Privacy Implications of Presence Sharing in Mobile Messaging Applications

Andreas Buchenscheit,[1,3] Bastian Könings,[2] Andreas Neubert,[3]
Florian Schaub,[4] Matthias Schneider,[3] Frank Kargl[2]

[1]Ulm University of Applied Sciences
Ulm, Germany
buchenscheit@hs-ulm.de

[2]University of Ulm
Ulm, Germany
bastian.koenings, frank.kargl@uni-ulm.de

[3]Cortex Media GmbH
Ulm, Germany
a.neubert, m.schneider@cortex-media.de

[4]Carnegie Mellon University
Pittsburgh, PA, USA
fschaub@cmu.edu

ABSTRACT

Mobile messaging applications, such as WhatsApp, provide a free alternative for mobile texting on smartphones. Mobile messengers typically also share presence information about users to indicate when a user is online. We investigated the privacy implications of such presence updates, using WhatsApp as an example. We conducted a user study with two independent groups (19 participants in total), in which we collected and analyzed their presence information over four weeks of regular WhatsApp use and conducted follow-up interviews. Our results show that presence information alone is sufficient to accurately identify, for example, daily routines, deviations, times of inappropriate mobile messaging, or conversation partners. We discuss resulting privacy implications of presence information and potential solutions to mitigate these issues.

1. INTRODUCTION

Mobile messaging applications like WhatsApp[1] or Line[2] have emerged as mostly free alternatives to conventional SMS messaging. Besides text messaging they also support the exchange of images, videos, or voice records. WhatsApp is one of the most popular messaging applications with more than 500 million users [15] and more than 20 billion messages sent per day in 2014 [13].

Mobile messaging apps commonly support sharing of presence information. Presence information serves to convey a user's online status to others (e.g., being *online*, *busy*, *offline*), whether the user is typing, or if a particular message has been read by the recipient. Thus, sharing presence information aims to enhance social interactions between sender and recipient, e.g., for deciding when to send a message. While conventional desktop messaging applications, such as Skype, Jabber, or Facebook messenger, typically rely on users manually selecting their presence status, manually setting the presence status is uncommon in mobile messaging applications. Presence information is instead inferred from interaction with the messaging app. For example, WhatsApp automatically sets the presence status to "online" when the WhatsApp app is in the foreground, and to "offline" when it is put in the background. The "last seen" feature further reports when a user was last online [20]. WhatsApp's different presence messages are shown in Figure 1. A user's presence information can be seen by any other WhatsApp user as long as the phone number is known and was added to the phone's address book [20]. While some WhatsApp applications allow users to deactivate the "last seen" feature, the automatic transmission of the presence status (i.e., online or offline) can not be deactivated.

Figure 1: WhatsApp presence information showing online status a) and last seen timestamp b).

In this paper, we investigate the privacy implications of presence sharing in such mobile messaging applications. Automated interaction-triggered status updates combined with the fact that mobile messengers are often used throughout the whole day make the resulting presence information potentially much more privacy sensitive than similar presence information in desktop messaging applications. In order to investigate the assumed privacy implications of presence sharing information from mobile devices we conducted a user study in which we collected and analyzed presence information of two independent groups of WhatsApp users (19 participants in total) over a four week period of regular WhatsApp use, resulting in over 27,000 presence updates and reflecting 545 hours of WhatsApp use in total. In our analysis, we considered privacy issues that emerge between users who may have only limited access to each other's presence information, as well as more potent adversaries that

[1]http://www.whatsapp.com/
[2]http://line.me/en/

Accepted for the International Conference on Mobile and Ubiquitous Multimedia (MUM '14)
Author version. Definitive version at http://dl.acm.org

are able to monitor groups of users or the whole network, such as service providers or intelligence agencies. Our results show that presence information of mobile messaging apps is sufficient to infer behavior patterns (e.g., bedtimes, work hours or school hours) with high accuracy. Furthermore, specific conversation patterns can be identified which facilitate detection of communication partners and duration of conversations. We further conducted interviews with our participants to verify their behavior in the study period and understand their perceptions of the uncovered privacy implications. While we use WhatsApp as an example, identified privacy implications also pertain to other mobile messaging applications that automatically share presence information.

We first discuss related work in Section 2. In Section 3 we describe our methodology for collecting real user's presence information, before presenting our analysis of associated privacy implications in Section 4. Based on follow-up interviews, we report how our participants perceived identified privacy implications in Section 5. The paper concludes with a discussion of the implications of our results in Section 6.

2. RELATED WORK

Our work is related to general studies of messaging application use, as well as experiments and studies investigating the inference of behavior patterns from user activity, and resulting privacy implications.

2.1 Usage of Messaging Applications

Avrahami and Hudson [3] studied instant messages of 16 participants and created statistical models to predict whether a message would receive a response within a specific time window or not. Avrahami et al. [2] further found that responsiveness is affected by multiple factors, such as how the message is presented. They also found that the time to respond not only affects the dynamics of a messaging conversation but also the participants' perceptions of each other.

A number of studies more specifically investigated smartphone usage. Do et al. [10] proposed a probabilistic framework for extracting usage patterns from the interaction with smartphone apps. They find that such patterns correspond to the users' interests but also reveal how they utilize their phone. In a nine-month longitudinal study, they further identified dependencies between smartphone usage and the user's location and social context [9]. Böhmer et al. [7] find that smartphone interactions often last less than one minute, on average. Certain apps are more popular at specific times of the day, e.g., news apps in the morning. However, communication apps, including messaging, dominate throughout the day and are almost always the first apps used after the device was in sleep mode. Shirazi et al. [19] investigated how users react to smartphone notifications, such as an incoming message, by tracking click interactions. In their sample population, WhatsApp generated by far the largest number and most frequent notifications, followed by emails and SMS. WhatsApp also exhibited one of the shortest delays between notification and user click (15 seconds median delay).

Concerning privacy and mobile messaging apps, Schrittwieser et al. [18] describe privacy issues related to friend finder features that upload a user's complete address book. Such features violate the privacy of primary users as well as their contacts, but also facilitate probing queries to determine whether someone is registered with a service. Schrittwieser et al. generated a phone book with ten million phone numbers that they used to verify 20,000 phone numbers of registered WhatsApp users. In a survey with 131 WhatsApp users, Church and Oliveira [8] found that the majority of users was more concerned about privacy when using WhatsApp compared to SMS. A major privacy concern was the "last seen" feature.

These findings led to our hypothesis that presence information from mobile messengers, which is sent out whenever the app is opened or closed, may reveal a user's activities and behavior patterns, because mobile messaging apps are used almost continuously throughout the day [7, 19].

2.2 Inferring Behavior from Activity

Inferring behavior patterns from user interactions has been studied in different contexts. Begole et al. [5, 4] analyzed mouse and keyboard activity in relation to the users' location, calendar information, and email activity. Based on data visualizations, they identified recurring patterns and noted variations dependent on location, time, and day. They conclude that online presence and awareness information involve privacy issues and provide suggestions for addressing them [5]. Raij et al. [17] conducted a similar study focused on wearable sensors. Their participants were most concerned about the potential of inferring information about their conversations, commutes, and psychological states. Bell et al. [6] argue that ethical implications require more serious consideration as such information not only allows to infer information about the user but is also subject to interpretation and potential misinterpretation. They conducted a study in which they showed participants visualizations of their data and elicited their comfort of sharing such information with different groups of recipients. Their participants showed less concern than originally expected, but the authors caution that participants may not have had the time to ponder implications and also volunteered to be logged.

In our study design, we placed great care on unbiased data collection while fully adhering to ethical research conduct. Beyond analysis of activity patterns, we conducted a deeper investigation of the actual privacy implications by interviewing participants on perceived risks associated with inferred behavior information.

3. METHODOLOGY

Our goal in this work was to investigate privacy implications stemming from the combination of mobile messaging apps automatically sharing presence updates and the frequent use of such apps throughout the day. We based our investigations on WhatsApp, as we were able to collect presence information of WhatsApp users by exploiting a design flaw in the WhatsApp protocol, as discussed below. This enabled us to gather presence information from our participants that reflect realistic user behavior and interactions.

We recruited two independent groups of 10 and 9 WhatsApp users. After we obtained their initial consent, we collected their presence data over a period of four weeks. After concluding data-collection, we conducted semi-structured interviews with all participants to obtain demographic and back-

ground information, as well as ground truth data for aspects of their behavior, such as sleeping hours or typical activities, as well as behavior at specific occasions identified in the collected data.

We took particular care to prevent priming participants about privacy in order to avoid influencing their WhatsApp usage behavior. We opted for a deception study design in which participants provided their phone number and consent to *"collect usage information everybody else could collect as well"* without being aware of what specific data we collect. We invited participants for semi-structured interviews, in which we showed them a visualization of their own presence information in relation to the anonymized presence information of the other participants in their group. As part of the interview, each participant received a full debriefing detailing what data was collected, what we learned from it, and for what purpose the data would be used (anonymized data analysis and use in research publications). After the debriefing participants could freely choose to explicitly affirm or withdraw their consent. All participants consented to the use of their data in our research.

3.1 Collection of Presence Information

WhatsApp presence information was collected for participants based on their phone number. While any WhatsApp user can check someone's online status by opening a chat conversation with that person (see Fig. 1), simultaneously observing several users would require switching between multiple conversation screens. Moreover, a manual observation of a larger group over an extended period of time is virtually impossible. In fact, we were initially working on WhatsApp automation and the WhatsApp protocol when we discovered how to automatically collect presence information of WhatsApp users and wondered about the related privacy implications.

WhatsApp uses a customized version of XMPP [22], an open XML-based communication protocol for message-oriented middleware, near real-time instant messaging, sharing of presence information, and contact list maintenance. WhatsApp provides the same features plus the ability to upload multimedia data such as images, audio or video content.

The following description of the WhatsApp protocol was gained using a man-in-the-middle proxy similar to Schrittwieser et al. [18]. We intercepted SSL connections between an iPhone 5 and the WhatsApp server by routing all traffic through our proxy. Three basic steps are required to automatically collect a WhatsApp user's presence information: *One-time registration*, *user login*, and *subscription to presence information*.

3.1.1 One-time WhatsApp Registration

WhatsApp accounts are identified by a unique username which is based on users' international phone numbers (e.g., 491511111110@s.whatsapp.net). The following steps are necessary to create a new user account:

1. The phone number and a SHA1-hashed device identity have to be POSTed to an HTTPS URL to request a six digit authentication code.
2. The authentication code is received via SMS or automated voice call.
3. The authentication code, device identity and phone number have to be combined into a HTTPS URL to obtain a server-generated password.

For our tests we registered a landline phone number with a randomly SHA1-hashed device id. After receiving the automatic phone call and password, we could perform the client authentication handshake to execute a valid user login.

3.1.2 User login

The WhatsApp authentication handshake uses a custom SASL [16] mechanism called WAUTH-1. After a "hello message" follows the customized XMPP protocol using a challenge response authentication mechanism for the user login initiated by the client. The server responds with the CHALLENGE_DATA used by the client to generate a PBKDF2 key with the private password obtained during the one-time registration and SHA1 as the hash function.

After the handshake, the client is authenticated to the server and each subsequent message is encrypted. The RC4 key consists of the first 20 bytes of the PBKDF2 result combined with a hash over the concatenated CHALLENGE_DATA, RESPONSE_DATA, and the current timestamp.

3.1.3 Subscription to Presence Information

The WhatsApp protocol supports the exchange of two different types of presence information. The *"online status"* which refers to the current availability of a user, and *"last seen"* which refers to the last time a user was online and actively using WhatsApp. Both presence information types can be requested by any registered user for any other registered user whose phone number is known. Because there are no bidirectional contact relationships on WhatsApp, neither a prior contact request nor an entry in the target's contact list are required to subscribe to presence information.

After subscribing to the *"online status"*, one receives an *available* message if the user comes online, i.e., opens the WhatsApp app on her smartphone, and an *unavailable* message when the app is closed again. In our tests, these subscriptions never expired or timed out except due to WhatsApp system maintenance or failure. Thus, a continuous data collection over a long period is potentially feasible.

Furthermore, it is also possible to request the *"last seen"* timestamp of a user directly, as long as that user has not explicitly disabled this feature. A request for the last seen timestamp of a user who has disabled this option, results in an error code 405 from the server.

Many different unofficial APIs for different programming languages have emerged that can handle registration and presence subscriptions for developers which make it relatively easy to obtain and leverage presence information of users. Examples of such APIs are WhatsApi[3] for PHP or WhatsPoke[4] for Python.

[3] `https://github.com/venomous0x/WhatsAPI`
[4] `https://github.com/koenk/whatspoke`

In summary, it is essentially possible to continuously collect the complete WhatsApp presence information of any given phone number. While the transmission of the "last seen" timestamp can be explicitly disabled by users, the transmission of the online status can not be disabled and thus subscribing to this kind of presence information does always allow a continuous monitoring.

We made use of this to record presence information of our participants. Each available and unavailable event was saved with a timestamp. Collected presence information was instantly anonymized by replacing phone numbers with pseudonyms. The mapping of phone numbers to pseudonyms was stored separately and was only used during exit interviews in order to provide participants an individual visualization of their presence information and to collect ground truth data for the validation of our inferred information. Analysis of the data from our two participant groups are discussed in Section 4. The analysis of two independent groups served to corroborate the external validity of identified privacy issues and implications.

3.2 Semi-structured Interviews

Data collection was followed by semi-structured interviews. Participants first provided demographic information, as well as information about their habits concerning mobile communication and WhatsApp usage. Further questions served to validate privacy implications without mentioning privacy explicitly. Then participants were shown a visualization of their presence information over the collection period. Participants were asked about specific characteristics of their usage and to interpret their presence patterns. Different aspects of the individual participant's visualization were jointly explored and discussed. The interview concluded with questions concerning the participant's perception of uncovered privacy risks. The 19 conducted interviews lasted about one hour on average.

3.3 Participant Groups

Our dataset and analyses are based on two independent participant groups. The second group was recruited later on to validate results from the first group with participants from a different social background.

3.3.1 Group 1

Our first group was recruited randomly from a contact list available to the authors. Phone numbers were selected without any information about associated names, and then contacted to obtain initial consent. Presence information for this group was collected in the fall of 2013.

Group 1 consisted of 10 participants (P1 to P10), 6 females and 4 males, with diverse backgrounds (6 employed in different domains like arts, journalism and engineering; 2 students, 1 trainee, 1 completing a voluntary social year). Their age ranged from 17 to 29 years (median=22). Six participants owned an Android smartphone, four an iPhone. Some participants knew each other and maintained stronger social relationships as revealed in the exit interviews. All group members stated to use WhatsApp as their primary way of mobile communication.

3.3.2 Group 2

Our second group was recruited from the population of a mid-sized cooperative state university. Undergraduates, all of the same course (Business Administration, Management in Media and Communication), were invited to voluntarily participate in a research study by providing their phone number. Presence information of group 2 was collected in January and February 2014.

Group 2 consisted of 9 students (P11 to P19), 6 females and 3 males. Their age ranged from 20 to 28 years (median=22). Five used an Android smartphone, four an iPhone. While all participants knew each other, they stated to rarely use WhatsApp to communicate with each other, because they met everyday on campus. There also existed a WhatsApp group chat for their semester, which was rarely used to send important information concerning the whole course. All participants stated that WhatsApp was their primary mobile messaging application.

4. ANALYSIS OF PRIVACY RISKS

In total we collected 13,805 single events of WhatsApp usage for the first group and 13,777 usage events for the second group. Together for both groups these events result in a total usage time of 545 hours during the 4 weeks of presence information collection.

We visualized the collected presence data to ease the identification of potential behavior patterns and privacy implications. Figure 2 shows this visualization for a complete day of the first group. For each participant (rows) a bar on the time-line represents a period of active WhatsApp usage. Each use period is defined by a pair of corresponding available/unavailable events. While the visualization allows to see when participants were actively using WhatsApp, we were interested in further information that could be extracted from this data. Thus in addition to calculating usage statistics we investigated the feasibility of inferring daily routines (e.g., bedtimes or working hours), as well as communication partners.

4.1 WhatsApp Usage

The collected presence information allows to calculate detailed usage statistics of WhatsApp users. For each participant of both groups we calculated the average number of uses ($\#_{avg}$), usage time (d_{avg}), as well as time periods between application usage. The combined statistics for both groups are shown in Table 1. Statistics of both groups are very similar, which indicates that the observed usage behavior is likely not group-specific.

The average number of uses (i.e., how often WhatsApp was opened) varied from 14 (P7: SD=17) to 98 (P17: SD=27) times per day. The average usage time per day (i.e., how long WhatsApp was used) varied from 15 minutes (P16: SD=9) to 164 minutes (P17: SD=72). The individual average number of uses ($\#_{est}$) and the usage time per day (d_{est}) was estimated by participants in the exit interviews. Eight participants were quite accurate in estimating their daily number of uses (e.g., P5: $\#_{avg}$=19, $\#_{est}$=20; P1,P14: $\#_{avg}$=50, $\#_{est}$=50) and varied only a few minutes from the calculated average usage time (e.g., P3: d_{avg}=141, d_{est}=150; P11: d_{avg}=83, d_{est}=90). Interestingly, most participants

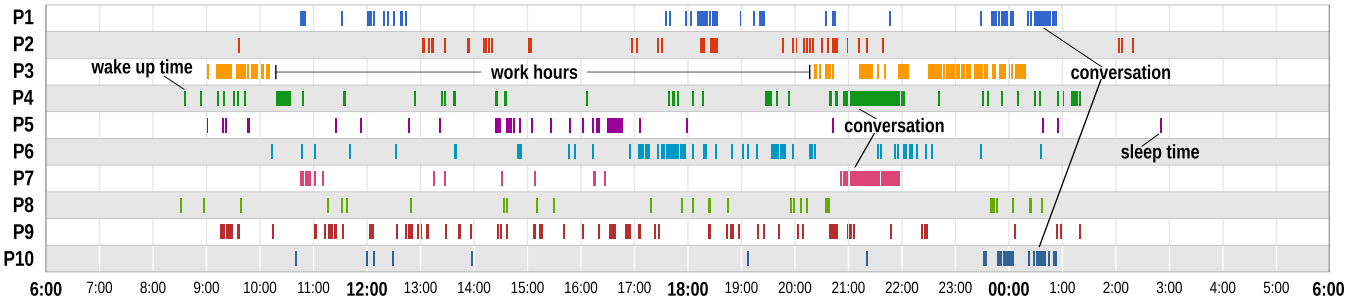

Figure 2: Presence information showing the WhatsApp activity of group 1 during one day. A bar on the timeline represents the amount of time WhatsApp was actively used. Tagged examples show that presence information can reveal wake up and sleep times, daily routines, and conversations.

Table 1: WhatsApp usage statistics for group 1 and group 2 after 4 weeks.

	group 1			group 2		
	AVG	MD	SD	AVG	MD	SD
# total uses	1381	1461	717	1531	1370	639
# uses per day	45	38	35	51	42	44
total usage [h]	28	26	14	30	23	22
usage per day [m]	54	43	49	60	44	69
time per use [s]	72	38	113	72	43	96
time betw. use [m]	32	3	195	28	5	90

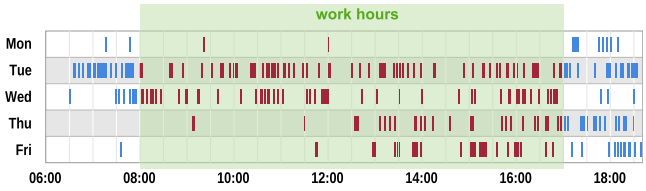

Figure 3: Presence information revealing that participant P9 used WhatsApp during work hours and thus violated policy prohibiting mobile phone use.

underestimated how often they opened WhatsApp per day on average (e.g., P3: $\#_{avg}$=70, $\#_{est}$=30; P16: $\#_{avg}$=36, $\#_{est}$=10) but overestimated their duration of use (e.g., P1: d_{avg}=57, d_{est}=240; P14: d_{avg}=48, d_{est}=90). This indicates that they are quite aware of how much time they spend with the app but not how often they look at it. High overestimates might be caused by days of less WhatsApp use that participants did not consider in their estimates. For instance, P7 (d_{avg}=14, d_{est}=150) stated that he was on holiday during the collection period and did not use WhatsApp as much as usual.

The average time per usage (i.e., the duration of a single WhatsApp session) varied from 51 seconds (P12: SD=61; P16: SD=67) to 120 seconds (P3: SD=149). The average time between two sessions varied from 13 minutes (P17: SD=49) to 104 minutes (P7: SD=289), which highlights the diverse usage frequency between participants. Note that considered pauses also include periods of sleep.

The collected presence information reveals a detailed picture of a user's WhatsApp behavior. In situations where use of messaging apps is inappropriate or even prohibited, this information could lead to several privacy implications. For instance, a superior could check whether employees are excessively using WhatsApp on workdays. Even more problematic is the particular time of use. In institutions that prohibit use of personal mobile devices during work or school hours, superiors or teachers could easily detect when policies have been violated. As this information could be collected surreptitiously by superiors, users may not be aware of this form of surveillance or how it would be used, e.g., in selecting employees for bonuses or promotion.

Ten participants (5 in each group) reported that they regularly use WhatsApp at work even though personal phone use is prohibited. Based on their presence information charts (example shown in Fig. 3), three participants (P5, P6, and P9) specifically identified their WhatsApp usage during work and school hours in violation of respective policies. Seven of them stated that their phone number is known to their superiors. Given our data collection approach, those superiors would be in the position to collect the required presence information and detect their employees' policy violation.

4.2 Daily routines

Our visualizations of presence information suggests that patterns may reveal participants' daily routines, such as bedtimes, working hours and variations thereof.

4.2.1 Inferring bedtimes

We calculated average wake up and sleep times for the complete 4 week datasets and compared those to times reported by participants in the exit interviews. Estimated bedtimes were based on the first event in the morning after a longer pause (est. wake up) and the last event in the evening or at night followed by a longer pause (est. sleep).

For both groups, we asked participants to estimate their usual bedtimes for weekdays and weekends. The combined results are listed in Table 2. The results of the first group are quite consistent for all estimated bedtimes and varied from an average difference of 47 minutes (sleep times on weekdays, SD=36) to 63 minutes (wake up times on weekdays, SD=48). For group 2 the calculated average wake up times were more accurate for weekends with an average difference of 55 minutes (SD=37) to participants' estimated average wake up times. Differences varied from -3 minutes (P17) to

Table 2: Estimated and reported bedtime statistics for group 1 and group 2.

	group 1			group 2		
	AVG	MD	SD	AVG	MD	SD
weekend: avg–est.						
diff. wake up [m]	54	44	48	55	51	37
diff. sleep [m]	62	65	45	98	83	55
weekdays: avg–est.						
diff. wake up [m]	63	58	48	105	94	47
diff. sleep [m]	47	39	36	45	39	33
final week: ET–real						
diff. wake up [m]	-	-	-	73	44	80
diff. sleep [m]	-	-	-	77	35	108

-101 minutes (P12). The calculated average sleep times were more accurate on weekdays with an average difference of 45 minutes (SD=33) to estimated sleep times. Differences varied from -3 minutes (P19) to 91 minutes (P14). Larger deviations and inaccurate results of the second group are likely caused by the fact that students had irregular class schedules and thus different wake up times on weekdays. Seven of them reported time periods of 1 to 3 hours as bedtimes instead of exact times. Likewise, eight participants of the first group reported larger time periods especially for sleep times on weekends. From all 46 bedtimes that were reported as time periods in both groups, 26 of the calculated average times were correct within those time periods (56.5%).

The second group was also asked to reconstruct their real bedtimes from the final week of data collection before the interviews took place. However, three interviews (P13, P15, and P19) had to be conducted two weeks after the collection phase due to limited availability of the participants. Those three participants had difficulty remembering their real bedtimes, and were therefore excluded from this part of the evaluation. The average difference for the seven remaining participants was 73 minutes (SD=80) for wake up times and 77 minutes (SD=108) for sleep times. Best results were achieved for P17 whose real bedtimes differed only by 17 minutes (SD=30) for wake up times and 11 minutes (SD=23) for sleep times. Large deviations were often caused by few single presence information events that skewed calculated bedtimes. For instance P11 used WhatsApp on a Sunday morning at 7am, but reported 11am as the wake up time. Thus, he apparently woke up, checked WhatsApp, and fell asleep again. However, for three of the six reported wake up times of P11 the difference was less than 30 minutes.

While the results exhibit some deviation, they largely reflect participants' stated behavior. More accurate results were achieved for participants who stated to usually check for new WhatsApp messages in the morning directly after getting up (16 participants) and regularly before going to bed (8 participants). One participant (P4) even stated that he always checks WhatsApp when waking up during nights (e.g., to use the restroom).

From our estimated bedtimes, we calculated the average awake duration of each participant and the average percentage of daily WhatsApp usage in relation to the average awake duration. Participants' usage time varies from

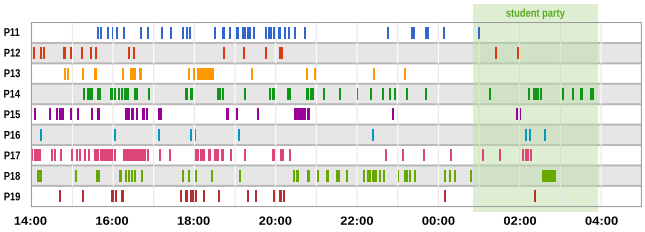

Figure 4: Presence information of group 2 showing that 8 participants where still online after 1am on a usual week day due to a student party.

2.5% (P16) to 16.8% (P17) of the respective awake duration, showing that P17 spends nearly a fifth of her day using WhatsApp.

4.2.2 *Deviations from routine*

The collected usage statistics further allow to identify deviations from daily routines, e.g., longer usage at nights or unusually long pauses during the day. Recurring variations, e.g., long pauses at particular days of the week led us to assume that participants performed weekly activities in which mobile devices could not be used, e.g., during sports or rehearsals. In the exit interviews participants were asked to explain variations in their usage patterns of the last week before the interviews.

All participants were impressed how well their activities could be depicted: *"Yes, here I was out at Saturday night [...] here you see how long I went to the gym on Tuesday and that I played Basketball on Wednesday"* (P1), *"Friday I was bar-hopping till 5 am"* (P2), *"here you can see my afternoon nap and on Saturday our work party"* (P3), *"yes, on Friday I was out till 3:46am"* (P4), and *"here I was skiing and had no network connection"* (P5). Participant P4 further stated that she was on vacation in a different country for a weekend which was clearly revealed by an irregular long pause compared to her usually intensive WhatsApp usage.

All participants of the second group confirmed their attendance of a student party on Wednesday night which was clearly visible by many presence information events of 8 participants after 1am on a usual week day, see Figure 4.

While the correct inference of activities often depends on further knowledge about a person (e.g., hobbies or profession), our results highlight the feasibility of automatically detecting daily routines and their variations, which could have several privacy implications. Knowing that someone usually does sports on Wednesdays and is not using messaging applications during this time, it is trivial to observe whether this person is exercising regularly or not (e.g., when no pause of messaging activity occurred). While the success of identifying such daily routines and variations depends on the usage frequency, our statistical results of the usage analysis show that most participants are using WhatsApp very frequently with an average time between sessions of less than 40 minutes. Only two participants (P5,P7) used WhatsApp less frequently. However, even for users with less frequent use, variations at night are likely identifiable, because they are typically caused by social events (e.g., parties, pub tours,

or cultural events) and often involve the exchange of several messages and taken pictures as confirmed by participants in our interviews.

As a further privacy implication, monitoring of bedtimes and variations thereof could be used to assess the productivity of employees or students. Someone who regularly goes to bed at 10pm is presumably more well rested and productive than someone who is partying until 4am twice a week.

4.3 Communication Partners

One of the most concerning findings was the ability to identify communication partners solely based on presence data. The feasibility of this approach depends on the employed conversation styles and the number of ongoing conversations with multiple others at the same time. According to Woodruff and Aoki [21] there are three conversation styles which we also identified in mobile messaging use:

- **Focused conversations:** Both communication partners are focused on the mobile messenger and exchange multiple messages with high frequency with each other.
- **Bursty conversations:** There are multiple short focused conversations with some gaps belonging to an overarching bursty conversation.
- **Intermittent conversations:** Messages of a conversation are irregularly exchanged with larger gaps between them.

Bursty conversations pose the highest risk of revealing conversation partners as they mark several start and end points. Focused conversations could be identified when less or no other conversations are ongoing in the observed group. Harder to identify are intermittent conversations where larger gaps occur between multiple messages of one conversation. The more additional conversations one of the communication partners is involved in at the same time, the harder it is to distinguish between them. Furthermore, the monitoring entity must already monitor potential communication partners in order to identify whether they converse with each other.

In the four week dataset of the first group, we identified 13 conversations which have been confirmed by the respective communication partners in the exit interviews. Figure 5 shows a bursty conversation between participants P1 and P10 (top) and a focused conversation between P7 and P10 (bottom). For the second group we could not identify any obvious conversations which could be caused by the fact that participants of that group exhibited weaker social ties.

Most participants (17) stated to prefer a bursty conversation style. However, P8 preferred a focused style and P18 an intermittent style during mobile messaging conversations. Five participants stated that they prefer focused conversations only with close friends or when fixing important appointments. This highlights the feasibility of determining conversation partners and the real risk associated with sharing presence information.

In addition to identifying likely conversations from the visualization of presence information, we also employed methods from social network analysis [12] to determine potential communication partners. We calculated conversation probabilities between all participants of each group for the complete data collection period. The probability is the ratio (in %) of the summarized duration of overlapping presence information events to the total duration of the specified time period. Probabilities were calculated for one hour time periods. Figure 6 shows the conversation probabilities of P10 as an example. Peaks indicate hours with high conversation probabilities between two participants. Seven of the twelve depicted probable conversations were confirmed in the interviews. While overlap in app use is a relatively naive proxy for a conversation, this approach is already sufficient to automatically identify bursty and focused conversations.

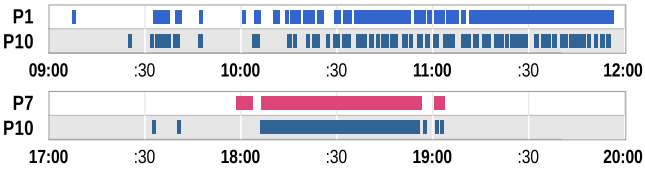

Figure 5: A bursty conversation between P1 and P10 of 144 minutes (top) and a focused conversation of 64 minutes between P7 and P10 (bottom).

Social, economical, or political implications of revealing conversation partners can be manifold. Someone might be interested in knowing whether two persons are communicating with each other for many different reasons. For example, a jealous person might want to assure that their partner is loyal and may observe their communication patterns with former partners or suspected love interests, of whom the phone numbers are known. Another person might want to know whether her best friend is chatting more often with someone else. Companies or financial institutions could be interested in collaborations between competing entities in order to gain market advantages. On a political level it may be even more critical when communication partners are revealed. Imagine an oppressive regime that can track whether someone is in contact with opposition groups or foreign agents. Also whistleblowers and their contact persons could be identified by matching their conversation patterns.

5. PERCEPTION OF PRIVACY RISKS

In addition to the analysis of the collected presence data, we also assessed the reactions and perceptions of our participants concerning what can be potentially inferred from their data. As outlined in Section 3.2, our exit interviews concluded with a series of respective questions concerning their perceptions after being confronted with visualizations of their presence data and our analysis, as well as the impact of this on their opinion on WhatsApp and similar mobile messaging applications.

5.1 (Ab)use of the Presence Feature

Asked about their awareness of the online status and "last seen" information, all participants were acquainted with these features and claimed to know how they work, i.e., that when someone is shown as online, that person is currently using the app. For instance, P1 stated a regularly usage of the "online" information: *"Sure, if you need an urgent answer*

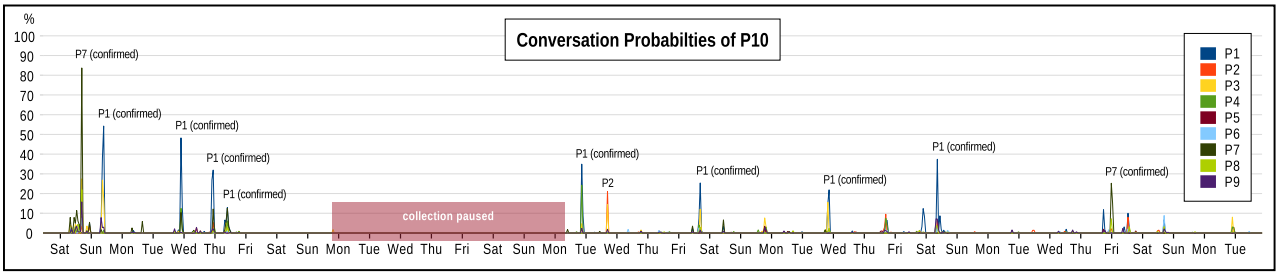

Figure 6: Conversation probabilities for P10 with other participants of group 1 during the complete period of data collection. Probabilities were calculated for one hour time windows. Ten of the eleven depicted probable conversations were confirmed in the exit interviews.

from someone who is not replying but you can see that he is online in WhatsApp." When asked further, whether they had activated or deactivated the "last seen" feature, 17 of the 19 participants had the feature turned on. One of the two participants who turned the feature off reported that he turned it off specifically to not get in trouble with his girlfriend anymore (*"Yes deactivated, because my girlfriend constantly confronted me about certain events like why I was online late last night."*, P10).

Six participants were also not aware that WhatsApp's "last seen" feature could be deactivated (*"Yes, the feature is activated. I didn't know that it could be deactivated."*, P14) while 4 of the 10 Android users stated that they would like to turn it off, if it would be possible in their version of the app[5] (*"Yes I know the feature. I have an Android phone and would like to deactivate it, but I don't know how."*, P18). This indicates that the default activation of the "last seen" feature results in a large number of users sharing their presence information without being aware of the possibility of an opt-out.

Interestingly, all four participants who stated that they would want to deactivate the "last seen at" information for themselves, actually used this feature to monitor other users. Thus, they had to keep the feature activated, because the "last seen" status of others cannot be seen anymore otherwise. The participants used the "last seen" feature to check the last time a conversation partner was online while they were awaiting a response for a prior message (16 participants who used the feature) or to check when someone went to bed or got up in the morning (6 participants). 8 of our participants also specifically checked how long others stayed at a party or event (*"If I see that X was online at 3am, I think: 'Wow! what's up with that?'"*, P6). P4 also stated that he used the online status and "last seen" feature to determine whether a friend is at home or not: *"A friend of mine only has internet access at her house, so I can determine if she is at home."*

This behavior implies that participants trade off a part of their own privacy by not deactivating the feature even if they would like to, because they want to monitor other people and get information about their behavior. Furthermore, this shows that WhatsApp users are already abusing presence information to determine similar aspects as discussed in Section 4—albeit on a smaller scale than enabled by an automated data collection method.

[5]At the time of our study the Android version of WhatsApp did not yet allow to deactivate the "last seen" feature. The transmission of the current online status, however, is not influenced by this option and thus can always be monitored.

5.2 Perceived Privacy Risks

When confronted with the visualization of their presence data and the results from our data analysis, most participants (13) were surprised or even shocked (6) by the ability to automatically gather this data (*"How is it possible for you to get this data although I have deactivated the feature?", P10*). This indicates that participants did not fully realize the privacy implications of their online status before, although all of them had a good understanding what it communicates and even used it to monitor people themselves. One of the six participants who was not surprised by the mere fact that this data is available, was surprised that his presence data was easily accessible to external parties. P6 stated *"I didn't think it was possible for you, I thought only WhatsApp could do that."*

While all but two participants would not use a tool to automatically monitor others, four participants exclaimed that they could imagine employers using it for surveillance (*"only interesting for employers"*, P18). A commonly stated reason against using automated monitoring (given by 3 of the 4 participants who gave a reason why they would not do it), was that they did not want to harm the privacy of others (*"No, that's too personal and does not concern me.", P18*). Thus, the perception of a privacy breach of a person changes with the scale of monitoring for most of the participants. An automated monitoring tool is seen as more intrusive than the occasional manual monitoring of someone's online or last seen status that nearly all participants admitted to.

However, voiced surprise about what can be inferred from presence information and WhatsApp usage did not translate into actual concern for their privacy. When asked whether any officials or superiors (e.g., teachers, superiors, or professors) knew their phone number, 14 participants confirmed that they frequently shared their number with persons who could potentially have an interest in monitoring them. Despite their confrontation with our results, 7 participants also did not see imminent privacy risks for themselves, even if someone would monitor their presence status *("No, doesn't concern me at all.", P17)*. These misconceptions of the privacy risks pertaining to themselves can to some extent be attributed to the valence effect, which is a well-known self-serving cognitive bias that causes individuals to overesti-

mate the likelihood of favorable events for them compared to others [1]. While our participants were aware of potential privacy implications, they did not perceive them as applying to themselves. Consequently most of our participants stated, that they did not know any alternative messengers (17 participants) and are reluctant to change to a more secure mobile messaging service (16 participants).

6. DISCUSSION OF RESULTS

While our data collection and analysis was focused on presence information of WhatsApp, our results are also of relevance for other mobile messaging applications. Our results show that presence information of mobile messaging applications facilitates surreptitious monitoring of user behavior and activity. We identified a number of privacy implications of presence information. Most notably, we were able to show that identified issues are not of a theoretical nature, but pose practical and immediate risks for users. Privacy implications become particularly prevalent in relationships with asymmetric dependency or power relations.

In work settings, superiors could gather presence information to learn about their employees' WhatsApp usage during work hours, make inferences of their work performance based on sleep hours and exercise regimen, and also detect deviations from daily routines that may be indicative of an unorthodox life style. Multiple of our participants confirmed that they often violate work and school policies against mobile texting when using WhatsApp, while at the same time their superiors know participants' phone numbers.

In personal settings, presence information can be analyzed to determine activity times and conversation partners. Such information can impact relationships and friendships as it may lead to the inference of flirting or cheating behavior. Our exit interviews revealed that many of our participants already engaged in occasional small-scale monitoring of others to learn about their activities. Thus, we argue that presence information should be seen and treated as sensitive information, somewhat comparable to metadata of call records. As a result, mobile messaging providers should strive to adequately protect presence information.

Presence information should only be available to a user's contacts to prevent abuse by third parties. This could be achieved by requiring mutual acceptance as contacts before messages can be exchanged. However, WhatsApp and other mobile messengers leverage the user's phone number as an identifier to enable message exchange without prior exchange of contact requests. A potential solution would be to enable message exchange without prior establishment of a contact relationship but only provide access to presence updates once a bidirectional message exchange occurred.

Mobile messaging applications should provide privacy settings that allow to disable presence updates for all or specific contacts. In contrast to our WhatsApp findings, such settings should effectively stop the communication of presence information rather than only hiding them in the client apps. Currently, WhatsApp users can only mitigate the identified privacy risks by switching off all Internet connections when using WhatsApp. Messages will then be delivered when Internet connectivity is available again. This method prevents the transmission of the last seen timestamp and is also used by some third party applications[6] that automatically handle the connectivity changes in the background.

Recently, a number of mobile messaging applications have emerged that support end-to-end encryption, such as Telegram[7] or Threema.[8] Yet, encryption alone does not prevent monitoring of when a user sends messages, which may suffice to identify communication partners. However, our results suggest that conducting multiple messaging conversations in parallel makes it more difficult to extract specific communication partners. Thus, multiple parallel messaging streams could enhance privacy. This could potentially be leveraged by generating a stream of fake messages in which genuine communication can be hidden.

6.1 Limitations

The privacy implications we have identified in this work are by no means exhaustive. Further issues and implications can likely be derived from presence information of mobile messaging apps, for example, with machine learning, data mining, or social network analysis techniques. While our approach for calculating conversation probabilities based on overlapping app use allows to automatically identify bursty and focused conversations, more sophisticated approaches could potentially also identify intermittent conversations. Future work could evaluate the utility of clustering [11] or dynamic-time-warping [14] approaches to recognize similar but non-overlapping conversation patterns.

Our results are based on the study of two independent groups of 9 and 10 users, which were monitored for four weeks. We believe that the consistency in results between groups provides good indication that our results can be generalized to other groups of WhatsApp users and potentially to users of other mobile messaging applications. However, further studies in different demographics and larger user groups are required to analyze further privacy implications of such presence information.

A potential limitation of our study is the qualitative, self-reported nature of user behavior, which we used as a confirmation for identified behaviors and routines. It would be desirable to collect more objective and fine-grained ground truth data of user behavior in order to compare it against derived behavior. We are investigating the potential of collecting such ground truth data by equipping participants with activity or fitness trackers.

Another aspect to consider are the ethical implications of conducting deception studies. Despite their consent, our participants were initially not aware of what data was collect to prevent biasing their behavior. We took great care that participants were debriefed extensively as part of the exit interviews and that all their questions and concerns were addressed. All participants agreed to the use of their presence data for our research, as long as it would be anonymized.

[6] e.g., WhatsAppGhost: `https://play.google.com/store/apps/details?id=com.albertoj.whatsappghost`
[7] `https://telegram.org`
[8] `https://threema.ch/`

7. CONCLUSIONS

WhatsApp has become one of the most prevalent mobile messaging applications on smartphones. It is widely used to exchange short messages, images, or videos. However, the use of WhatsApp and similar mobile messaging apps also poses new privacy risks for users. In this paper, we investigated WhatsApp's feature of sharing presence information which allows users to see each other's online status and "last seen" timestamp. While the online status and last seen timestamp are only updated each time WhatsApp is in the foreground on the phone and actively used, it accurately reflects a user's interaction with the app. We found that this presence information can be requested by everyone for any known number. While the "last seen" feature can be deactivated by users, sharing the current online status can not be deactivated and thus allows continous monitoring.

Through the analysis of captured online status data of 19 WhatsApp users for a period of four weeks, we identified several privacy risks stemming from presence information. Presence information alone is sufficient to create detailed usage statistics (i.e., when and how long someone is using WhatsApp), to derive a user's daily routines, and even to infer communication partners with sufficient reliability. As we discussed, social or even economic and political implications of misusing such information are manifold. Currently, the only immediate mitigations to this risk are the use of third party apps to prevent WhatsApp from updating the online status, or not using WhatsApp at all. We think that providers of mobile messaging applications should rethink their system architecture in the light of our findings and provide better control mechanisms to users to enable fine-grained control of the presence sharing feature that is actually implemented on the protocol level. We also advocate for conservative default settings following the "privacy-by-default" paradigm. In the case of WhatsApp, both recommendations are apparently not implemented.

As future extensions to our work, we foresee multiple ways to enhance the detection of conversations between communication partners using more advanced statistical approaches. Furthermore, we are also planning a larger study where we want to deepen our investigations on user concerns and the feasibility of automatically interpreting status data to infer communication partners and behavior patterns.

8. REFERENCES

[1] A. Acquisti and J. Grossklags. What Can Behavioral Economics Teach Us About Privacy? In *Digital Privacy: Theory, Technologies, and Practices*, chapter 18, pages 363–377. Auerbach Pub., 2008.

[2] D. Avrahami, S. R. Fussell, and S. E. Hudson. IM waiting: Timing and responsiveness in semi-synchronous communication. In *Proc. CSCW '08*. ACM, 2008.

[3] D. Avrahami and S. E. Hudson. Responsiveness in instant messaging: Predictive models supporting inter-personal communication. In *Proc. CHI '06*. ACM, 2006.

[4] J. B. Begole, J. C. Tang, and R. Hill. Rhythm modeling, visualizations and applications. In *Proc. UIST '03*. ACM, 2003.

[5] J. B. Begole, J. C. Tang, R. B. Smith, and N. Yankelovich. Work rhythms: Analyzing visualizations of awareness histories of distributed groups. In *Proc. CSCW '02*. ACM, 2002.

[6] M. Bell, M. Chalmers, L. Fontaine, M. Higgs, A. Morrison, J. Rooksby, M. Rost, and S. Sherwood. Experiences in logging everyday app use. In *Digital Economy*, 2013.

[7] M. Böhmer, B. Hecht, J. Schöning, A. Krüger, and G. Bauer. Falling asleep with angry birds, facebook and kindle: A large scale study on mobile application usage. In *Proc. MobileHCI '11*. ACM, 2011.

[8] K. Church and R. de Oliveira. What's up with whatsapp?: Comparing mobile instant messaging behaviors with traditional SMS. In *Proc. MobileHCI '13*. ACM, 2013.

[9] T. M. T. Do, J. Blom, and D. Gatica-Perez. Smartphone usage in the wild: A large-scale analysis of applications and context. In *Proc. ICMI '11*. ACM, 2011.

[10] T. M. T. Do and D. Gatica-Perez. By their apps you shall understand them: Mining large-scale patterns of mobile phone usage. In *Proc. MUM '10*. ACM, 2010.

[11] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer, 2001.

[12] M. Hennig, U. Brandes, J. Pfeffer, and I. Mergel. *Studying Social Networks*. Campus Verlag, 2012.

[13] K. Hong. Whatsapp hits new record after handling 64 billion messages in one day. *The Next Web*, apr 2014. `http://thenextweb.com/apps/2014/04/02/whatsapp-hits-new-record-handling-64-billion-messages-one-day/`.

[14] E. J. Keogh and M. J. Pazzani. Scaling up dynamic time warping for datamining applications. In *Proc. KDD '00*. ACM, 2000.

[15] J. Koum. 500,000,000. *WhatsApp Blog*, dec 2014. `http://blog.whatsapp.com/613/500000000`, accessed on 04.07.2014.

[16] A. Melnikov and K. Zeilenga. Rfc4422: Simple authentication and security layer (SASL), 2006. `https://tools.ietf.org/html/rfc4422`.

[17] A. Raij, A. Ghosh, S. Kumar, and M. Srivastava. Privacy risks emerging from the adoption of innocuous wearable sensors in the mobile environment. In *Proc. CHI '11*. ACM, 2011.

[18] S. Schrittwieser, P. Frühwirt, P. Kieseberg, M. Leithner, M. Mulazzani, M. Huber, and E. Weippl. Guess who's texting you? evaluating the security of smartphone messaging applications. In *Proc. NDSS '12*. Internet Society, 2012.

[19] A. S. Shirazi, N. Henze, T. Dingler, M. Pielot, D. Weber, and A. Schmidt. Large-scale assessment of mobile notifications. In *Proc. CHI '14*. ACM, 2014.

[20] WhatsApp. Whatsapp legal info, 2013. `http://www.whatsapp.com/legal/`, accessed on 04.07.2014.

[21] A. Woodruff and P. M. Aoki. How push-to-talk makes talk less pushy. In *Pro. GROUP '03*. ACM, 2003.

[22] XMPP Foundation. Extensible messaging and presence protocol (xmpp) standard, 2011. `http://xmpp.org/`, accessed on 22.02.2014.