Mobile Device Usage Characteristics: The Effect of Context and Form Factor on Locked and Unlocked Usage

Daniel Hintze
FHDW University of Applied Sciences
Fürstenallee 3 - 5
33102 Paderborn, Germany
daniel.hintze@fhdw.de

Rainhard D. Findling
University of Applied Sciences Upper Austria
Softwarepark 11
4232 Hagenberg, Austria
rainhard.findling@fh-hagenberg.at

Sebastian Scholz
FHDW University of Applied Sciences
Fürstenallee 3 - 5
33102 Paderborn, Germany
sebastian.scholz@fhdw.de

René Mayrhofer
University of Applied Sciences Upper Austria
Softwarepark 11
4232 Hagenberg, Austria
rene.mayrhofer@fh-hagenberg.at

ABSTRACT

Smartphones and tablets are an indispensable part of modern communication and people spend considerable time interacting with their devices every day. While substantial research has been conducted concerning smartphone usage, little is known about how tablets are used. This paper studies mobile device usage characteristics like session length, interaction frequency, and daily usage in locked and unlocked state with respect to location context. Based on logs from 1,585 Android devices (470 years of total usage time), we derive and analyze 23 million usage sessions. We found that devices remain locked for 60% of the interactions and usage at home occurs twice as frequent as at work. With an average of 58 interactions per day, smartphones are used twice as often as tablets, while tablet sessions are 2.5 times longer, resulting in almost equal aggregated daily usage. We conclude that usage session characteristics differ considerably between tablets and smartphones.

Categories and Subject Descriptors

K.6.2 [**Installation Management**]: Performance and usage measurement; D.4.6 [**Security and Protection**]: Authentication.

General Terms

Human Factors, Measurement

Keywords

Daily interactions, Device unlocking, Locked usage, Session length, Smartphone, Tablet, Usage session, User context

1. INTRODUCTION

Personal mobile devices have become ubiquitous today and people typically spend several hours using smartphones and tablet computers each day. Analyzing characteristics of user interactions with their devices can benefit many research areas [8]. Consequently, smartphones – being the most popular mobile device form factor today – have recently been the subject of handset-based studies analyzing characteristics of usage and interaction [2, 4, 8, 12, 13]. However, little is known about how users interact with tablet devices, which are becoming a mainstream phenomenon, reporting an annual market growth rate of 68% in 2013 [3]. Smartphones and tablets offer comparable technical capabilities like connectivity, computational power, operating systems and application ecosystem. The two form factors differ predominantly in screen size. As device size has an effect on both application and mobility, understanding how tablets are used in comparison to smartphones is worthwhile. In this work, we therefore analyze mobile device usage characteristics such as session length, interaction frequency and daily usage with respect to three dimensions:

1) As the majority of interactions with mobile devices do not include unlocking the device [4], we distinguish between locked and unlocked usage.

2) Since location context (e.g., being at home or at work) is known to have a significant effect on mobile device usage [12], we consider contexts classified as *home*, *office*, *other meaningful place*, and *elsewhere.*

3) With little previous knowledge about the impact of form factor on device usage, this work is to our best knowledge the first to analyze and compare usage characteristics of both smartphones and tablets.

Our objectives are two-fold: on one hand we aim to give a high level overview of mobile device usage characteristics. On the other hand we want to provide extensive multi-layered statistical information on device usage based on the dimensions stated above. Being first to consider not only one but three dimensions, we seek to answer our main research question: *How do context, form factor, and lock status effect mobile device usage session characteristics?*

The paper is organized as follows: First, previous mobile device usage studies and their results are discussed in section 2. Next, the concepts of locked and unlocked usage sessions as well as user context are explained and defined in section 3. Section 4 outlines the underlying dataset and how usage sessions are derived, the algorithms applied to detect locations based on Wi-Fi scan results and GSM cell-IDs, and how contextual meaning is assigned to discovered locations. We introduce and discuss our findings in section 5. The final section 6 concludes the paper.

2. RELATED WORK

Different aspects of mobile device usage have previously been studied. Falaki et al. [2] reported "immense diversity" in smartphone usage when analyzing user interaction on 255 Android and Windows Mobile smartphones, finding the average number of interactions to vary from 10 to 200. A study on 17,300 BlackBerry devices was conducted by Oliver [8], observing an average interaction time per day of 101 minutes with 80% of the usage sessions taking 90 seconds or less. The high percentage of rather short interaction can be explained with the emergence of what Oulasvirta et al. [9] refer to as *checking habits*: short but repetitive access of dynamic content providing some form of informational "reward", for example checking news, emails or Facebook. Verkasalo [14] examined contextual patterns in mobile device usage of 324 smartphone users, finding device usage to be noticeably diverse in *office* and *home* context. Soikkeli et al. [12] studied the relation between mobile device usage and end user context using 140 smartphones. While not distinguishing locked and unlocked usage, they found that usage sessions are longer in *home* context while more frequent in *office* context.

To our knowledge, the first work concerned with locked and unlocked mobile device usage was conducted by Truong et al. [13]. In order to motivate *Slide to X*, an alternative lock screen utilizing microtasks, they conducted a user study to analyze how often users unlock their devices. Studying mobile device usage based on an earlier version of the dataset used in this work, Wagner et al. [16] observed that a noticeably number of interactions occur without unlocking the device. Recent work closely related to our research focused on analyzing characteristics of locked and unlocked mobile device usage [4]. While their work is also based on an earlier version of the dataset used in our study and covers many similar aspects, their study does neither cover the effect of end user context nor form factor in relation to device usage. Except for Wagner et al. [16], who did not distinguish device types, all previous studies analyzed smartphone usage. Little work has been published, however, on tablet usage patterns. A rare example is Müller et al. [7], who conducted a multi-method based exploration of tablet usage, finding tablets to be mostly used at home and often while doing secondary activities such as watching TV, eating or cooking.

3. SESSION AND CONTEXT

This work studies locked and unlocked mobile device usage sessions with respect to user contexts and device form factors. These concepts are explained and defined in this section.

3.1 Mobile Device Usage Session

Mobile device usage sessions are consecutive periods of time wherein users interact with their devices. Unlike most previous studies, we distinguish between *locked usage sessions* and *unlocked usage sessions*. Locked usage sessions represent periods of usage throughout which devices are locked by some form of keyguard: for instance PIN, password, graphical pattern, face unlock, fingerprint, or swipe-to-unlock. To protect access to mobile device data and services as well as preventing accidental interaction with the device, interaction possibilities are heavily restricted while devices are locked (i.e., in an unauthenticated state). Interactions are limited to certain activities considered uncritical, like activating device screens to check for time, battery health, network connectivity, notifications, appointments, or taking pictures.

In contrast to locked usage sessions, an unlocked usage session begins with unlocking the device's keyguard, e.g., by entering the correct password, PIN, or pattern – and last as long as the device is used continuously. Hence device access is unrestricted in unlocked sessions. We further consider the period of time between users starting to interact with their devices and devices being unlocked (starting unlocked sessions) to be the *authentication time*: the time it takes to unlock devices.

3.2 User Context

People use their mobile devices in different ways, depending on their current situation. For instance, in an office situation people might be more likely to use their smartphones to make phone calls or check for next meetings, while at home devices might be used more to surf the Internet or watch movies. Research by Soikkeli et al. [12] reflects these different usage patterns by observing that usage sessions are 37% longer in home context over office context, but happen 56% more often in office context over home context.

Deriving context from aggregated information is often difficult. Nevertheless, information on time and location can be combined in order to derive contextual place information. Based on previous research by Jiménez [5] and Soikkeli [11] we distinguish four different place-related user contexts: *a) home b) office c) other meaningful*, and *d) elsewhere*. While *home* and *office* are self-explanatory, *other meaningful* refers to places that do not have the characteristics of *home* and *office*, but still a significant amount of time is spent there. A frequently visited gym, for instance, would be considered an *other meaningful* place. Any place that is not classified as one of these three contexts is assigned the *elsewhere* context. This includes, but is not limited to, less frequent visited places like restaurants as well as transitions between other contexts.

Unlike other studies [5, 11, 12], we do not assign an *abroad* context for places outside users' home country. The reason being that Soikkeli et al. [12] found that on average users spend only 2% of their time abroad, making this context neglectable for the analysis of average usage patterns.

3.3 Device Form Factor

We assume device form factors to have a considerable impact on device usage. We therefore analyze usage sessions characteristics with respect to device form factors – namely smartphone and tablet devices. One previously used approach to distinguish form factors in device logs is based on the device's ability to place or answer phone calls [4]. However, some tablet devices are capable of performing GSM voice call (e.g., the Galaxy Tab 10.1). Hence we chose the screen diagonal as a discriminator for form factors. Devices featuring a screen size of 7″ or higher are considered to be tablets while devices with smaller screens are regarded as smartphones. We calculate the screen diagonal from screen resolution and pixel density stated in the dataset.

4. METHODOLOGY AND DATA

Our analysis is based on more than 100 billion records of Android mobile device usage, collected from over 17,000 devices around the world by the Device Analyzer project [16]. The complete dataset we were granted access to by the University of Cambridge Computer Laboratory[1] is the largest and most detailed dataset on Android device usage publicly available today. It consists of 263 different features[2], spanning a broad range from raw sensor data to application usage, recorded either periodically or event based by a stand-alone application available in the Google Play Store. The set of devices in the dataset covers at least 1,277 unique device types in 175 countries, with 4,700 devices participating for more than one month and 321 devices participating in the project for more than one year. Further details about the underlying data used in this work can be found in [15, 16].

In order to ensure quality of data used in our study, the Device Analyzer dataset is cautiously revised, and records not suitable for our study are excluded. We disregard records from old versions of the Device Analyzer application that do not include all data required in our session detection model (see section 4.1). Devices not using any keyguard (with keyguard including slide-to-unlock) are omitted, as they do not allow distinguishing between locked and unlocked state. To enable per-day statistics, only days captured entirely in the available dataset are used. Thus, we drop records of days recorded only partially, for instance because the recording application crashed, was paused or not installed for the full time. Since the usage session model is based on display state as an indicator of user presence, devices configured to keep the display on while connected to a charger were also omitted. Furthermore, we only analyze devices which provide data for at least seven days.

4.1 Usage Session Extraction

Deriving usage sessions from device logs requires a technical concept of recognizing user interaction. Two approaches have been used in previous research. Since mobile device interactions usually involve an application (even simple actions such as making a phone call require an application), one approach by Soikkeli et al. [12] defines usage sessions as the time periods certain applications are running in the foreground on the device. However, this approach is not suitable to study locked usage, as there is not necessarily an application active in the foreground during locked interaction. Mobile device interaction almost entirely relies on touchscreen interaction, either to display information or to capture user input. Because energy consumption is an inherent concern with battery powered mobile devices, displays – which are energy-intensive – are usually switched off as soon as possible after usage. This is done either manually or automatically after a short idle timeout. Hence, the more frequently used approach to derive usage sessions from device logs is to define usage sessions as time periods in which the device's screen is switched on (screen power based models) [2, 4, 8, 9, 13].

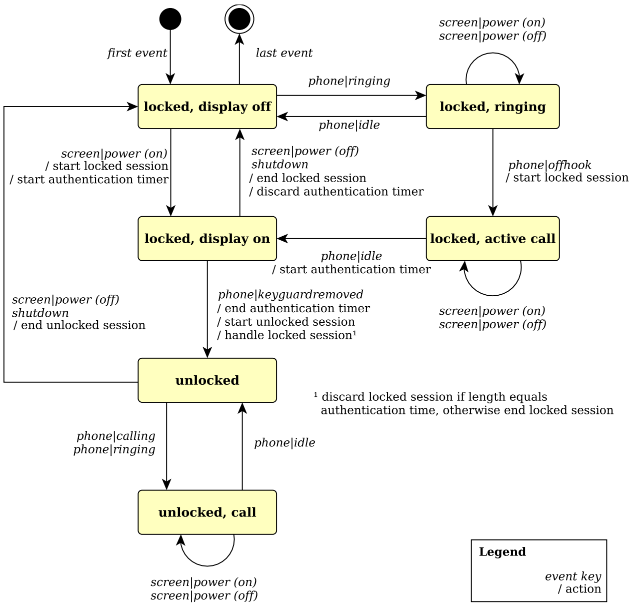

Figure 1: State machine for session detection [4]

Although naive screen power based usage session extraction comes fairly close to actual device interaction, some pitfalls exist which – in our experience – can distort the results noticeably if not considered carefully. Consider e.g., incoming phone calls, which activate the screen to display the caller's number and to allow the user to answer the call. If the call goes unanswered, a naive screen power based approach would falsely consider this a session of user interaction. Or consider phones with touchscreens, that utilize a proximity sensor to switch off the screen when the device is held closely to a user's head, e.g., during a call, in order to prevent accidental touch events caused by the user's ear. As users tend to slightly shift the phone's position during calls, this would result in naive screen power based models mistakenly recognizing multiple short usage sessions instead of one consecutive session. It has been observed that overall 12.7% of the changes in screen power state on smartphones are actually related to calls and hence do not constitute the boundaries of genuine user interaction sessions [4]. These findings are based on a state machine based usage session extraction approach capable of avoiding mentioned pitfalls – which we consequently incorporate in our approach (see fig. 1).

4.2 Context Detection Algorithm

Alongside extracting locked and unlocked sessions, we derive the context these sessions occurred in, based on time and location information. While the Device Analyzer dataset provides timestamps, obtaining location information requires some effort. The dataset does not contain GPS information,

[1] The University of Cambridge Computer Laboratory and Data Funder do not bear any responsibility for our analysis or interpretation of the Device Analyzer Dataset or data thereof.

[2] `http://deviceanalyzer.cl.cam.ac.uk/keyValuePairs.htm`

which would be of little use in indoor or urban environments anyway. The Device Analyzer application records coarse locations of devices as returned by the network provider. Since recording such information raises privacy concerns, participants were requested to opt-in for sharing their location for research purposes – which only 1.12% of the users chose to do, precluding further analysis due to sample size.

Hence we derive location information from two other sources of information which can be related to device locations indirectly: GSM cell IDs and Wi-Fi scan results. GSM cell IDs were anonymized by hashing in the dataset to protect participants' privacy. Further, while the option to opt-out from recording anonymized GSM cell IDs existed too, only 2.41% chose to do so – leaving records for 97.59% of participating devices. Wi-Fi scan results, including service set identifier (SSID) and MAC address of Wi-Fi access points within range are anonymized as well and are available for all capable devices in the dataset. An algorithm to extract location context information from handset-based GSM cell ID data has been proposed by Jiménez [5], extended to utilize Wi-Fi scan data by Soikkeli [11] and applied to a study of smartphone usage in [12]. For our research, we implemented the extended algorithm while applying some simplifications for the sake of computation time ([11, 12] applied the algorithm on a dataset of 140 devices while the dataset we use contains 17,103 devices). The algorithm consists of two parts: first, meaningful locations are identified, which requires different approaches for cell ID data and Wi-Fi scan results. Subsequently, contexts such as *home* or *office* are assigned to the identified locations based on time information.

4.2.1 Deriving Places from Cell Data

A mobile phone is almost always connected to a cell tower, uniquely identified by cell identifier (CID) and location area code (LAC). As these attributes are anonymized in the dataset used in this work, we cannot relate them to geographic coordinates by using a database like OpenCellID[3]. However, since a cell tower has a fixed position and a limited range, it could be considered to be one place in terms of user context detection. As cell tower placement aims to minimize areas without network coverage and enhance connectivity robustness, adjacent cells usually overlap each other. Devices may dynamically switch between cells if another one is considered "better" than the current cell. As a result, it is not unlikely for even a stationary mobile phone to be connected to several different cells over the course of time [18]. Moving the device, for instance in an office building, possibly even increases the number of different cells a device is connected to while still being in the same abstract place (e.g., *office* context). In order to obtain places from cell data, adjacent cells therefore need to be clustered. For our implementation, we apply a clustering algorithm based on *minimum circular subsequences* proposed by Yang [18]. Given a sequence of cell IDs a device has been connected to, Yang defines a *circular subsequence* as a subsequence starting and ending with the same cell ID and containing at least two different cell IDs with the *cardinality* being the number of different cell IDs it contains. A *minimum circular subsequence* is a circular subsequence that does not contain other circular subsequences and thus indicates that a device has "returned" to where it was in the beginning. Cells that appear in a minimum circular subsequence of low cardinality are assumed to be co-located and therefore assigned to the same cluster. To avoid the problem of "over-clustering" large areas in situations like stop-and-go traffic on a freeway, cells are clustered around "qualified" cells that appeared at least Q times for at least one day. For our work, we choose $Q = 10$ and a minimum circular subsequence cardinality threshold $S = 2$, as suggested by Yang [18]. Further details on deriving places from cell data are found in [5, 12, 18].

4.2.2 Deriving Places from Wi-Fi Scan Results

Wi-Fi-enabled mobile devices periodically scan for Wi-Fi access points within range. The result contains a list of access points, each described by its MAC address, SSID, received signal strength indicator (RSSI), and frequency. The interval between individual scans ranges from a few seconds to several minutes, depending on factors like OS build, hardware, device state, and connectivity state. The dataset used in this work features an average scan frequency of 129 scans per day.

Since Wi-Fi access points are typically stationary, Wi-Fi scan results are frequently used for location-based services such as indoor positioning and navigation systems. A popular approach is to construct a unique Wi-Fi "fingerprint" of a certain location based on observed unique access point identifiers and corresponding signal strengths and an extensive body of literature exists on various fingerprinting techniques. While previous studies used a fingerprinting-based approach to derive meaningful places from Wi-Fi data [11, 12], we choose a less complex method. Taking the available history of scan results for a single device as input, the following steps are applied to derive contextual places, each identified by a cluster of adjacent access points:

1. Create a sequence A of all known access points, sorted descending by the number of occurrences.
2. The first access point from A constitutes the root R of a new cluster C.
3. For each scan result containing R, add all other discovered access points to C.
4. Remove from A each access point contained in C.
5. If A is not empty, proceed with step 2.

While this approach is less sophisticated and assumable less accurate than a fingerprinting-based approach, it is also less complex and computational intensive which has to be taken into account, given we are processing approximately 10 TB of raw data. Assuming a Wi-Fi access point has a maximum indoor range of 50 meters, a cluster spans at most a circular area with a diameter of 150 meters (imagine a cluster containing three access points with the root R located in the middle and the other two access points opposed to each other as far away as possible while still maintaining an overlap with R). Given we are trying to identify places such as home and office (and keeping in mind that in contrast, GSM cells can have a range of several kilometers), we argue that the granularity of our approach is sufficient for the study at hand, allowing us to avoid a more computationaly expensive fingerprinting-based approach.

[3] http://opencellid.org

4.2.3 Context Detection

Time information is one of the most important information available to detect user context [1]. Making some basic assumptions about standard users' diurnal patterns allows us to draw educated guesses on *home* and *office* contexts: In order to put a contextual meaning to the places derived from cell and Wi-Fi scan data, we assume that under normal circumstances a standard user

- does not sleep in the office,
- is at home during night hours (12 a.m. and 6 a.m.),
- works between 10 a.m. and 4 p.m. on workdays,
- and does not go to work on weekends.

While these assumptions are obviously fuzzy and oversimplified considering e.g., night shifts, home workers, holidays, unemployment, or traveling salesmen, previous research shows that results are still fairly accurate. Based on similar assumptions, Jiménez [5] was able to detect *home* contexts with an accuracy of 66% and *office* contexts with an accuracy of 74% (n = 578) while Verkasalo [14] reported classifying 70% of contexts correctly (n = 87), both solely using places derived from cell information. Due to the lack of ground truth in our dataset, we are unfortunately not able to state accuracy measures of detected contexts. However, as we combine both Wi-Fi and cell data for context location detection and since places derived from Wi-Fi are considerably more accurate than cell-based places we assume to achieve similar or better accuracy than reported in [5, 14].

To detect *home* and *office* context we apply an algorithm based on Soikkeli [11] to both cell-based and Wi-Fi-based places. At first, places that have been visited more often than the average number of visits across all derived places are considered to be *meaningful* places. Place not classified as *meaningful* places are assigned the *elsewhere* context. Further, meaningful places are considered to be *office* context if both

$$\frac{\text{Visits during weekends}}{\text{Total visits}} < 0.2 \qquad (1)$$

$$\frac{\text{Visits during weekday working hours}}{\text{Visits during weekdays}} > 0.5 \qquad (2)$$

Home context is assigned to meaningful non-office places if both

$$\frac{\text{Visits on weekday night hours}}{\text{Visits during weekdays}} > 0.25 \qquad (3)$$

$$\frac{\text{Visits on weekdays during non-working hours}}{\text{Visits during weekdays}} > 0.7 \qquad (4)$$

Other meaningful is assigned to all meaningful places neither considered *home* or *office.*

Cell-based and Wi-Fi-based context detection is applied to classify the context of a usage session, depending on which information is available. If for one place contexts derived cell-based and Wi-Fi-based differ, we choose the most specific context in the following order: *1) home 2) office 3) other meaningful*, and *4) elsewhere.*

5. RESULTS AND DISCUSSION

In this section we present and discuss our results. As device unlocking is related to mobile device security, we first take a look at security-relevant device configuration features available within the dataset not covered by previous empirical research. To evaluate the feasibility of Wi-Fi scan based context detection and similar applications, we discuss the yield of Wi-Fi scans based on 88 million scan results and compare the findings to Bluetooth scans. Subsequently, we outlines the results of the context detection approaches.

Studying locked and unlocked usage sessions for certain characteristics constitute the core results of this work. Examined characteristics include: average device usage time per day, average usage session duration and average amount of usage sessions per day. For each locked, unlocked, and overall usage sessions we compute mean and median number of daily interactions as well as mean and median daily usage time in regard to context and form factor. For each device, this is done by calculating the mean and median for each feature over all observed days. The mean and median locked, unlocked, and overall session durations are calculated across the entire observation period for each device, again in relation to context and form factor. We then calculate the grand mean (mean of the means of all devices) and the grand median (median of the medians of all devices). Table 1 summarizes our results and compares them to findings of previous mobile device usage studies.

Finally, we present results concerning device unlocking duration with a focus on graphical pattern unlock.

5.1 Security-Related Device Configuration

A fundamental aspect of the Android security model is to strictly limit root privileges to the kernel and a small subset of services. This intentionally constrains the capabilities of applications. One way to overcome those restrictions is to gain root privileges by *rooting* the device. Since rooting undermines the security model, it introduces significant security threats [10]. We found that 17.4% of the devices in the dataset are rooted (n = 2,698), which seems to be rather high. However, given the scientific nature of the Device Analyzer project, one has to keep in mind that the dataset potentially includes more devices of technology enthusiast than the overall population, therefore the share of rooted devices might be lower outside the dataset.

While the primary source of Android applications is Google's Play Store, a number of alternative markets exist. As these markets are known to frequently host ad-aggressive applications, plagiarisms, and malware [6], Android prevents installing Android Application Package (APK) files from all sources except Google's Play Store by default. However, Android provides an option to allow installing these so-called "third party applications". We found installation of third party applications to be allowed on 66% of all devices in our dataset (n = 10,883). While allowing third party applications to be installed might be necessary to access services of other companies, like the Amazon App Store, it also raises mentioned security concerns.

5.2 Wi-Fi and Bluetooth Scan Results

As mobile phones are almost always connected to a base station, cell ID-based context detection is usually feasible. Wi-Fi scan-based context detection, however, relies on Wi-Fi being enabled on devices and the presence of visible Wi-Fi access points in close vicinity (around 50 meters indoors). In order to evaluate spatial coverage and achievable accuracy of Wi-Fi scan-based applications such as indoor positioning, navigation, and context detection, we analyzed 88 million Wi-Fi scan results from 8,485 devices. For comparison, we also evaluated 219,967 Bluetooth scan results from 4,686 devices, as Bluetooth is another wireless protocol frequently used for similar applications (e.g., social context detection [11]).

We found that Wi-Fi scans occurred on average 129 times per day while Bluetooth scans were performed only 6 times per day. Single devices detected up to 55 Wi-Fi access points in one scan on average, while across all devices on average 4.7 access points were in range. Bluetooth scans discovered one other bluetooth device on average. For the context detection used in our work, the probability of seeing at least one access point to determine the current location is key. We found the probability of seeing at least one Wi-Fi access point to be high with a mean of 0.90 and a median of 0.96, compared to the probability of seeing another Bluetooth device, which is 0.55 for both mean and median, as depicted in fig. 2.

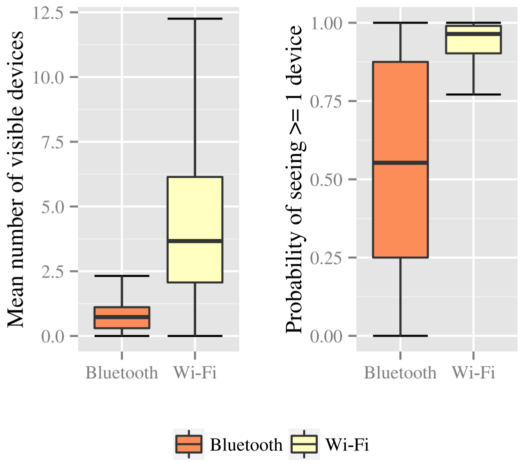

Figure 2: Average number of surrounding devices and probability of seeing at least one other device

5.3 Context Detection

Comparing GSM and Wi-Fi-based location detection, as expected we found Wi-Fi-based location detection to yield better results in most situations. Quality of results was measured by the amount of distinctly detected home and office contexts. We argue reasons therefor are twofold: First, Wi-Fi signals have a smaller signal range compared GSM signals, hence allowing a more precise detection of locations. Secondly, parameterizing the clustering of cell IDs is a trade off between under- and over-clustering, in which either multiple clusters exist for one abstract location or multiple locations are falsely grouped together. Moreover, cell ID information are not available for around half the analyzed tablet devices. However, Wi-Fi-based location detection as well does not always yield results, for instance at work places without any Wi-Fi access points in range.

Hence, combining both location sources improved the overall result in every situation. In particular, *home* context could be detected for 91% of the phones and 79% of the tablet devices, as outlined in fig. 3. For 77% of the phone-type devices, *office* context was detected while only for 48% of the tablet devices an *office* context was found. This is to be expected, considering that tablet devices are less handy and thus less often brought to work, compared to smartphones. To not distort results, we excluded devices for which no *home* context could be detected from consecutive usage session analysis.

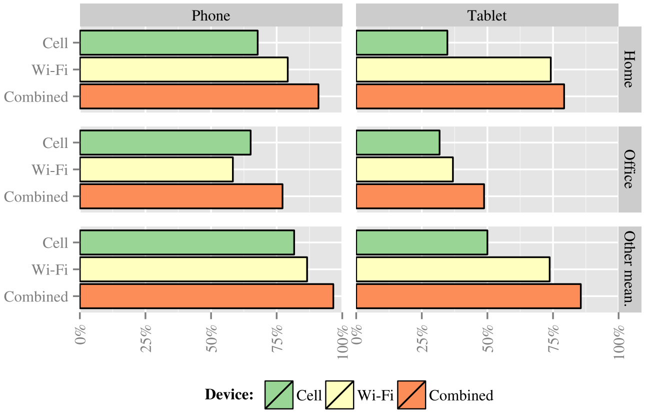

Figure 3: Context detection results

5.4 Number of Daily Interactions

When looking at the number of daily interactions, we confirm observations from [4] that the majority of interactions does not include unlocking the device. Overall, people used their phones on average 57 times per day but only unlocked them for 40% of the interactions. Tablet devices are used about half as often, namely 27 times per day on average with a similar unlocked usage share of 39%. Since locked usage only allows for a limited set of actions, mainly checking information, the high proportion of locked sessions can be explained by *checking habits* as described by Oulasvirta et al. [9].

In the distribution of interactions across the different contexts (see fig. 4 (a)), we see that the number of interactions is fairly evenly distributed across *home*, *other meaningful* and *elsewhere* for both phone and tablet as well as locked and unlocked usage. However, *office* context accounts only for roughly half as many interactions as each of the other contexts, indicating that people use their devices less frequently in work situations compared to leisure activities. Our results in respect of *office* usage are well in line with findings by Soikkeli et al. [12], who reported that 12% of smartphone usage sessions occur in *office* context and 29% *elsewhere*. Our results indicate that the share of smartphone sessions in *office* context is 16% while 25% occur *elsewhere*. However, Soikkeli et al. [12] found that 47% of the sessions take place in a *home* situation while *other meaningful* places only account for 9% of the sessions. We found, though, the share of sessions in *home* context to be 28% while *other meaningful* places

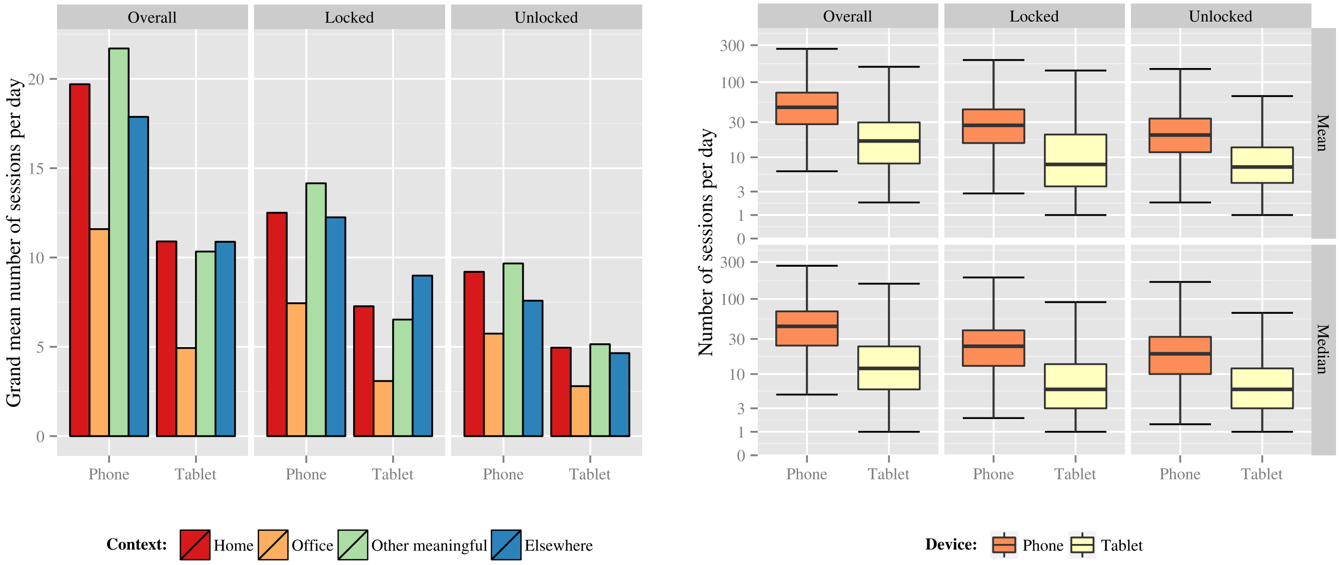

(a) Grand mean of number of sessions per day by context

(b) Distribution of sessions per day across devices

Figure 4: Daily interactions

accumulate 31% of the usage sessions. This effect might be introduced by different user panels: the Device Analyser dataset we use contains users from 175 different countries and is not limited to specific professions, age groups, or life styles, while the panel used in [12] consists mainly of Finnish male university students.

All previous studies observed high diversity across users and usage sessions in terms of both frequency and duration [2, 4, 8, 12, 16]. Since the average number of interactions therefore is limited in adequately describing the panel as a whole (as it is biased by a few rather high values), fig. 4 (b) illustrates the distribution of both mean and median of daily interactions per device.

5.5 Session Duration

Regarding session duration, we found that in general, usage sessions on tablet devices last more than twice as long as phone usage sessions. Locked sessions on average last 94 seconds on phones (median 11 seconds) while spanning 206 seconds on tablet devices (median 15 seconds). As locked usage sessions are usually short, they are more prone to distortion caused by display timeouts counted towards usage time in cases in which the user does not manually switch off the device's screen, which technically marks the end of the usage session [4]. Locked sessions being longer for tablet devices compared to smartphones are hence arguably caused by the fact that tablets are configured with an average display timeout of 6,6 minutes while smartphones feature a mean display timeout of only 2.8 minutes.

Average unlocked sessions span 5 minutes on phones (median 1.2 minutes) while lasting for 11.6 minutes on tablet devices (median 3.3 minutes). Interestingly, context does have a noticeable effect on session duration (see fig. 5 (a)): In *home* context, sessions on both tablet and phone devices are considerably longer than in other contexts while sessions in *office* context are usually the shortest. On tablet devices, for instance, unlocked sessions in home context have an average duration of 11.4 minutes while in office context, unlocked sessions would only last 6.7 minutes.

Again, session duration is highly diverse, both across sessions and across users by more than an order of magnitude. For example, the median across the mean unlocked session lengths of tablet devices is 8.8 minutes, compared to a mean of 11.6 minutes, which is biased by a mean session length of up to 89 minutes on some devices. Figure 5 (b) therefore again depicts the distribution of both mean and median of the session duration per device for locked, unlocked and overall usage.

5.6 Daily Usage Duration

We found that the average locked device usage per day for phones and tablets is nearly equal (43 minutes vs 36 minutes), as the tablets' longer sessions compensate for the higher number of sessions on phones. Unlocked usage of tablet devices sums up to 88 minutes per day (median 53 minutes), while phones are used on average 86 minutes per day (median 58 minutes). Overall, phone usage amounts to 117 minutes per day (median 82 minutes) while tablets feature an overall usage of 112 minutes (median 67 minutes). As with individual session length, *home* context accounts for the largest share of usage while *office* has the smallest share per context of daily usage.

The average device usage per day is again dominated by a small amount of devices accumulating an excessive amount of daily usage. Some phones featured an average usage per day of almost 15 hours while the maximum average usage of tablet devices is 7 hours. The median of the overall mean daily usage is, however, 106 minutes for phones and 99 min-

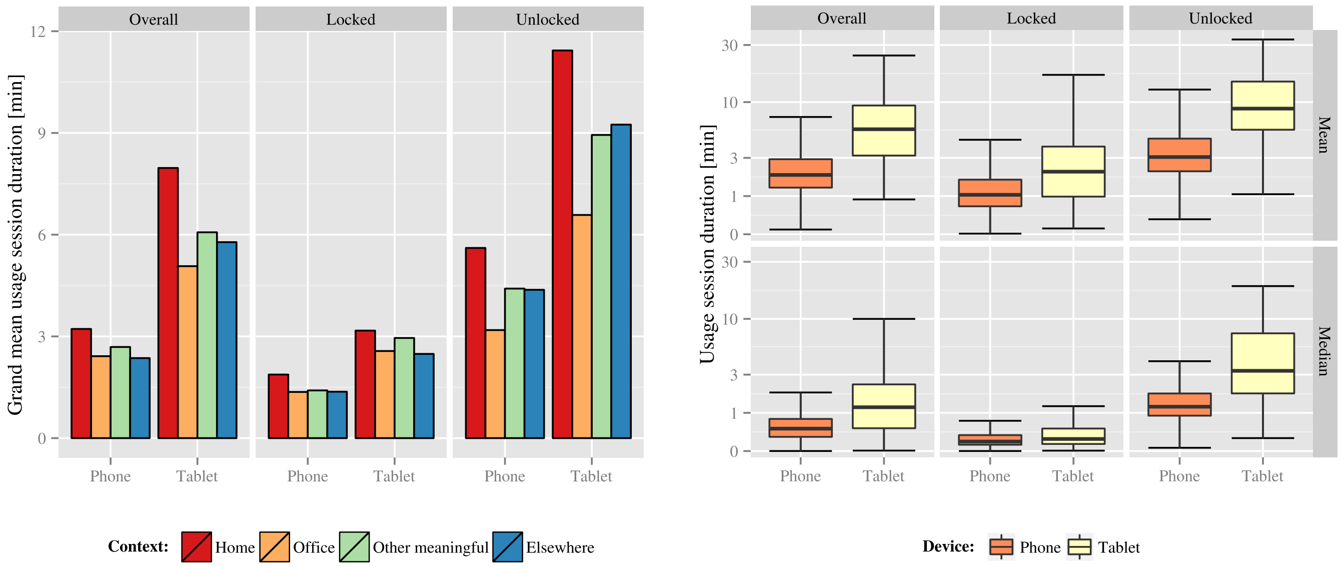

(a) Grand mean of session duration by context

(b) Distribution of session duration across devices

Figure 5: Session duration

utes for tablet devices. Figure 6 (b) depicts the distribution of both mean and median of daily usage for locked, unlocked, and overall device usage.

5.7 Device Unlocking

Apart from analyzing unlocked device usage, we analyzed how users lock their devices based on all devices within the original dataset featuring the required information. Unlocking a device requires either slide-to-unlock or some form of authentication like PIN, password, or graphical pattern. Since the underlying dataset unfortunately only labels graphical pattern-based unlocking explicitly, means of comparing different authentication methods are limited. However, pattern unlock seems to be quite popular, as it is enabled on 42% of the smartphones (n = 4,152) and on 28% of the tablet devices (n = 418). Of these devices, 74% are configured to provide visual feedback while entering the pattern, increasing the vulnerability to so-called shoulder surfing attacks, i.e., capturing the secret pattern by looking over the user's shoulder during device unlocking [17]. On 5% of the phones and 9% of the tablets no form of device locking, not even slide-to-unlock, is enabled (see fig. 7).

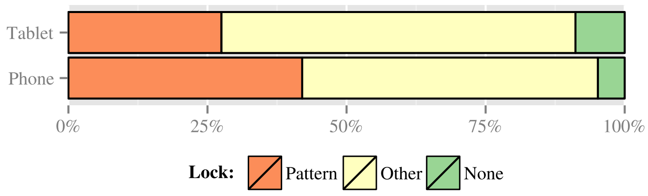

Figure 7: Usage of different locking mechanisms

One aspect of the usability of unlocking mechanisms is the speed at which the device can be unlocked. Using the state machine approach described in fig. 1, we measure the time between turning the device's screen on and unlocking the device, indicated by a *USER_PRESENT* intent broadcasted by the Android system when the device is unlocked. The 8.6 million unlocking sessions we extracted that way, however, also contain sessions in which the user turns the device on but only attempts to unlock it after several minutes (given a long display timeout is configured). From a comprehensive real world study conducted by Zezschwitz et al. [19] we know that unlocking takes on average 1.5 seconds for PIN-based mechanisms and 3.1 seconds for pattern-based unlocking. We therefore reasonably choose an arbitrary upper limit of 10 seconds and only take shorter unlocking sessions into account, which leaves us with 6.8 million sessions.

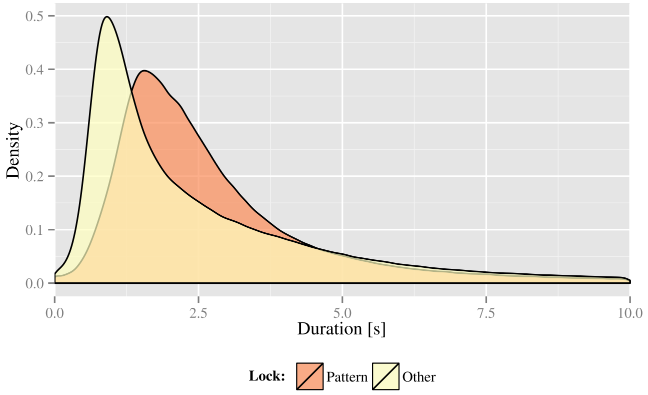

Figure 8: Density of unlocking session duration

Our results confirm the observation of Zezschwitz et al. [19] that pattern unlock requires notably more time than other unlocking mechanisms like PIN, as the unlocking duration distribution (see fig. 8) illustrates. Looking at sessions shorter than 10 seconds, we find that pattern unlock requires on av-

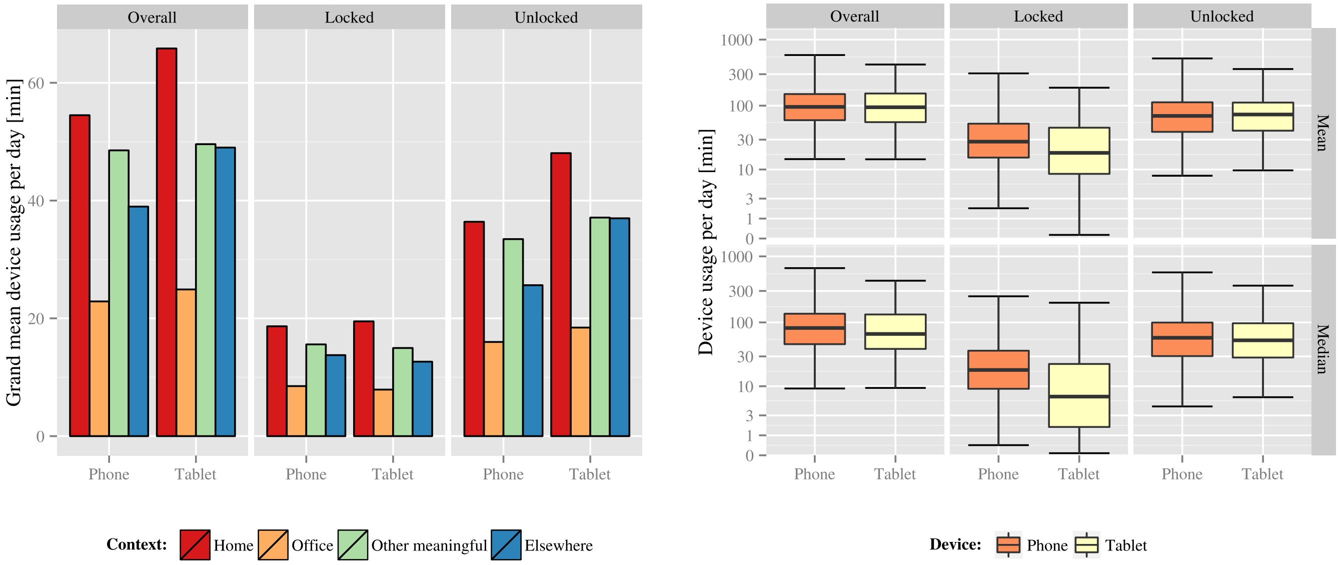

(a) Grand mean of device usage per day by context

(b) Distribution of daily usage across devices

Figure 6: Daily device usage

erage 2.7 seconds (median 2.2 seconds) while other unlocking methods take only 2.5 seconds on average (median 1.7 seconds).

6. CONCLUSION

In this work we studied locked and unlocked mobile device usage with respect to device form factor and user context. For our study we extracted a total of 23 million usage sessions from 100 billion mobile device usage records using a sophisticated screen power state machine-based approach. Using anonymized GSM cell IDs, Wi-Fi scan results and timestamps of records we derived location information for usage sessions. By making reasonable assumptions about standard users' diurnal patterns, we were able to draw educated guesses about users' locational context, identifying *home* context for 91% and *office* context for 77% of the smartphone devices.

Consistent with previous studies we found high diversity in device usage characteristics, both across sessions and users, vary with more than an order of magnitude. We observed that on average, smartphones are used around twice as much per day as tablet devices (58 times vs. 27 times). However, devices are unlocked in only 40% of the interactions. Given the limited forms of interaction available in locked state, the high share of locked usage indicates that the majority of usage constitutes some form of short information checking. Our results show that 16% of smartphone usage occurs in *office* context and 28% in *home* context. Contrary to the number of interactions, we found that the duration of usage sessions is in general more than twice as long for tablets compared to smartphones: on average, unlocked sessions on phones last 5 minutes while tablet usage sessions account for 11,6 minutes. As a result, the daily usage of both smartphones and tablets is nearly equal (117 minutes vs. 112 minutes). Again, home context accounts for the largest share of usage while office has the smallest share per context of daily usage. Alongside the usage session analysis we investigated the process of unlocking a mobile device in terms of speed, which we found to be on average 2.7 seconds for graphical pattern and 2.5 seconds for other unlocking mechanisms. Furthermore, in terms of security-related device configuration, we noted that with 17.4% a fairly high number of devices in the dataset is rooted while 66% of the devices allow the installation of applications from untrusted sources known to frequently host malware.

Our work shows that despite offering similar technical capabilities, smartphones and tablets are used quite differently. While substantial research has been conducted in regard of smartphone usage, little work has been done to analyze tablet usage. With the increasing ubiquity of mobile devices, people tend to simultaneously own and use several devices of different form factors like phones, tablets, and smartwatches. Further research is needed, e.g., on when and why users transition between different device types.

7. ACKNOWLEDGEMENTS

This work has been carried out in part within the scope of u'smile, the Josef Ressel Center for User-Friendly Secure Mobile Environments. We gratefully acknowledge funding and support by the Christian Doppler Gesellschaft, A1 Telekom Austria AG, Drei-Banken-EDV GmbH, LG Nexera Business Solutions AG, and NXP Semiconductors Austria GmbH.

We also thank the University of Cambridge for providing access to the Device Analyzer dataset and Andrew Rice for instantaneously answering all our questions.

	Devices	**Daily interactions**						**Session length [sec]**						**Daily usage [min]**					
		overall		locked		unlocked		overall		locked		unlocked		overall		locked		unlocked	
Smartphones		mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.
Falaki et al. [2]	255	10-250	-	-	-	-	-	10-250	-	-	-	-	-	30-500	-	-	-	-	-
Oliver [8]	17,300	87	76	-	-	-	-	68	20	-	-	-	-	101	79	-	-	-	-
Soikkeli et al. [12]	140	-	-	-	-	20	-	-	-	-	-	207	45	-	-	-	-	73	-
Truong et al. [13]	10	-	-	-	-	5-105	-	-	-	-	-	-	-	-	-	-	-	-	-
Hintze et al. [4]	1,969	57	44	33	22	25	20	177	-	88	56	285	192	117	97	33	21	87	71
Our study	1,487	58	44	37	24	25	19	165	30	94	11	299	74	117	82	43	18	86	58
Tablets		mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.
Our study	98	27	12	17	6	11	6	414	73	206	15	694	197	112	67	36	7	88	53
Smartphones & Tablets		mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.	mean	med.
Wagner et al. [16]	16,000	57	-	-	-	-	-	116	-	-	-	-	-	123	79	-	-	-	-

Table 1: Comparison of usage session characteristics in different mobile device usage studies

References

[1] G. Chen and D. Kotz. A survey of context-aware mobile computing research. Technical Report TR2000-381, Dartmouth College, 2000.

[2] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, and D. Estrin. Diversity in Smartphone Usage. *Proc. MobiSys 2010*, 2010.

[3] Gartner. Gartner says worldwide tablet sales grew 68 percent in 2013, with android capturing 62 percent of the market, 2014. URL `http://www.gartner.com/newsroom/id/2674215`.

[4] D. Hintze, R. D. Findling, M. Muaaz, S. Scholz, and R. Mayrhofer. Diversity in Locked and Unlocked Mobile Device Usage. In *Proceedings of the 2014 ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication*, 2014. To appear.

[5] B. Jiménez. *Modeling of Mobile End-User Context.* Master's thesis, Helsinki University of Technology, 2008.

[6] M. Lindorfer, S. Volanis, A. Sisto, M. Neugschwandtner, E. Athanasopoulos, F. Maggi, C. Platzer, S. Zanero, and S. Ioannidis. Andradar: Fast discovery of android applications in alternative markets. In S. Dietrich, editor, *Detection of Intrusions and Malware, and Vulnerability Assessment*, volume 8550 of *Lecture Notes in Computer Science*, pages 51–71. Springer International Publishing, 2014.

[7] H. Müller, J. L. Gove, and J. S. Webb. Understanding tablet use: A multi-method exploration. In *Proceedings of the 14th Conference on Human-Computer Interaction with Mobile Devices and Services (Mobile HCI 2012)*, 2012.

[8] E. Oliver. The Challenges in Large-Scale Smartphone User Studies. *Proc. HotPlanet 2010*, 2010.

[9] A. Oulasvirta, T. Rattenbury, L. Ma, and E. Raita. Habits make smartphone use more pervasive. *Personal and Ubiquitous Computing*, 16(1):105–114, 2011.

[10] Y. Shao, X. Luo, and C. Qian. Rootguard: Protecting rooted android phones. *Computer*, 47(6):32–40, June 2014.

[11] T. Soikkeli. *The effect of context on smartphone usage sessions.* Master's thesis, Aalto University School of Science, 2011.

[12] T. Soikkeli, J. Karikoski, and H. Hämmäinen. Diversity and End User Context in Smartphone Usage Sessions. In *Proc. NGMAST 2011*, pages 7–12, 2011.

[13] K. N. Truong, T. Shihipar, and D. J. Wigdor. Slide to X: Unlocking the Potential of Smartphone Unlocking. In *Proc. CHI 2014*, pages 3635–3644, 2014.

[14] H. Verkasalo. Contextual patterns in mobile service usage. *Personal and Ubiquitous Computing*, 13(5):331–342, 2008.

[15] D. T. Wagner, A. Rice, and A. R. Beresford. Device Analyzer: Large-scale mobile data collection. In *Big Data Analytics workshop, ACM Sigmetrics 2013*, 2013.

[16] D. T. Wagner, A. Rice, and A. R. Beresford. Device Analyzer: Understanding smartphone usage. In *Proc. MobiQuitous 2013*, 2013.

[17] S. Wiedenbeck, J. Waters, L. Sobrado, and J.-C. Birget. Design and evaluation of a shoulder-surfing resistant graphical password scheme. In *Proceedings of the Working Conference on Advanced Visual Interfaces*, AVI '06, pages 177–184, 2006.

[18] G. Yang. Discovering significant places from mobile phones – a mass market solution. In R. Fuller and X. Koutsoukos, editors, *Mobile Entity Localization and Tracking in GPS-less Environnments*, volume 5801 of *Lecture Notes in Computer Science*, pages 34–49. Springer Berlin Heidelberg, 2009.

[19] E. V. Zezschwitz, P. Dunphy, and A. D. Luca. Patterns in the Wild: A Field Study of the Usability of Pattern and PIN-based Authentication on Mobile Devices Emanuel. *Proceedings of the 15th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 261–270, 2013.