

Clear All: A Large-Scale Observational Study on Mobile Notification Drawers

Dominik Weber
VIS, University of Stuttgart
Stuttgart, Germany
dominik.weber@vis.uni-stuttgart.de

Alexandra Voit
VIS, University of Stuttgart
Stuttgart, Germany
info@alexandra-voit.de

Niels Henze
University of Regensburg
Regensburg, Germany
niels.henze@ur.de

ABSTRACT

Notifications are an essential feature of smartphones. The notification drawer is the central place to view and attend notifications. Although a body of work already investigated how many and which types of notifications users receive and value, an in-depth analysis of notification drawers has been missing. In this paper, we report the results of a large-scale observational in-the-wild study on mobile notification drawers. We periodically sampled the notification drawer content of 3,953 Android devices, resulting in over 8.8 million notification drawer snapshots. Our findings show that users have, on average, 3.4 notifications pending in the notification drawer. We saw notifications accumulate overnight and being attended to in the morning. We discuss the prominent positioning of messaging notifications compared to other notification types. Finally, inspired by prior work on the management of email inboxes, we propose the three user types “Frequent Cleaners”, “Notification Regulators”, and “Notification Hoarders” and discuss implications for future notification management systems.

CCS CONCEPTS

• **Human-centered computing** → **User studies; Field studies; Empirical studies in HCI.**

KEYWORDS

Notification, drawer, center, panel, interruptions, smartphones, in-the-large, in-situ study.

ACM Reference Format:

Dominik Weber, Alexandra Voit, and Niels Henze. 2019. Clear All: A Large-Scale Observational Study on Mobile Notification Drawers. In *Mensch und Computer 2019 (MuC '19)*, September 8–11, 2019, Hamburg, Germany. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3340764.3340765>

1 INTRODUCTION

When users of current smartphones turn on the display of the device, they are usually greeted with a lock screen, consisting of a large clock and a list of notifications below it. On the dominant mobile operating systems Android and iOS, this list of notifications can also be accessed at any time by swiping down from the top of the screen. This universally accessible list is an important feature of current smartphones, as it enables asynchronous communication and provides users with proactive information. The notification list is commonly referred to as the notification *drawer* (Android), notification *center* (iOS), notification *tray*, or notification *panel*. We use the term notification drawer in the following.

The number of notifications exploded in the last couple of years. Mobile devices began to be connected to the Internet around the clock, and app stores allowed all kinds of applications and services to run on devices. Recent work investigated how many and which kind of notifications users receive on their smartphones and how users perceive different kinds of notifications [35, 40]. Messaging notifications have shown to be of high importance to users, as they enable users to stay connected with their contacts [39]. However, not all notifications are of equal importance. Research has shown that notifications can cause interruptions, which can induce adverse effects such as decreased work performance and inattention [13, 26, 31, 32]. A field of research, therefore, focused on reducing these adverse effects. Various approaches were explored, from disabling notifications altogether [38], to using context-aware models to delay notifications to opportune moments [19, 20, 34]. Although a large body of prior work on notifications exists, the notification drawer on smartphones as the central place to view and attend notifications has not been explored in detail so far. However, this is a crucial aspect for a complete understanding of mobile notifications.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MuC '19, September 8–11, 2019, Hamburg, Germany

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-7198-8/19/09...\$15.00

<https://doi.org/10.1145/3340764.3340765>

In this paper, we complement prior work by reporting the results of a large-scale observational study on notification drawers of current smartphones. Using a research-in-the-wild approach, we periodically sampled the contents of notification drawers on Android devices. We collected 8.8 million notification drawer snapshots from almost four thousand devices. Our contribution is three-fold. 1) We present a novel analysis on the number of notifications in notification drawers and the positioning of different notification categories. 2) Based on prior work on the management of email inboxes [48], we propose three different user types regarding the management of notifications. 3) We published the collected data set of notification drawer snapshots to foster research on mobile notifications.

2 RELATED WORK AND BACKGROUND

In the following, we provide an overview of related work on mobile notifications, disruptive effects, notification management, and research in-the-wild. Subsequently, we briefly provide background information on notifications in the Android mobile operating system.

Mobile Notification and Messaging

In a study comparing different kinds of “smart” devices, Weber et al. found that the smartphone was the preferred device to be notified on [45]. Sahami Shirazi et al. conducted a large-scale assessment of mobile notifications [40]. The researchers collected approximately 200 million notifications from 40,000 users. They found that users value messaging notifications and notifications about people and events. Further, the researchers compared click times for different types of notifications. They found that messaging and system notifications were clicked on the fastest, and news notifications the slowest. In a smaller scale study with 15 participants, Pielot et al. found that participants received 63.5 notifications per day [35]. Most of these notifications were related to messaging and email. The researchers found that participants were fast to attend to these notifications and typically did not let notifications accumulate. Pielot et al. recently revisited mobile notifications in a study with 278 participants [39]. The results again showed the importance of messaging notifications. Participants were fast to attend messaging notifications, while other types of notifications were either removed quickly or left unattended for longer periods.

Messaging is a recurrent topic in prior work. Instant messaging is a flexible way of communication, that can vary between synchronous and asynchronous discussions [5, 27, 33]. Researchers investigated “traditional” SMS usage [6] and compared it with modern instant messaging (IM) apps such as *WhatsApp* [11]. For instance, Church et al. found that cost and social influence are reasons for *WhatsApp* overtaking SMS messaging [11].

Disruptive Effects and Digital Well-being

Being always connected and reachable massively shifted our attentiveness towards messaging [17, 36]. Birnholtz et al. investigated “unavailability” in an always-connected world [8]. Lee et al. investigated smartphone “overuse” and the role of messaging [28]. Aranda and Baig discussed how users are more and more dependent on smartphones, difficulty to disconnect, and “the fear of missing out” [3].

While notifications allow us to be connected, they can also cause interruptions. The disruptive nature of interruptions and task switching has been an important research topic for many years [13, 26]. While not all interruptions are disruptive [22], Adamczyk and Bailey showed that different timings of interruptions have different effects on users [1]. Mehrotra et al. found that the perceived disruption of a notification is influenced by several factors, including the notification’s presentation, the relationship of the sender and receiver, and the task the user is engaged in [31, 32].

Mobile Notification Management

A body of prior work has explored how mobile notifications can be better managed. Researchers investigated what users do when they sense notifications [10], and which strategies users apply to cope with notifications. Gallud and Tesoriero suggest a movement from sound to visual notifications [21]. Weber et al. explored the idea of a notification dashboard that allows users to reflect on how many notifications they receive on a daily basis [46]. The researchers also investigated “snoozing” of notifications [43], i.e., allowing users to temporarily dismiss notifications from the notification drawer and re-triggering them after a user-defined duration or point-in-time. In a study with 295 participants, the researchers found that notifications related to people and events were snoozed most often. Auda et al. explored a system for rule-based notification deferral by suppressing, summarizing or automatically snoozing notifications [4]. Mehrotra et al. took this a step further by automatically suggesting rules based on usage patterns [30]. The researchers found that the notification’s title and the user’s location can be used as features to determine whether a message will be dismissed.

Anderson et al. recently published a survey on attention management systems [2]. A number of research projects are focusing on the approach to deliver notifications at opportune moments, instead of delivering them immediately [19, 20]. With *Attelia*, Okoshi et al. developed a middleware that defers notifications to so-called breakpoints - times between two consecutive activities [34]. Deferring notifications to these breakpoints has been shown to lessen disruptive effects [20]; however, this has to be balanced with social expectations to reply quickly [43].

Research in the Wild

A challenge of research on mobile notifications is the fact that they are highly context-dependent and received around the clock. To overcome this challenge, most prior work in this area moved from lab studies to in-the-wild studies. Henze and Pielot explored how app stores can provide external validity for mobile HCI [24]. Henze et al. also discussed the trade-off between opt-in and opt-out for consent in in-the-wild studies [25]. While opt-out allows for greater data collection, it also poses legal and ethical challenges. Using multiple case studies, the researchers found that many users may use apps only for short periods and that users expect research apps to offer a similar user experience to commercial products.

An example of such a research app was the *Desktop Notifications* service that allowed users to synchronize notifications across devices while enabling researchers to gain insights on notifications in-the-wild from a large user base [42].

Notifications on Android

Notifications were an integral feature of the Android mobile operating system since the first version. Notifications are opt-out, meaning that all installed apps can post notifications by default without asking the user for permission. Notifications may use visual, tactile or sound cues to gain the user's attention. All notifications end up in the notification drawer that is accessible by swiping down from the top of the screen (see Figure 1). Since Android 5.0, notifications are shown on the lock screen by default as well. Notifications can contain action buttons [18], expandable text, and images. Users can click on notifications to take action or swipe to the left or right to dismiss them. By clicking "clear all", users can dismiss all notifications at once.

Summary and Research Motivation

To summarize, notifications are an important part of how users interact with smartphones. They are prominently featured on the lock screen and notification drawer. A body of prior work investigated which notifications users receive, how they are valued, interruptions, and means to reduce adverse effects. However, the notification drawer as the central place to view and attend notifications has yet to be investigated. To fill this gap in prior work and to create a more complete understanding of mobile notifications, we explored notification drawers in an in-the-wild study.

3 STUDY

We conducted a large-scale observational in-the-wild study on the content of notification drawers. Our research question was how many and which kind of notifications can be found in notification drawers, and whether different notification management approaches exist.




Figure 1: The Android 9.0 (Pie) notification drawer showing four different kinds of notifications about a new message, news articles, traffic updates, and the current weather.

Apparatus

We developed an Android app that allowed us to snapshot the content of notification drawers in-the-wild in an unobtrusive manner. Our goal was for users to install the app on their own, without explicitly recruiting participants. Inspired by prior work of Weber et al. [43], we developed an Android app that allows users to log and explore their notifications in a local history [44]. The added value for users is the option to look up accidentally dismissed notifications or reflect on notifications they received throughout the day.

The app supports the Android versions 5.0 - 9.0, with Android 9.0 being the most recent Android version at the time of writing. According to the Android distribution dashboard [14], this covers 96.50% of all active Android devices. The app uses the *Notification Listener Service* API [15] to access the notifications on the device. The notification access for the app has to be explicitly enabled by the user in the device settings. Once enabled, all newly created notifications are stored in a local *SQLite* database. Users can then browse their notifications in a list and select individual notifications to read the text in detail.

Data Collection and Consent

After the user installed the app and permitted the app to access the device's notifications, the app displayed a dialog asking the user to opt into the anonymous data collection.

Inspired by prior research on asking for consent in in-the-wild studies [37], the dialog contained the options “Agree” and “No thanks”. The dialog was only shown once. Users could also enable or disable the data collection at any later point in time in the app’s settings. Declining the anonymous data collection did not negatively impact the main functionality of the app in any way. If the user consented to the data collection, the app would periodically snapshot all pending notifications in the notification drawer. We used the *Android-Job* library [12] to schedule the sampling. The library abstracts from version differences in the Android SDK. We set the sampling job to be executed every 15 minutes, which is the minimum amount of time between two jobs. In later versions of Android, these jobs might be deferred if the device uses battery saving features such as the *Doze Mode*, which defers background processes if the device was not used and not moved for a certain amount of time. Each snapshot contained the following features:

- A randomly generated unique ID for the device (UUID) to associate multiple snapshots with a specific device.
- The current Android version, device model, product name, and device manufacturer.
- The current timestamp and timezone.
- Meta-data of all notifications in the notification drawer, such as the package name, timestamp of creation and position in the drawer.

Snapshots generated by the app were limited to meta-data and did not contain text or images. The snapshots were stored in a separate local *SQLite* database.

Procedure

We published the study app on the Google Play Store. Users from all over the world were able to download it for free. We did not advertise the app in any way. Instead, users found the app using the Google Play Store search or by reading articles and watching videos that reported on the app. If a user decided to opt into the data collection, the locally stored snapshots were periodically sent to a server hosted at our university using a secure connection. To avoid negatively impacting the device, the app only sent data over WiFi and if the battery was not low. If the server did not acknowledge the data, the app would re-try sending the data.

Data Filtering

We defined a set of filter rules on the collected snapshots and excluded all devices that did not match the rules:

- (1) The time delta between the first and last snapshot is at least one week (7 days).
- (2) There are at least 672 snapshots for the device. This assumes a snapshot every 15 minutes (4 per hour), for each hour of the day (24), for each day of a week (7).
- (3) The maximum time delta between two snapshots is less than 48 hours. Larger deltas might happen if a device is turned off for extended periods.
- (4) No snapshot is missing. On the device, each snapshot is assigned an ascending ID. This ID is sent to the server along with the snapshot, which allows us to identify missing snapshots.
- (5) No snapshot has invalid timestamps, i.e., the timestamp associated with a snapshot is within the data collection period. Invalid timestamps might happen because of malfunctions of the clock of the device, failed synchronizations with timeservers or incorrect time/date set by the user.
- (6) At least one snapshot contains at least one notification.

This set of filter rules ensures a valid and consistent data set. It is robust against typical problems of in-the-wild data collection, such as unknown hardware and unstable network connections. In addition to the filter rules, we excluded all snapshots from *Huawei* devices. Our testing showed that many *Huawei* devices have an aggressive battery saving feature that interfered with the notification logging.

4 RESULTS

After filtering the collected data, we ended up with 8, 830, 112 notification drawer snapshots from 3, 953 devices.

Demographic Background

While we did not ask the users about their demographic background directly, we can infer some information from the devices. We found that the language of the devices was set to Turkish most of the time (57.22%), followed by English (17.76%), Spanish (9.89%), and German (5.34%). Overall, we saw 31 different languages (grouped language variants). We also looked at the time zones configured on the devices as reported by the Android system (e.g., “Europe/Istanbul”). Most devices were set to a European timezone (64.99%), followed by Asia (18.54%), America (13.84%), Africa (1.82%), Australia (0.30%), and other (0.52%). In total, we saw 158 different time zone configurations, which shows the international user base of the app. In terms of devices, most devices were manufactured by Samsung (63.22%), followed by LG Electronics (5.59%), and General Mobile (5.21%). We saw 74 different manufacturers in total. Compared to the global average [14], the devices used more recent versions of the Android operating system. The Android versions used were Android 5.x (7.89%), Android 6.0 (21.35%), Android 7.x (23.13%), Android 8.x (44.06%), and Android 9.0 (3.57%).

Collected Snapshots

Overall, we collected snapshots for a minimum of 7 days and a maximum of 110 days ($Md = 20$ days). In this time frame,

Clear All

we collected between 673 and 14,257 snapshots per device ($Md = 1,631$). We first investigated whether we managed to collect an even distribution of snapshots across the day. We collected $M = 93.14$ ($SD = 8.93$) snapshots per device for each hour of the day. The number of snapshots per hour decreases slightly at night. This was expected, as battery saving mechanisms in modern Android smartphones delayed the execution of our background process when the devices were idle. The average time delta between two subsequent snapshots was 17.74 minutes ($SD = 3.56$), which is close to our target of a snapshot every 15 minutes.

Number of Notifications in the Notification Drawer

Each notification drawer snapshot contains zero or more notifications. Counting all notifications of all snapshots revealed that we collected a total of 40,836,340 notifications. The same notification may appear in multiple snapshots if it has not been dismissed by the user, notifying app, or Android system. Thus, we identified how many unique notifications we were able to capture. For each notification in each snapshot, we extracted the PACKAGE_NAME, NOTIFICATION_ID, NOTIFICATION_TAG, and CREATION_TIMESTAMP. The combination of these values allowed us to identify unique notifications across snapshots. We collected between 65 and 55,703 ($Md = 1,514$) unique notifications for each device, with a total of 10,928,880 unique notifications.

Snapshots without Notifications. About one-fifth of all snapshots (20.53%) did not contain notifications. The other 79.47% snapshots contained between 1 and 160 notifications.

Average Number of Notifications (Total). Next, we calculated the average number of notifications per device. The left side of Figure 2 shows, that when considering all snapshots, we saw $M = 4.30$ ($SD = 5.86$) notifications in the notification drawer. The median number of notifications was 2.68. While most devices had less than five notifications on average, we also saw a number of outliers. 77.64% of devices had between $[0, 5)$ notifications on average, 14.87% between $[5, 10)$, and 7.49% more than 10 with a maximum of 70.53.

Average Number of Notifications (Grouped). The previously reported numbers represent the total number of notifications as reported by the Android system. However, in Android multiple notifications can be visually grouped, reducing the number of notifications actually shown to the user. For instance, Figure 1 shows a single *Google News* notification with two headlines and an indicator about two additional headlines. The notification can be expanded to allow the user to explore the four individual headlines, and to click or dismiss them individually. Internally, this single notification is represented as four individual notifications for the headlines and a summary notification to visually group them, totaling in

MuC '19, September 8–11, 2019, Hamburg, Germany




Figure 2: The mean number of notifications in the notification drawer over all snapshots. Total: All notifications, as reported by the Android system. Grouped: Visually grouped notifications. Outliers omitted.




Figure 3: Histogram of the mean number of notifications (grouped) in the notification drawer per device. 62 devices had >20 notifications on average ($max = 61.28$).

five notifications. Many instant messaging and email apps make use of this feature to group conversations. To find out how many notifications are actually visually shown to users, we processed all snapshots and counted each notification group as one. Thus, the notification count in Figure 1 would be reduced from 8 to 4.

With this calculation in place, the average number of grouped notifications was 3.40 ($SD = 4.59$), with a median number of 2.17 (see right side of Figure 2). As Figure 3 shows, 84.59% of devices had between $[0, 5)$ grouped notifications on average, 10.35% between $[5, 10)$, and 5.06% more than 10 with a maximum of 61.28. For the remaining analysis, we report on the visually grouped notifications as this better reflects how users see notifications.

Average Number of Notifications (Per Hour). Previous work has shown that the number of notifications users receive throughout the day drops significantly between midnight and 6am [35, 43]. Interestingly, as shown in Figure 4, the average number of notifications in the notification drawer *increases* in this time frame, with a peak at 6am. While users receive fewer notifications at night, they are also likely asleep and therefore do not dismiss notifications. Consequently, the number decreases again as users wake up and start attending the notifications. This shows an opportunity of assisting




Figure 4: The mean number of visually grouped notifications in the notification drawer by the hour of the day. The number increases between midnight and 6am.

users in managing their notifications in the morning to ease the start in the day.

Average Number of Notifications (Per Weekday). Looking at the number of notifications in the notification drawer for each day of the week, we saw little differences, with only a slight drop on Sunday ($M = 3.31$) compared to the overall average ($M = 3.40$).

Number of Apps

Looking at the notifications in more detail, we saw between 4 and 111 ($Md = 28$) different apps per device that created at least one notification. In total, we saw 8,823 different apps that triggered at least one notification. Only 24 apps were used on $\geq 1,000$ devices and 908 apps on ≥ 10 devices. A long tail of apps was used on < 10 devices, with over half of the apps (56.24%) only being used on one device.

Of the ten apps used on most devices, five were system apps including the *Google Play Store* (3,416 devices). The other five apps were the instant messaging app *WhatsApp* (3,591 devices), the social media network *Instagram* (2,962), the video-sharing app *YouTube* (2,773), the *Google Chrome* web browser (2,603), and the *Google Maps* app (2,381).

App Categorization

In line with prior work, we categorized the apps. We based the categories on the 12 categories used by Weber et al. [43], which in return is based on the work by Böhmer et al. [9] and Sahami Shirazi et al. [40]. Additionally, we introduced the category *Navigation* and extended the categories *Social* to *Social & Dating* and *News* to *News & Weather*.

We focused on the 908 apps with ≥ 10 devices and left the long tail of apps with fewer devices uncategorized. Still, with this number of apps, we were able to categorize 92.0% of the 10,928,880 unique notifications in the data set. Similar

Table 1: This table shows the number of unique notifications and apps per category, and the median number of devices per app for each category.

Category	# Notif.	# Apps	Median Devices/App
Calendar & Rem.	170,284	23	43.0
Email	728,464	14	24.5
Game	102,164	111	19.0
Health & Fitness	137,404	24	17.5
Media	637,850	150	24.5
Navigation	235,631	24	18.5
News & Weather	118,797	48	18.0
Phone	454,586	18	133.0
Shopping & Fin.	85,758	127	21.0
SMS & IM	3,692,077	45	42.0
Social & Dating	843,004	40	35.5
System	1,671,025	94	32.0
Tool	1,177,235	190	21.5
Uncategorized	874,601	7,915	1.0
Σ	10,928,880	8,823	-

to prior work, we automatically extracted the app category from the Google Play Store. It is important to note that the developers of the apps provide the categories on the Google Play Store. The categories might not necessarily reflect which kind of notifications an app creates. Further, 178 apps were not available on the Google Play Store, e.g., due to them being pre-installed by device manufacturers or by being manually installed by users. Two researchers independently went through the apps and manually categorized them. The categories provided by the Google Play Store were used as guidelines. For apps not available on the Google Play Store, the researchers searched the web for more information. Finally, the researchers compared the labeled categories and discussed conflicts until an agreement was reached. Table 1 shows the number of notifications and apps assigned to each category. The categories with the most notifications were *SMS & IM*, *System*, and *Tool*. The categories with most apps were *Tool*, *Shopping & Finance*, and *Media*.

Notification Ranking

Based on the notification categories, we investigated what users typically see when they unlock their phones or open the notification drawer.

Background. Notifications in the Android notification drawer are not simply displayed in chronological order. Instead, the Android system uses a number of signals to rank notifications. The used signals differ between Android versions and




Figure 5: Distribution of which kinds of notifications are shown in the first five positions of the notification drawer.

might be modified by device manufacturers. Some of the most prominent signals are as follows:

- The time when the notification was triggered and how much time has passed since then.
- The priority level which can be set by the notifying app. The priority level values range from *MIN*, *LOW*, *DEFAULT*, *HIGH*, to *MAX*. In Android 8.0 and newer, the priority level has been replaced by the importance level, which features the same values but allows users to overwrite them in the settings.
- Contacts associated with a notification and whether the contacts are marked as favorites by the user.

The idea behind the ranking is that the most relevant notifications for the user are shown at the top of the notification drawer. In 2014, Sahami Shirazi et al. found that “*notifications are for messaging*” and that “*important notifications are about people and events*” [40]. In 2018, Pielot et al. found that messaging notifications have a much higher conversion rate than notifications from other types [39]. This was also reflected in the Android 8.0 update released in 2017. The update introduced a “visual hierarchy” for notifications by first assigning notifications to one of four sections and then ranking the notifications within each section [16]. Notifications in the *Major Ongoing* section are about time-sensitive content. Examples include ongoing phone calls, navigation, timers, and media controls. The *People to People* section focuses on instant messaging notifications and notifications about missed calls. The *General* section contains most other notifications, including reminders and email notifications. Finally, the *By the Way* section includes non-urgent content,

such as weather and traffic updates. On recent versions of Android, these notifications are visually muted by reducing them to a single line and graying them out.

Analysis. Since almost half (47.63%) of the devices in the data set were running Android version 8.0 and newer, we expected the notification ranking to be influenced by these new sections. Indeed, Figure 5 shows the distribution of notification categories for the first five positions in the notification drawer. We limited the Figure to five positions, as this is typically the maximum amount of notifications a user sees on the lock screen or the notification drawer before having to scroll down. We see five dominant notification categories: *SMS & IM*, followed by *System*, *Social & Dating*, *Tool*, *Media*, and *Phone*. *Media* notifications are prominent in the first position, due to playback control notifications that end up in the *Major Ongoing* section. However, we also categorized many apps as media that likely do not show playback controls. *SMS & IM* and *Phone* are focused around the first three positions in the drawer. This is likely due to them being in the *People to People* section. *Social & Dating* notifications might be part of the *People to People* or *General* sections, resulting in a more even distribution across the position in the notification drawer. Finally, *System* and *Tool* notifications made up a large number of apps and notifications and were therefore presented across the first five positions as well.

Notification Priority Levels. Figure 7 shows the priority level distribution of the notifications per category. A large number of *Navigation*, *System*, and *Tool* notifications were assigned the minimum priority level. Those notifications are displayed




Figure 6: The age in hours of the notifications in the snapshots, normalized per device. Outliers omitted.




Figure 7: The notification priority level distribution per category. The priority level is set by the notifying app.

visually muted on recent Android versions. Notifications related to people and events (*Calendar & Reminder*, *Phone*, *SMS & IM*, *Social*) were assigned the high and maximum priority levels more often. More than half of the *Phone* notifications were assigned the maximum priority value. Interestingly, *SMS & IM* notifications were less often assigned the maximum priority level than *Social & Dating* notifications, however, due to *SMS & IM* being assigned in the *People to People* section, they are likely to be ranked higher. A notable exception of people and events related notifications are *Email* notifications, as almost all had the default priority value.

Notification Age. Figure 6 shows the age of the notifications in the drawer when a snapshot was taken. We can see that *Calendar & Reminder*, *Navigation*, *Phone*, and *SMS & IM* notifications tend not to stick around as long as the other categories. This might be either because of users reacting faster on these categories of notifications or because the app is

often updating the notification. While the nature of our data set does not allow us to know the reason, we know from prior work that users tend to attend messaging notifications faster and more often [39, 40].

Non-clearable Notifications. Another reason for some notification sticking around longer than others is that Android notifications can be marked as non-clearable. These notifications cannot be dismissed by users, even when clicking on *Clear All*. Half of the snapshots (51.39%) contained at least one non-clearable notification. We saw a median of one non-clearable notification per snapshot. Of the 10,928,880 unique notifications, 71.27% were clearable and 28.73% were non-clearable. Most of the non-clearable notifications were from the category *System* (35.15%), followed by *Tool* (16.13%), *Media* (10.77%), and *Phone* (9.18%). A typical example are active media playback notifications and notifications about ongoing phone calls.

User Types

So far we mainly looked at the data set in an aggregated manner. However, our earlier results on the mean number of notifications in the notification drawer indicated differences in how users manage notifications. To explore this further, we turned to prior work by Whittaker and Sidner who investigated the management of email inboxes and found three user types [48]. *Frequent Filers* constantly tried to reduce the number of items in their inbox, *Spring Cleaners* made “clean-up” passes in larger intervals of time, and *No Filers* did not make use of filing emails and relied on search instead. Inspired by these user types, we clustered the snapshots according to the mean number of notifications, i.e., [0, 5), [5, 10), and 10+ mean notifications. Within those clusters, we found similar usage patterns regarding the notifications in the drawer over time.




Figure 8: Examples for the three user types. *Frequent Cleaners* try to keep notifications out of the notification drawer. *Notification Regulators* have an increased number of notifications in the drawer but keep the overall number in check. *Notification Hoarders* accumulate notifications and dismiss them all at once, by pressing *Clear All* or restarting the device.

Frequent Cleaner: Figure 8a shows snapshots from 36 days of usage with 24 apps and $M = 0.48$ notifications ($SD = 0.62$, $Md = 0$). 37.67% of the snapshots contained a non-clearable notification. Similar to *Frequent Filers*, *Frequent Cleaners* try to minimize the number of notifications in the drawer. Even with one-third of the snapshots containing a non-clearable notification, the median number of notifications is zero. This also reminds of the “inbox zero” email management approach of keeping the inbox empty [29].

Notification Regulator: Figure 8b shows snapshots covering 39 days and 23 apps with $M = 9.69$ notifications ($SD = 4.19$, $Md = 9$). 99.73% of the snapshots contained a non-clearable notification. *Notification Regulators* have a higher number of notifications in the notification drawer, but they take action before the number gets too high.

Notification Hoarder: Figure 8c shows snapshots covering 34 days of usage with 97 apps and $M = 45.0$ notifications ($SD = 20.01$, $Md = 45$). 99.7% of the snapshots contained a non-clearable notification. This type of user does not seem to dismiss notifications regularly. Instead, they let notifications accumulate and presumably only take action on the notifications that are important to them. In the shown time frame, we can see the number of (grouped) notifications reaching 100 multiple times. We can also see multiple drops were all notifications were cleared, presumably from the user pressing *Clear All* or restarting the device. However, right after the drop, the number of notifications starts to accumulate again. In a recently published work on the importance of notification content, Visuri et al. were surprised by a participant of a pilot study not clearing their notifications [41], a characteristic of *Notification Hoarders*.

Most users we have seen in our data set can be categorized as a *Frequent Cleaner* or *Notification Regulator*. While the number of *Notification Hoarders* is rather small, this behavior seems to be alarming from a notification overload perspective. Grevet et al. suggested a link between high email unread

counts and feelings of disorganization [23], something that future work should investigate for notifications.

Summary

We conducted a large-scale observational study to gain an understanding of notification drawers in-the-wild. By periodically sampling almost four thousand devices, we showed the average number and the positioning of notifications in notification drawers, and the existence of different user types.

5 DISCUSSION

Prior work has mostly focused on the *arrival* of notifications, e.g., by developing models for automatically deferring notifications until breakpoints [19, 20, 34]. While this is an important aspect of notification management, it is not the whole story. Even when notifications are deferred, they eventually end up in the notification drawer. The same is true for “silent” notifications that do not trigger vibrotactile or sound feedback or when the user silenced the device. In the end, the user is presented with an ever-filling list of notifications on the lock screen and notification drawer. This list has somehow to be managed; otherwise, the advantages of providing proactive information are lost.

Notification Management. Ranking the notifications not in chronological order but based on signals already helps the notification drawer management on Android. In recent Android versions, *Major Ongoing* notifications that often require user interaction (ongoing phone calls, media controls) have a secured spot at the top of the list [16]. Messaging notifications, that were shown again and again to be the most important kind of notifications, are hoisted to the top as well. Still, we argue that this can be improved further. Last year we saw first work towards improving the interaction in the notification drawer. Pielot et al. investigated the dismissal behavior of users [39], and Weber et al. explored new interactions by enabling users to snooze notifications,

i.e., temporarily removing and re-triggering them from the notification drawer [43].

Notification Middleware. We see many parallels between notification drawers and email inboxes. Users receive many different kinds of emails and notifications, e.g., personal messaging, reminders, promotions, and spam. However, while it is common to filter, label and categorize emails, notification controls are currently mostly limited to muting and disabling specific apps. Nowadays, using email without a spam filtering middleware is uncommon, and we argue that there is a need for a similar middleware for notifications.

User Types and Digital Well-being. Our data set revealed that most users seem to be able to keep their notifications in check. Most users had between zero and ten pending notifications in the notification drawer. We suggested the two user types *Frequent Cleaners*, who try to keep the notification drawer clean, and *Notification Regulators*, who have an increased number of notifications in the drawer but overall keep them in check. However, we also saw a small set of users “hoarding” notifications. On first sight, it seems like these users have given up managing their notifications. The implications of these user types are not yet known. Future work should investigate different notification management strategies and their effects on the users’ digital well-being. Possible research questions are whether *Notification Hoarders* are feeling more overwhelmed or feel like they are missing more information than the other user types. The opposite could also be hypothesized. Since those users are spending less time managing notifications, they could feel less stressed than the other user types.

The Importance of Messaging. Finally, as shown again and again in prior work, we saw the importance of messaging in the data set. By far most of the unique notifications were of the category *SMS & IM*, they were prominently positioned in the notification drawer, and were quickly attended to, implying a high turnover rate. However, other categories should not be neglected. For future work, we suggest exploring new tools for managing notifications. Notifications could be automatically cleared after a particular time has passed or based on a context change, e.g., for location-based notifications.

Limitations and Future Work

In this work, we focused on Android devices since prior work on mobile notifications primarily used Android devices as well [17, 30, 34–36, 39, 40]. Future work should also consider the other current dominant smartphone operating system iOS. While the notification drawer on iOS is similar to Android, with notifications shown on the lock screen and by swiping down from the top of the screen, the two operating systems differ in important ways. For instance, notifications

in iOS are opt-in and opt-out on Android [7, 47]. iOS also makes heavy use of notification badges on app icons that allow apps to gain the user’s attention more subtly without posting a notification in the notification drawer.

A second limitation is that we did not record the user interaction in-between snapshots. Therefore, users who receive few notifications and users who receive many notifications but act upon them quickly likely have similar characteristics in this data set. Future work should consider this as well.

Open-Source Data Set

We published the data set and Jupyter notebooks for analysis on our project page¹ under the MIT license. We are confident that this will allow the community to further explore the data set and foster future research on mobile notifications.

6 CONCLUSION

A body of work investigated how many and which types of notifications users receive on a daily basis, and which notifications are valued by users [40]. So far, most research on mobile notifications focused on the moment when a notification is triggered. However, the notification drawer as the central place to view and attend notifications has yet to be explored in detail. This is a crucial aspect and required for a complete understanding of mobile notifications.

In this paper, we complemented prior work by reporting the results of a large-scale observational in-the-wild study, in which we sampled the contents of notification drawers. We collected 8,830,112 notification drawer snapshots from 3,953 devices. We systematically analyzed the data set and found users have, on average, 3.4 notifications in the notification drawer. Although users receive significantly fewer notifications at night, notifications accumulate overnight, resulting in more notifications for users to handle in the morning. We found that *SMS & IM* notifications dominate the number one position in the notification drawer and discussed reasons for this. Finally, we suggested the existence of three different types of users regarding the management of notification drawers. *Frequent Cleaners* aim to dismiss all pending notifications in the drawer quickly, *Notification Regulators* receive an increased number of notifications but keep them under control, and *Notification Hoarders* accumulate notifications in the drawer over time and dismiss them all at once. Future work should look further into these user types and investigate the effects of different notification management strategies on users.

¹<https://github.com/interactionlab/android-notification-drawers>

Acknowledgments

This work was supported by the German Research Foundation (DFG) within project C04 of SFB/Transregio 161 and within the SimTech Cluster of Excellence (EXC 310/2).

REFERENCES

- [1] Piotr D. Adamczyk and Brian P. Bailey. 2004. If Not Now, when?: The Effects of Interruption at Different Moments Within Task Execution. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 271–278. <https://doi.org/10.1145/985692.985727>
- [2] Christoph Anderson, Isabel Hübener, Ann-Kathrin Seipp, Sandra Ohly, Klaus David, and Veljko Pejovic. 2018. A Survey of Attention Management Systems in Ubiquitous Computing Environments. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 2, Article 58 (July 2018), 27 pages. <https://doi.org/10.1145/3214261>
- [3] Julie H. Aranda and Safia Baig. 2018. Toward "JOMO": The Joy of Missing out and the Freedom of Disconnecting. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '18)*. ACM, New York, NY, USA, Article 19, 8 pages. <https://doi.org/10.1145/3229434.3229468>
- [4] Jonas Auda, Dominik Weber, Alexandra Voit, and Stefan Schneeeggass. 2018. Understanding User Preferences Towards Rule-based Notification Deferral. In *Extended Abstracts of the 2018 CHI Conference on Human Factors in Computing Systems (CHI EA '18)*. ACM, New York, NY, USA, Article LBW584, 6 pages. <https://doi.org/10.1145/3170427.3188688>
- [5] Daniel Avrahami, Susan R. Fussell, and Scott E. Hudson. 2008. IM Waiting: Timing and Responsiveness in Semi-synchronous Communication. In *Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work (CSCW '08)*. ACM, New York, NY, USA, 285–294. <https://doi.org/10.1145/1460563.1460610>
- [6] Agathe Batestini, Vidya Setlur, and Timothy Sohn. 2010. A Large Scale Study of Text-messaging Use. In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '10)*. ACM, New York, NY, USA, 229–238. <https://doi.org/10.1145/1851600.1851638>
- [7] Frank Bentley. 2018. Losing the Power of Defaults: Exploring Changing Patterns in Enabling Mobile Notifications. In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers (UbiComp '18)*. ACM, New York, NY, USA, 1208–1214. <https://doi.org/10.1145/3267305.3274109>
- [8] Jeremy Birnholtz, Jeff Hancock, Madeline Smith, and Lindsay Reynolds. 2012. Understanding Unavailability in a World of Constant Connection. *Interactions* 19, 5 (Sept. 2012), 32–35. <https://doi.org/10.1145/2334184.2334193>
- [9] Matthias Böhmer, Brent Hecht, Johannes Schöning, Antonio Krüger, and Gernot Bauer. 2011. Falling Asleep with Angry Birds, Facebook and Kindle: A Large Scale Study on Mobile Application Usage. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '11)*. ACM, New York, NY, USA, 47–56. <https://doi.org/10.1145/2037373.2037383>
- [10] Yung-Ju Chang, Yi-Ju Chung, Yi-Hao Shih, Hsiu-Chi Chang, and Tzu-Hao Lin. 2017. What Do Smartphone Users Do when They Sense Phone Notifications?. In *Proceedings of the 2017 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2017 ACM International Symposium on Wearable Computers (UbiComp '17)*. ACM, New York, NY, USA, 904–909. <https://doi.org/10.1145/3123024.3124557>
- [11] Karen Church and Rodrigo de Oliveira. 2013. What's Up with WhatsApp?: Comparing Mobile Instant Messaging Behaviors with Traditional SMS. In *Proceedings of the 15th International Conference on Human-computer Interaction with Mobile Devices and Services (MobileHCI '13)*. ACM, New York, NY, USA, 352–361. <https://doi.org/10.1145/2493190.2493225>
- [12] Evernote Corporation. 2019. Android-Job. <https://github.com/evernote/android-job>. [Online; accessed 20-January-2019].
- [13] Mary Czerwinski, Mary Czerwinski, Eric Horvitz, and Susan Wilhite. 2004. A Diary Study of Task Switching and Interruptions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '04)*. ACM, New York, NY, USA, 175–182. <https://doi.org/10.1145/985692.985715>
- [14] Android Developers. 2019. Android Distribution Dashboard. <https://developer.android.com/about/dashboards/>. [Online; accessed 20-January-2019].
- [15] Android Developers. 2019. NotificationListenerService. <https://developer.android.com/reference/android/service/notification/NotificationListenerService>. [Online; accessed 20-January-2019].
- [16] Android Developers. 2019. Notifications UX: What's New for Android O (Google I/O '17). <https://youtu.be/vwZi56l0Mi0?t=1234>. [Online; accessed 20-January-2019].
- [17] Tilman Dingler and Martin Pielot. 2015. I'll Be There for You: Quantifying Attentiveness Towards Mobile Messaging. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '15)*. ACM, New York, NY, USA, 1–5. <https://doi.org/10.1145/2785830.2785840>
- [18] Tilman Dingler, Dominik Weber, Martin Pielot, Jennifer Cooper, Chung-Cheng Chang, and Niels Henze. 2017. Language Learning On-the-go: Opportune Moments and Design of Mobile Microlearning Sessions. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, USA, Article 28, 12 pages. <https://doi.org/10.1145/3098279.3098565>
- [19] Joel E. Fischer, Chris Greenhalgh, and Steve Benford. 2011. Investigating Episodes of Mobile Phone Activity As Indicators of Opportune Moments to Deliver Notifications. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '11)*. ACM, New York, NY, USA, 181–190. <https://doi.org/10.1145/2037373.2037402>
- [20] Joel E. Fischer, Nick Yee, Victoria Bellotti, Nathan Good, Steve Benford, and Chris Greenhalgh. 2010. Effects of Content and Time of Delivery on Receptivity to Mobile Interruptions. In *Proceedings of the 12th International Conference on Human Computer Interaction with Mobile Devices and Services (MobileHCI '10)*. ACM, New York, NY, USA, 103–112. <https://doi.org/10.1145/1851600.1851620>
- [21] Jose A. Gallud and Ricardo Tesoriero. 2015. Smartphone Notifications: A Study on the Sound to Soundless Tendency. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. ACM, New York, NY, USA, 819–824. <https://doi.org/10.1145/2786567.2793706>
- [22] Tony Gillie and Donald Broadbent. 1989. What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research* 50, 4 (01 Apr 1989), 243–250. <https://doi.org/10.1007/BF00309260>
- [23] Catherine Grevet, David Choi, Debra Kumar, and Eric Gilbert. 2014. Overload is Overloaded: Email in the Age of Gmail. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 793–802. <https://doi.org/10.1145/2556288.2557013>

- [24] Niels Henze and Martin Pielot. 2013. App stores: external validity for mobile HCI. *interactions* 20, 2 (2013), 33–38.
- [25] Niels Henze, Martin Pielot, Benjamin Poppinga, Torben Schinke, and Susanne Boll. 2011. My app is an experiment: Experience from user studies in mobile app stores. *International Journal of Mobile Human Computer Interaction (IJMHCI)* 3, 4 (2011), 71–91.
- [26] Edward Cutrell Mary Czerwinski Eric Horvitz. 2001. Notification, disruption, and memory: Effects of messaging interruptions on memory and performance. In *Human-Computer Interaction: INTERACT*, Vol. 1. 263.
- [27] Yoram M Kalman, Gilad Ravid, Daphne R Raban, and Sheizaf Rafaeli. 2006. Pauses and response latencies: A chronemic analysis of asynchronous CMC. *Journal of Computer-Mediated Communication* 12, 1 (2006), 1–23.
- [28] Uichin Lee, Joonwon Lee, Minsam Ko, Changhun Lee, Yuhwan Kim, Subin Yang, Koji Yatani, Gahgene Gweon, Kyong-Mee Chung, and Junehwa Song. 2014. Hooked on Smartphones: An Exploratory Study on Smartphone Overuse Among College Students. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 2327–2336. <https://doi.org/10.1145/2556288.2557366>
- [29] Merlin Mann. 2019. Inbox zero. http://sagolla.net/media/jpg/merlinmann_adobetalk_2007-03.pdf. [Online; accessed 20-January-2019].
- [30] Abhinav Mehrotra, Robert Hendley, and Mirco Musolesi. 2016. PrefMiner: Mining User’s Preferences for Intelligent Mobile Notification Management. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 1223–1234. <https://doi.org/10.1145/2971648.2971747>
- [31] Abhinav Mehrotra, Mirco Musolesi, Robert Hendley, and Veljko Pejovic. 2015. Designing Content-driven Intelligent Notification Mechanisms for Mobile Applications. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '15)*. ACM, New York, NY, USA, 813–824. <https://doi.org/10.1145/2750858.2807544>
- [32] Abhinav Mehrotra, Veljko Pejovic, Jo Vermeulen, Robert Hendley, and Mirco Musolesi. 2016. My Phone and Me: Understanding People’s Receptivity to Mobile Notifications. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 1021–1032. <https://doi.org/10.1145/2858036.2858566>
- [33] Bonnie A. Nardi, Steve Whittaker, and Erin Bradner. 2000. Interaction and Outeraction: Instant Messaging in Action. In *Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (CSCW '00)*. ACM, New York, NY, USA, 79–88. <https://doi.org/10.1145/358916.358975>
- [34] Tadashi Okoshi, Julian Ramos, Hiroki Nozaki, Jin Nakazawa, Anind K Dey, and Hideyuki Tokuda. 2015. Attelia: Reducing user’s cognitive load due to interruptive notifications on smart phones. In *2015 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 96–104.
- [35] Martin Pielot, Karen Church, and Rodrigo de Oliveira. 2014. An In-situ Study of Mobile Phone Notifications. In *Proceedings of the 16th International Conference on Human-computer Interaction with Mobile Devices & Services (MobileHCI '14)*. ACM, New York, NY, USA, 233–242. <https://doi.org/10.1145/2628363.2628364>
- [36] Martin Pielot, Rodrigo de Oliveira, Haewoon Kwak, and Nuria Oliver. 2014. Didn’t You See My Message?: Predicting Attentiveness to Mobile Instant Messages. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3319–3328. <https://doi.org/10.1145/2556288.2556973>
- [37] Martin Pielot, Niels Henze, and Susanne Boll. 2011. Experiments in app stores-how to ask users for their consent. In *Proc. workshop on Ethics, logs & videotape*.
- [38] Martin Pielot and Luz Rello. 2017. Productive, Anxious, Lonely: 24 Hours Without Push Notifications. In *Proceedings of the 19th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '17)*. ACM, New York, NY, USA, Article 11, 11 pages. <https://doi.org/10.1145/3098279.3098526>
- [39] Martin Pielot, Amalia Vradi, and Souneil Park. 2018. Dismissed!: A Detailed Exploration of How Mobile Phone Users Handle Push Notifications. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '18)*. ACM, New York, NY, USA, Article 3, 11 pages. <https://doi.org/10.1145/3229434.3229445>
- [40] Alireza Sahami Shirazi, Niels Henze, Tilman Dingler, Martin Pielot, Dominik Weber, and Albrecht Schmidt. 2014. Large-scale Assessment of Mobile Notifications. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 3055–3064. <https://doi.org/10.1145/2556288.2557189>
- [41] Aku Visuri, Niels van Berkel, Tadashi Okoshi, Jorge Goncalves, and Vassilis Kostakos. 2019. Understanding Smartphone Notifications’ User Interactions and Content Importance. *International Journal of Human-Computer Studies* (2019).
- [42] Dominik Weber, Alireza Sahami Shirazi, and Niels Henze. 2015. Towards Smart Notifications Using Research in the Large. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. ACM, New York, NY, USA, 1117–1122. <https://doi.org/10.1145/2786567.2794334>
- [43] Dominik Weber, Alexandra Voit, Jonas Auda, Stefan Schneegass, and Niels Henze. 2018. Snooze! Investigating the User-defined Deferal of Mobile Notifications. In *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services (MobileHCI '18)*. ACM, New York, NY, USA, Article 2, 13 pages. <https://doi.org/10.1145/3229434.3229436>
- [44] Dominik Weber, Alexandra Voit, and Niels Henze. 2018. Notification Log: An Open-Source Framework for Notification Research on Mobile Devices. In *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers (UbiComp '18)*. ACM, New York, NY, USA, 1271–1278. <https://doi.org/10.1145/3267305.3274118>
- [45] Dominik Weber, Alexandra Voit, Philipp Kratzer, and Niels Henze. 2016. In-situ Investigation of Notifications in Multi-device Environments. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '16)*. ACM, New York, NY, USA, 1259–1264. <https://doi.org/10.1145/2971648.2971732>
- [46] Dominik Weber, Alexandra Voit, Huy Viet Le, and Niels Henze. 2016. Notification Dashboard: Enabling Reflection on Mobile Notifications. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '16)*. ACM, New York, NY, USA, 936–941. <https://doi.org/10.1145/2957265.2962660>
- [47] Tilo Westermann, Sebastian Möller, and Ina Wechsung. 2015. Assessing the Relationship Between Technical Affinity, Stress and Notifications on Smartphones. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. ACM, New York, NY, USA, 652–659. <https://doi.org/10.1145/2786567.2793684>
- [48] Steve Whittaker and Candace Sidner. 1996. Email Overload: Exploring Personal Information Management of Email. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '96)*. ACM, New York, NY, USA, 276–283. <https://doi.org/10.1145/238386.238530>