cases the purpose to repair related to the need of balancing a situation where less collaborative behaviours were used (e.g., not using video or not responding a message).

5.4.2. *Signaling*

Coordination signals were described as lightweight depending on the brief, background, and simultaneous track-II characterizations, which reflected the effort needed to produce a signal: we consider an automatic status change lighter than a text message. The distinctiveness characterization of the signal reflected the effort needed to understand whether a signal was meant as a coordination signal or not.

Following the classification of the 501 coordination signals as being produced in content or coordination channels, we identified four different characterizations regarding lightweight: brief (br), background (ba), simultaneous (si) and distinctiveness (di) (see Table 4).

In coordination channels, participants communicated their intentions to interact based on two types of presence information: system generated signals (video broadcast and automatic changes of online status) and user generated signals based on self-presentation status (presence message, online status, and picture). The former were considered lightweight signals (brief, in the background, and simultaneous), while the latter were less lightweight since they could not be generated simultaneously to content signals (changing one's status required participants to select a new status from a checkbox or adding or removing text). Both system and user generated signals were characterized as not distinctive. In the analysis, the distinctiveness of such signals was related to the existing common ground among the parties involved. In cases when common ground was insufficient, participants engaged in extra effort to understand the coordination purpose of a signal. For example, in some cases a blurred video required explicit coordination to understand one's intentions of interaction (e.g. as in the example

Table 4. Signalling characterizations classified in coordination and content channels.

Coordination Channel (58 %)				Content Channel (52 %)			
Br	Ba	Si	Di	Br	Ba	Si	Di
+	+	+	−	+	−	+	−
System generated presence representations that are used to coordinate the initiation or end of an interaction (30 %)				Text signals that initiates interaction by directly contributing to the conversation (13 %)			
Br	Ba	Si	Di	Br	Ba	Si	Di
+	+	−	−	−	−	−	+
User generated presence representations that are used to coordinate the initiation or end of an interaction (28 %)				Text signals that seek establishing a share understanding of intentions to interact (29 %)			

extract in Table 5). Yet in other cases that blurred video was sufficient for one party to understand that the other party was using it to signal conversation unavailability.

In content channels, participants created two types of text messages to communicate their intentions to interact. If no explicit coordination was considered necessary, communicators initiated their interaction directly and thus contributed to the content of the conversation. If explicit coordination was needed, participants typically exchanged text messages before content was communicated to first ensure a shared understanding of their availability. The ‘no explicit coordination’ was characterized as brief and simultaneous to the content signals but not in the background, as the initiator pushed the initiation without using track-II signals (or ignoring existing ones). Similar to what was observed in coordination channels, these types of signals were only distinctive when sufficient common ground existed. The explicit coordination signals were described as distinctive but non-lightweight contributions, used by participants primarily to ground existing system generated coordination signals. Table 5 shows an example, where text is used as explicit coordination to ground the meaning of the video representations.

Distinctiveness seemed to be necessary to guarantee the effectiveness of coordination signals, as the meaning of signals varied over time and between participants. For example, in some occasions an ‘away’ status meant ‘I am not in my office’, therefore a delay was most likely to be expected. In other situations it meant ‘I am in my office but away from my screen’ so a sooner answer could be

Table 5. Explicit coordination to ground the meaning of system generated representations.

Date/Time	Action	Content	User	Channel	Br	Ba	Si	DI
08:41:00	Video	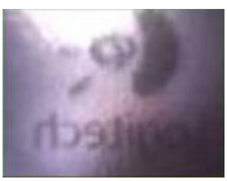	P1	Presence	+	+	−	−
09:10:22	Msg	Are you at home? Is that why shields are up?	P9	Chat	−	−	−	+
09:10:33	Msg	yah at home but actually forgot about shield	P1	Chat	−	−	−	+
09:10:40	Video	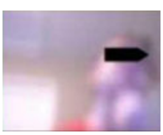	P1	Presence	+	+	−	−
09:10:43	Msg	There you are!	P9	Chat	−	−	−	+
09:10:47	Msg	Still a tad blurry	P9	Chat	−	−	−	+
09:10:59	Msg	yah I adjusted the focus on the camera	P1	Chat	−	−	−	+

expected. From interviews, participants reported the need of distinctive signals, though they also needed these signals to be more lightweight:

> "Because people can't see if there was a visitor [in my office] when people are interacting with me and all of the sudden my status changes, what I need is a very lightweight way for me to acknowledge them … that the reason why … I'm away is because there is someone else here or my computer is crashing or …" (Interview Participant 9).

> "If I'm at home because I don't have webcam I will tend to use this [presence message] more, and it's more to give extra information in terms of where I am, if I'm away or something like that and that is basically because they don't have access to video" (Interview Participant 14)

As illustrated in Table 5, the signals described as (− − − +) represented this behaviour, mostly occurring in content channels as text messages.

5.4.3. *Grounding*

We identified two types of grounding, implicit and explicit, characterized by people's need to provide a distinctive signal or not. The success of implicit grounding (without distinctive signalling) depended on the existing common ground of the parties involved. For example, delays and omissions were considered by the CB community as acceptable lightweight and implicit grounding mechanisms, as in most cases no visible repairs were observed afterwards.

Yet conflicting evidence was found in the interviews. Some participants reported that 'waiting' and 'delaying' was perceived as a frustrating experience, indicating the need to ground such signals explicitly. Common problems for the initiator related to misinterpreting whether their initiation has been ignored, missed or postponed, and therefore misinforming the decision of the next accepted practice: insist or wait. Similarly, for the receiver the uncertainty that a delay could be misunderstood could mislead them to explicitly ground it, even if it proved unnecessary. Therefore explicit grounding shows that participants needed to ground their meaning as privacy borders representations, as the distinctiveness of signals is highly dynamic and subjective. Without grounding, the meaning of existing signals as privacy border representations was likely to be misunderstood or unnoticed when coordinating a particular instance of interaction.

Two types of explicit grounding were observed with the used of distinctive signals (− − − +): to ground existing signals (like delays or video representations, see Table 5 for an example) or to overwrite existing signals. The latter was mostly observed with the use of texts messages that did ignore the presence of existing signals, mostly to represent an urgent or very specific type of interactions. Table 6 provides an example.

Table 6. Explicit grounding using text messages.

Date/Time	Action	Content	User	Channel	Br	Ba	Si	Di
09:56:45	Person	I hate macs	P9	Presence	+	+	−	−
09:56:58	Manual	Away	P9	Presence	+	+	−	−
09:57:38	Msg	Hey P9, working on paper. Do you want to chat for a bit or …	P13	Chat	−	−	−	+
10:55:11	Msg	Hi P13, I was at a meeting	P9	Chat	−	−	−	+
10:55:21	Manual	Online	P9	Presence	+	+	−	−
11:08:50	Msg	Hi P9. Actually I haven't sent it. But if you want to talk I am at XXXX. Maybe 5 min duration?	P13	Chat	−	−	−	+
11:08:58	Buzz	To:P9	P13	Presence	+	+	−	−
11:09:03	Msg	OK. Just a sec	P9	Chat	−	−	−	+

5.4.4. *Regulation*

Regulation was characterized by people's needs to control their interaction (solitude), information about themselves (confidentiality), and their identity within the community (autonomy).

The most common example of confidentiality control was participants' use of presence signals (video and presence messages) to adapt their self-disclosure, i.e., where they increased or decreased the amount of information about themselves made available via the system to others. Tables 5 and 6 contain examples of this.

As we have described earlier, solitude borders were coordinated with different levels of distinctiveness. For example, in many occasions participants added text under their presence representations (whether a picture or video), which gave the opportunity for others to make an educated guess of when to interact with them. Similarly, the way participants use their video could give away their intentions for interaction. A more distinctive way of coordinating solitude was the use of the buzz function, which was often used to repair a previous initiation that had been ignored or omitted, and thus emphasizing the need to get someone else's attention.

The flexibility to not use video, or the use of delays and omissions provided evidence that autonomy control was necessary by their members. This is so even though some members characterized this behaviour (to some extent) as 'non-collaborative' practices. As previously reported, in most cases autonomy control was accompanied by mechanism (explicit grounding) to compensate such 'insufficient' social practices within the community.

5.5. Conclusions

The presented study analysed three sources of data: conversational content, participants' interaction behaviours using Community Bar, and their attitudes and reflections about their interactions within it. Data analysis followed Yin's (1994) framework, and included two coders and additional data triangulation to increase the reliability of the findings. Our approach used social theories, where collaboration was assumed as fundamental to describe interpersonal privacy coordination as a border regulation process based on signalling and grounding mechanisms. Alternatives to PGM were not explicitly considered: most well-known alternatives (like legal and security privacy solutions) did not seem to fit the context of interpersonal privacy studied from a social perspective.

Our results not only confirmed the presence of social collaboration in privacy coordination practices, but also allowed us to identify the nuances in people's need for coordination. We saw that participants tried to find a balance in their collaboration between their needs for coordination and the social and individual effort they were willing to invest.

As mentioned earlier, our purpose was not to critique nor provide a full account of that community's practices when using Community Bar. Rather, the objective of the analysis was to identify confirming and disconfirming evidence regarding key aspects of the PGM. Our analysis, as exemplified in the previous section, has shown that PGM components can provide a detailed account of the process by which (existing) representations of privacy borders are grounded. Signalling privacy borders ensured that its representation became shared knowledge, so it could be recognized by everyone as a privacy border. Grounding a privacy border ensured that its representation was mutually understood, so that everyone could use it as common ground in the coordination process.

In particular, we demonstrated how individuals resolved different strategies for their signalling and grounding. Although Community Bar provided sufficient coordination channels to support lightweight mechanisms to signal privacy borders, in many cases communicators needed more distinctive mechanisms to explicitly ground the privacy implications of their interaction intentions. This finding addresses the definition provided by Altman (1975) and later extended by Petronio (2002) emphasizing the dynamic and dialectic nature of privacy borders regulation. Even further, we found evidence of the trade-off between 'least collaborative effort' (Clark 1996) and the need to invest extra effort to both assess and announce interaction moments when opportune (Patil and Kobsa 2004). Our first conclusion indicates that to support dynamic and dialectic signalling and grounding, lightweight-interactive coordination channels will provide opportunities for communicators to reach a shared understanding of their coordination signals dynamically and with low effort. This is in line with the earlier discussed works (Lederer et al. 2004; Patil and Lai 2005, and others), who opposed fully automatic privacy management solutions.

The collaboration component of PGM characterized coordination signals as sequences of presentations and reactions that varied in their level of distinctiveness representation of privacy borders. If the needs for coordination required special actions, a collaborative effort was represented by a sequence of distinctive signals to develop the necessary common ground representations (e.g. Table 5). If there was no need for such distinctiveness, the required effort was minimized by a sequence of less distinctive contributions used as existing common ground (e.g. Table 2).

As mentioned in the analysis, some participants considered the use of less distinctive contributions (e.g., delays and omissions) as less collaborative when missing a minimum shared understanding. Others considered the use of distinctive contributions as less collaborative when existing common ground was perceived as sufficient. This finding confirms earlier reported conflicts in mediated settings, where social tensions were aggravated by the limitations of the mechanisms supported (Woodruff and Aoki 2003; and Wiberg and Whittaker 2005). Looking back to the notion of ambiguity as defended by Aoki and Woodruff (2005), delays and omissions would need a grounding scheme: multiple interpretations can be offered and one preferred explanation could be agreed upon if necessary. This would support minimizing the uncertainty of one party (they would know why no response is coming), while the other party can use the delay to prevent unwanted interaction. Thus the delay becomes more socially acceptable. Our second conclusion argues that grounding the intentions of interactions (i.e. the meaning of the representations used to signal privacy borders) must be supported by the system and not left outside of it, as argued by Bellotti et al. (1993) and Palen et al. (2003). Assuming that external common ground or established social practices would suffice to resolve possible conflicts, rejects the notion of privacy as a dynamic regulation process. In addition, mechanisms that allow for multiple interpretations and a joint agreement to a preferred one, support grounding even in cases when social conflicts are present (as no truth has to be revealed as soon as a common interpretation of each other's borders is established).

6. Discussion and future work

Interpersonal privacy management has become a growing concern in the design of technologies supporting mediated communication. Users are often confronted with difficulties regarding interruptions and managing their availability for communication. A common approach to address these issues has been primarily concerned with predicting users' availability and automating the interaction control based on their privacy preferences. From a social interaction perspective, such solutions appear ill fated, as privacy should be treated as a dynamic and collaborative regulation process in which communicators open and close their interaction borders to represent and understand their mutual privacy needs.

Following the above conception of privacy, we have investigated, proposed and illustrated the Privacy Grounding Model as a way to describe and analyse the process by which communicators cooperate to convey, adjust and mutually agree on a desired level for solitude or for socializing. The model identifies signalling and grounding mechanisms as collaborative and dynamic practices in the coordination of interpersonal privacy borders. Further, we have shown that lightweight and distinctive mechanisms are needed to support the signalling and grounding of privacy borders.

Overall, this research contributes a better understanding of the dynamic and collaborative process of interpersonal privacy coordination. It provides evidence to emphasize the importance of designing mechanisms that support individuals to ground their intentions of interaction when opportune, thus minimizing social conflicts and cognitive load (Wiberg and Whittaker 2005).

While not described in this paper, PGM can also inform possible design explorations. For example, we developed new features called 'one-click' and 'drag-and-drop phrases' in the chat box of the CB application (Romero and Markopoulos 2009; Romero et al. 2009). With 'one-click', a user can just click a message to say that one has read it. With 'drag-and-drop', one can drag a predefined phrase under a selected message to more explicit express an intention. Each mechanism represented different levels of lightweight and distinctive signalling that people could use to ground existing text representations in the chat box. First insights show that participants preferred the 'one-click' mechanism as it was lighter weight (though less distinctive) than 'drag-and-drop'. An interesting collaborative behaviour that emerged from the use of the one-click was the grounding of passive participation: participants could acknowledge with one-click that they were aware about a particular conversation but without having to actively participate in it.

Our analysis suggests interactive coordination of privacy borders. Further research could explore the role of automation in supporting this process. As observed in the analysis of CB, the most recurrent reason for explicit grounding corresponded to the need to explicitly ground automated representations (such as status or video updates). In other words, it is clear that explicit grounding in systems supporting mediated communication requires exploring a range of disambiguation signals, such as ambiguities that may arise when automatically constructed status updates are used for privacy border regulation. Belloti et al. (1993) offered a list of requirements that could be used to check whether grounding mechanisms support things such as: appropriate feedback timing (as when control is most likely to be required and effective), perceptibility, unobtrusiveness, flexibility, meaningfulness, and learnability (solutions should be sensible to existing psychological and social mechanisms). Palen et al. (2003) defined genres of disclosure to describe the expectation and response around socio-technology arrangements that are involved in everyday privacy management. Genres of disclosure could be used to check whether grounding representations address tensions in disclosing public and private information, identity of the self and other, and temporality (past and future).

It is important to note that, regardless the level of distinctiveness, contributions were considered in this analysis as collaborative practices as soon as both communicators reached the same understanding. As reported earlier, delays were sometimes considered collaborative practices, but at other times were considered ambiguous and thus not collaborative. In future analysis, more attention could be dedicated to understanding the effect of grounding on ambiguity. This admits that a signal can have multiple interpretations, and assumes that a positive interpretation will be used if necessary. The main questions then becomes: how could the process of grounding not hinder ambiguity when desired? And vice versa, how can protecting ambiguity not hinder the grounding process? If we look at the main benefits of grounding, we observed that communicators collaborate to minimize uncertainty. Our findings showed no evidence of participants questioning the validity of others' privacy borders after an understanding was reached. In line with Aoki and Woodruff's (2005) results, this gives insights on the possibilities to acquire grounding by helping minimize uncertainty without compromising undesired interaction. Therefore, explicit grounding could 'preserve' ambiguity by supporting the possibility to associate fake or multiple meanings to contributions of privacy borders. Supporting flexible and lightweight grounding mechanisms decreases uncertainty and increases collaboration and interactive ways to achieve efficient and effective coordination of interpersonal privacy.

To summarize, lightweight and distinctive characteristics of signalling and grounding describe the dynamic and dialectic strategies that people develop when establishing shared common ground to regulate interpersonal privacy. Communicators at some moments have the need for more lightweight mechanisms, while at other moments a more distinctive mechanism is necessary.

Acknowledgment

We thank our experimental participants for their participation in our studies.

Background The reported research is based on Romero's doctoral thesis done at TU Eindhoven.

Support The GroupLab research group (Computer Science Department, University of Calgary) and the User Centered Design research group (Industrial Design Department, Eindhoven University of Technology) for providing all the facilities needed to conduct this research

Notes

1. We recognize that this separation may be somewhat artificial. The pilot study showed that participants signaled coordination contributions (track-II) using content channels (e.g., a text message 'I'm back'), most likely to overcome the lack of appropriate track-II channels for coordination purposes.
2. Diary entries were not considered as interaction signals for the coding analysis, therefore they are not present in Table 1. The response rate of the prompts was 15 %, with a highest number of entries in the first week: 29.3 average per day.
3. The extracts presented here were adapted to improve the legibility: a header row was added, participants' name were encrypted as 'P X' were X was a number from 1 to 15 to identify each participant, which sometimes was followed by an '@' and a text to indicate the location of the telecommuters, and a black bar on the video snapshots was added to protect the participants' identity when necessary (note that a blurred image was an effect intended by the participants, not a manipulation of the experimenter). For later analysis the coding results of the signaling characterizations are represented in the last four columns: "Br", "Ba", "Si", and "Di" correspond to brevity, peripheralness (background), simultaneity and distinctiveness respectively, using the codes '+' to indicate that the characterization was present and '−' to indicate the contrary.
4. The distance to the community was defined considering their work/social and physical distance from the core group. A person could be physically further away but socially close.

References

Altman, I. (1975). *The Environment and Social Behaviour—Privacy, personal space, territory, crowding*. Monterey, Calif.: Brooks/Cole Pub. Co., Inc.

Aoki, P. M. and Woodruff, A. (2005). Making space for stories: ambiguity in the design of personal communication systems. In *CHI'05. Proceedings of the Conference on Human Factors in Computing Systems, Portland, Oregon, USA, 2-7 April 2005*. New York: ACM Press, pp. 181–190.

Begole, J. B., Matsakis, N.E, and Tang, J. (2004). Lilsys: Sensing Unavailability. In *CSCW'04. Proceedings of the Conference on Computer Supported Cooperative Work, Chicago, Illinois, 6-10 November 2004*. New York: ACM Press, pp. 511–514.

Bellotti, V. and Sellen, A. (1993). Design for privacy in ubiquitous computing environments. In *ECSCW'93. Proceedings of the European Conference on Computer-Supported Cooperative Work, Milano, Italy, 13-17 September 1993*. Dordrecht: Kluwer Academic Publishers, pp. 77–92.

Boyle, M., and Greenberg, S. (2005). The Language of Privacy: Learning from Video Media Space Analysis and Design. *ACM Transactions on Computer–Human Interaction*, vol. 12, no. 2, June 2005, pp. 328–370.

Boyle, M., Neustaedter, C. and Greenberg, S. (2009) Privacy Factors in Video-based Media Spaces. In S. Harrison (ed.): *Media Space: 20+ Years of Mediated Life*. London: Springer, pp. 97–122.

Clark, H. (1996). *Using language*. New York: Cambridge University Press

Clark, H. and Brennan, S. (1991) Grounding in Communication. In L.B. Resnick, J.M. Levine, and S.D. Teasley (eds): *Perspectives on Socially Shared Cognition*. Washington: APA Books, pp. 127–149.

Dourish, P. (1993). Culture and Control in a Media Space. In *ECSCW'93. Proceedings of the European Conference on Computer-Supported Cooperative Work, Milano, Italy, 13–17 September 1993*. Dordrecht: Kluwer Academic Publishers, pp. 125–137.

Erickson, T., Smith, D. N., Kellogg, W. A., Laff, M., Richards, J. T., and Bradner, E. (1999). Socially translucent systems: social proxies, persistent conversation, and the design of "Babble". In *CHI'99. Proceedings of the Conference on Human Factors in Computing Systems, Pittsburgh, Pennsylvania, USA, 15–20 May 1999*. New York: ACM Press, pp. 72–79.

Greenberg, S., and Rounding, M. (2001). The notification collage: Posting information to public and personal displays. In *CHI'01. Proceedings of the Conference on Human Factors in Computing Systems, Seattle, Washington, USA, 31 March – 5 April 2001*. ACM Press, pp. 514–521.

Lederer, S., Hong, J., Dey, A., and Landay, J. (2004). Personal privacy through understanding and action: five pitfalls for designers. *Journal of Personal and Ubiquitous Computing*, vol. 8, no.6, November 2004, pp. 440–454.

Lewis, K., Kaufman, J., and Christakis, N. (2008). The taste of privacy: An analysis of colleague student privacy settings in an online social network. *Journal of Computer-Mediated Communication,* vol. 4, no. 1, October 2008, pp. 79–100.

Mackay, W. (1999). Triggers and barriers to customizing software. In CHI'99. *Proceedings of the Conference on Human Factors in Computing Systems, Pittsburgh, Pennsylvania, 15–20 May 1999*. New York: ACM Press, 1999, pp. 153–160.

Markopoulos, P., Romero, N., van Baren, J., IJsselsteijn, W., de Ruyter, B., and Farschian, B. (2004). Keeping in touch with the family: Home and away with the ASTRA awareness system. In *CHI'04. Proceedings of the Conference on Human Factors in Computing Systems Vienna, Austria, 24–29 April 2004*. New York: ACM Press, pp. 1351–1354.

McEwan, G. and Greenberg, S. (2005). Supporting social worlds with the community bar. In *GROUP'05. Proceedings of the International Conference on Supporting Group Work, Sanibel Island, Florida, USA, 6–9 November 2005*. New York: ACM Press, pp. 21–30.

McEwan, G., Greenberg, S., Rounding, M., and Boyle, M. (2006). Groupware plug-ins: A case study of extending collaboration functionality through media items. In *CollabTech'06. Proceedings of the International Conference of Collaboration Technologies, Tsukuba, Japan, 13-14 July 2006*. Tokyo: Information Processing Society of Japan, pp. 42–47.

McFarlane, D. C. (2002). Comparison of Four Primary Methods for Coordinating the Interruption of People in Human–Computer Interaction. *Human Computer Interaction*, vol. 17, no. 1, March 2002, pp. 63–139

Metaxas, G., Metin, B., Jutta Schneider, J., Markopoulos, P., and de Ruyter, B. (2007). Daily Activities Diarist: Supporting Aging in Place with Semantically Enriched Narratives. In *INTERACT'07. Proceedings of the International Conference on Human-Computer Interaction, Rio de Janeiro, Brazil, 10–14 September 2007*. Berlin: Springer, pp. 390–403.

Miles, M. and Huberman, M.(1994). *Qualitative data analysis: an expanded sourcebook.* London: Sage Publications.

Miller, T. and Stasko, J. (2001). The InfoCanvas: information conveyance through personalized, expressive art. In *CHI'01. Proceedings of the Conference on Human Factors in Computing Systems, Seattle, Washington, USA, 31 March – 5 April 2001*. ACM Press, pp. 305–306.

Monk, A. (2003). Common ground in electronically mediated communication: Clark's theory of language use. In J.M. Carroll (ed.): *HCI Models, Theories, and Frameworks*. San Francisco: Morgan Kaufmann Publishers, pp. 265–290.

Nardi, B., Whittaker, S., and Bradner, E. (2000). Interaction and outeraction: instant messaging in action. In *CSCW'00. Proceedings of the Conference on Computer Supported Cooperative Work, Philadelphia, Pennsylvania, 2–6 December 2000*. New York: ACM Press, pp. 79–88.

Palen, L. (1999). Social, individual and technological issues for groupware calendar systems. In *CHI'99. Proceedings of the Conference on Human Factors in Computing Systems, Pittsburgh, Pennsylvania, USA, 15–20 May 1999*. New York: ACM Press, pp. 17–24.

Palen, L., and Dourish, P. (2003). Unpacking "privacy" for a networked world. In *CHI'03. Proceedings of the Conference on Human Factors in Computing Systems, Fort Lauderdale, Florida, USA, 5–10 April 2003*. New York: ACM Press, pp. 129–136.

Patil, S., and Kobsa, A. (2004). Instant messaging and privacy. In *HCI'04. Proceedings of Human Computer Interaction, Leeds, UK, 6–10 September 2004*. Berlin: Springer, pp. 85–88.

Patil, S., and Lai, J. (2005). Who gets to know what when: configuring privacy permissions in an awareness application. In *CHI'05. Proceedings of the Conference on Human Factors in Computing Systems, Portland, Oregon, 2–7 April 2005*. New York: ACM Press, pp. 101–110

Petronio, S. (2002). *Boundaries of privacy: dialectics of disclosure.* Albany, New York: State University of New York Press.

Reichenbach, M., Damker, H., Federrath, H., and Rannenberg, K. (1997). Individual management of personal reachability in mobile communication. In *SEC'97. Proceedings of the Conference on Information Security in Research and Business, Copenhagen, Denmark, 14–16 May 1997*. UK: Chapman and Hall, pp. 164–174.

Romero, N. (2008). *Coordination of Interpersonal Privacy in Mediated Communication*. Ph.D. dissertation. Eindhoven University of Technology, Eindhoven: Industrial Design Department.

Romero, N., and Markopoulos, P. (2005). Common Ground to Analyse Privacy Coordination in Awareness Systems. In *INTERACT'05. Proceedings of the International Conference on Human-Computer Interaction, Rome, Italy, 12–16 September 2005*. Berlin: Springer, pp. 1006–1009.

Romero, N. A. and Markopoulos, P. (2009). Grounding interpersonal privacy in mediated settings. In *GROUP'09. Proceedings of the International Conference on Supporting Group Work, Sanibel Island, Florida, USA, 10–13 May 2009*. New York: ACM Press, pp. 263–272.

Romero, N., Markopoulos, P., van Baren, J., de Ruyter, B., Ijsselsteijn, W. and Farshchian, B. (2007a). Connecting the family with awareness systems. *Personal and Ubiquitous Computing* [Special issue on "Memory and Sharing of Experiences"], vol. 11, no. 4, April 2007, pp. 299–312.

Romero, N., Matysiak, A., Kaptein, M., and Markopoulos, P. (2007b). Behaviours and preferences when coordinating mediated interruptions: Social and system influence. In *ECSCW'07. Proceedings of the European Conference on Computer-Supported Cooperative Work, Limerick, Ireland, 24–28 September 2007*. London: Springer, pp. 351–370.

Romero, N., Boer, L., and Markopoulos, P. (2009). Interactive and Lightweight Mechanisms to Coordinate Interpersonal Privacy in Mediated Communication. In *INTERACT'09. Proceedings of the International Conference on Human–Computer interaction, Uppsala, Sweden, 24–28 August 2009*. Berlin: Springer, pp. 832–833.

Roussopoulos, M., Maniatis, P., Swierk, E., Lai, K., Appenzeller, G., and Baker, M. (1999). Person-level routing in the mobile people architecture. In *USITS'99. Proceedings of the second USENIX Symposium on Internet Technologies and Systems, Boulder, Colorado, 11–14 October 1999*. Berkeley: USENIX Association, pp. 15–15.

Garde-Perik, E.M. van de, Markopoulos, P., de Ruyter, B, Eggen, B., IJsselsteijn, W.A. (2008). Investigating privacy attitudes and behavior in relation to personalization. *Social Science Computer Review*, vol. 26, no. 1, February 2008, pp. 20–43.

Wiberg, M. and Whittaker, S. (2005). Managing availability: supporting lightweight negotiations to handle interruption*s. ACM Transactions on Computer Human Interaction,* vol.12, no.4, December 2005, pp. 356–387.

Woodruff, A. and Aoki, P. M. (2003). How push-to-talk makes talk less pushy. In *GROUP'03. Proceedings of the International Conference on Supporting Group Work, Sanibel Island, Florida, 9–12 November 2003*. New York: ACM Press, pp. 170–179.

Yin, R. (1994). *Case Study Research, Design and Methods*. Newbury Park, Calif.:Sage.