

An In-Situ Study of Mobile Phone Notifications

Martin Pielot
Telefonica Research
Barcelona, Spain
martin.pielot@telefonica.com

Karen Church^{*}
Telefonica Research
Barcelona, Spain
kchurch@yahoo-inc.com

Rodrigo de Oliveira[†]
Telefonica Research
Barcelona, Spain
oliveirar@google.com

ABSTRACT

Notifications on mobile phones alert users about new messages, emails, social network updates, and other events. However, little is understood about the nature and effect of such notifications on the daily lives of mobile users. We report from a one-week, in-situ study involving 15 mobile phones users, where we collected real-world notifications through a smartphone logging application alongside subjective perceptions of those notifications through an online diary. We found that our participants had to deal with 63.5 notifications on average per day, mostly from messengers and email. Whether the phone is in silent mode or not, notifications were typically viewed within minutes. Social pressure in personal communication was amongst the main reasons given. While an increasing number of notifications was associated with an increase in negative emotions, receiving more messages and social network updates also made our participants feel more connected with others. Our findings imply that avoiding interruptions from notifications may be viable for professional communication, while in personal communication, approaches should focus on managing expectations.

Author Keywords

Notifications; Asynchronous Communication; Interruptions; Mobile Phones; Empirical Study

ACM Classification Keywords

H.5.2 Information interfaces and presentation: Miscellaneous; H.4.3 Communications Applications: Other.

INTRODUCTION

Smartphones have evolved to become a constant, pervasive companion to almost 1 billion people across the globe. Applications like messaging, email, and online social networks, allow smartphone users to connect to their family, friends and even co-workers wherever they go. Most of these communication apps are asynchronous, i.e. the receiving person is not expected or required to respond to a message/email/social network update right-away. However, the majority of these

^{*}The author is currently affiliated with Yahoo Inc., USA.

[†]The author is currently affiliated with Google Inc., USA.

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 ACM 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.
MobileHCI '14, September 23 - 26 2014, Toronto, ON, Canada
Copyright © 2014 ACM 978-1-4503-3004-6/14/09...\$15.00.
<http://dx.doi.org/10.1145/2628363.2628364>




Figure 1. Notification drawer on Android OS, which can be opened by pulling down the notification area. Here, users can see pending notifications and preview sender and content.

apps use proactive, push-based notifications, i.e., visual, auditory, and/or haptic alerts, to inform mobile users about new, unattended messages or events, even when the user is not actively using the application in question. On arrival, these notifications alert users via sounds, vibrations, icons, and badges on the app's icon.

This forms a stark contrast to the *quiet and invisible servants* that Weiser foresaw in his vision of ubiquitous, calm technology [31]. Push notifications have the ability to turn our mobile and ubiquitous communication devices into loud, disruptive machines. And the increased volume of communication exchanges among smartphone users highlights just how disruptive smartphones have become. For example, WhatsApp, a free messaging service that uses the internet to relay SMS-like exchanges, currently handles over 50 billion messages per day¹. Finally, at the end of 2012, more users connected to Facebook every day via their mobile phones than via the desktop². Given that notifications are at the core of many of

¹On Dec 31st [2013] [WhatsApp] users sent 18B msgs and received 36B = 54 Billion total messages in a day. See <https://twitter.com/WhatsApp/status/420373902980689920>

²Facebook Now TRULY a Mobile Company, Beats Expectations, see <http://www.cnbc.com/id/100420893>

these applications, it's likely that mobile users have to deal with dozens and even hundreds of notifications per day.

While the effect of notifications in desktop environments, particularly in work settings, has been studied thoroughly [8, 9, 16, 18, 20, 22, 32], little is understood about the nature and effect of *mobile notifications on everyday life*. Since we carry our phone with us throughout the day, mobile notifications continually cross the boundaries of work and private life and as such have the potential to interrupt us in a wider range of situations and contexts. To shed light on what notifications phone users have to cope with and how these affect them, we conducted a one-week, in-situ study involving 15 mobile phone users, where we collected objective and subjective data about mobile notifications. The contribution of this work is:

- A quantitative analysis of the notifications that our participants receive and how they dealt with them,
- the fusion of this data with daily subjective ratings and stories describing events around notifications, and
- a discussion of previously explored approaches to improve notifications and asynchronous communication.

RELATED WORK

Notifications and their Interruptive Nature

Iqbal and Bailey [19] define a *notification* as a visual cue, auditory signal, or haptic alert generated by an application or service that relays information to a user outside her current focus of attention. On mobile phones, notifications are typically delivered instantly, *e.g.*, when the user receives a call or a message. In general, they only arrive when a corresponding application is closed, *e.g.*, when the user has an email application open, no notification will be generated by the OS if a new email arrives.

To date, the majority of research on notifications has focused on information workers in desktop environments. This previous work has repeatedly highlighted the negative effects of notifications on work efficiency. On the basis of the results of a study with 11 information workers, Czerwinski *et al.* [10] conclude that people find it difficult to return to a previous task after having been interrupted by *e.g.*, instant messages, calls, or an engagement with a colleague. Cutrell *et al.* [8] found that the effect is more pronounced when the task is cognitively demanding. Leiva *et al.* [21] found that phone calls interrupting the use of an app, though occurring rarely, significantly increase the time a user spends completing the initial task. De Vries *et al.* [11] showed that, depending on the mental workload, the level of politeness of a notification affects how annoyed and disrupted users feel. In their study, the participants preferred polite notifications during low-workload tasks and neutral notifications during high-workload tasks.

Mark *et al.* [22] conducted a deprivation study where they completely cut off email access to 13 information workers for 5 work days. Their results show that without email, the workers multitasked less, spent more consecutive time on tasks, and had decreased levels of stress. Similarly, Iqbal *et al.*

[20] found that disabling email notifications leads to less frequent opportunistic email checking. However, while participants were aware of the disruptive effects, they appreciated the awareness provided by the notifications [20].

In mobile phones, notifications play an even more central role to “notify” users of new messages, events or actions. Previous work has shown that the perception towards mobile notifications varies strongly. If apps which are not perceived as useful keep sending notifications, users become annoyed and consider deleting those apps [12]. According to a field study with 11 co-workers by Fischer *et al.* [14], the user's receptiveness is determined by message content, *i.e.* how interesting, entertaining, relevant, and actionable a message is. The time of delivery, in contrast, did not affect receptivity. Sahami *et al.* [28] conducted a large-scale study with more than 40,000 Android phone users, and found via experience sampling that notifications from communication applications the most important.

Computer-Mediated Communication and Social Pressure

Communication via such channels has become essential for young people in particular [3, 29]. Skierkowski and Wood [29] report that when they restricted messaging as part of a user study, the participants not only showed increased anxiety, but many also did not comply. Communication applications, such as messengers, email, or social networks, rely on notifications to draw its users' attention to new messages or content. Holtgraves *et al.* [17] describes the unspoken rule of immediacy and pressure felt by people to carry their personal mobile devices at all times. Notifications are an essential part to supporting immediacy, as otherwise people would have to check their phone frequently. And despite the availability of notifications, Oulasvirta *et al.* [24] found that mobile phones force user's to adopt the habit of frequently checking for new messages and notifications, which is perceived as annoying rather than useful. Birnholtz *et al.* [5] found that one way to counter this social pressure is through “butler lies”, *i.e.*, pretending that a message wasn't seen on time.

Mediating Interruptions and Expectations

To counter the disruptive effects on notifications, previous work has investigated to delay their delivery until a suitable point. In a study with 16 graduate students, Adameczyk *et al.* [1] found that their participants felt higher workloads when notifications were delivered while they were in the middle of a work task, such as correcting or writing text, or conducting a web search. Horvitz *et al.* [18] proposed deferring email notifications by 1 to 5 minutes to minimize their disruptive nature. Iqbal and Bailey [19] showed that delivering emails at so-called *breakpoints*, *i.e.* events when a person has just mentally finished a task, reduces frustration and makes users react to them faster.

Another approach in the case of communication is to shape expectations by estimating and/or communicating the availability of the recipient. Nardi *et al.* [23] explored the use of desktop instant messengers in work settings and highlight that when starting a conversation, time and topic are convenient for the initiator, but not necessarily for the recipient.

Fogarty *et al.* [15] conducted a series of studies that show that in an office environment, simple, low-costs sensor, such as a microphone, can infer the interruptibility of an information worker with the same accuracy as humans. Avrahami and Hudson [2] showed that it is possible to predict how fast a receiver of a message is going to respond by utilizing the *status* of the user and general activity events, such as key presses or mouse clicks. While Teevan and Hehmeyer [30] found that for desktop instant messaging, “busy” workers are more likely to accept a call, because they believe that the call is important.

Research in this regard has also been conducted in the mobile space. For example, Ho and Intille [16] explored delivering notifications at the transition between physical activities. Their results suggest that notifications might be considered more positively when being delivered between two physical activities, such as sitting and walking. Fischer *et al.* [13] conducted an experience-sampling study, where they created notifications that asked users to provide feedback about their context, such as taking a picture of the current activity. The participants reacted faster to notifications, when they are delivered right after finishing a phone call or reading a text message. However, Fischer *et al.* could not show whether the subjective experience improves, presumably because the artificial notifications created extra work. Rosenthal *et al.* [27] contributed to a method to learn, in a personalized way, when to mute a phone on a call or a notification to avoid embarrassing interruptions.

In summary, most of prior work has focused on desktop environments and office settings. While some past work has studied notifications in the context of mobile devices, none of those studies investigated *real-world* notifications in natural, un-biased settings. Given that mobile phone notifications are often related to communication, and thus highly personal, what is missing from the literature are insights on the basis of unobtrusive, in-situ data collection, which explore how mobile phone users deal with the fact that notifications now reach them everywhere and all the time.

METHOD

In order to study the nature and effect of notifications on the lives of mobile phone users, we conducted a one-week, in-situ field study. Given the goal of this work was to investigate mobile notifications in natural settings, we needed to employ a study methodology that would (1) avoid generating extra notifications that would potentially bias the nature of the phenomenon being studied — as would have been the case with experience sampling, and (2) would enable us to gather *actual* notification data in real-world contexts and in an unobtrusive manner. Thus, we used a mixed-method approach, where we collected both objective and subjective data.

For the objective data, we collected and analyzed real-world logs of all notifications that our participants received during the study period and how fast the participants viewed them after they arrived to their phone. For the subjective data, we conducted an online-diary study, where our participants reflected on the notifications they received each day, and they

provided insights into their perceptions regarding those notifications.

Participants

For security reasons, the ability to monitor notifications is restricted on most popular mobile platforms. At the time the study was conducted, the only way was to use Android OS and intercept notifications through its accessibility API. Hence, we recruited Android participants via announcements on social networks and community forums. In total, 15 English-speaking participants (6 female, 9 male) aged 24-43 ($Mdn = 28$, $M = 30.46$, $SD = 6.04$) successfully completed all aspects of the study³. All of the participants lived in either North America or Europe, and their professions included engineering, teaching, graduate and undergraduate students. In terms of incentives, participants were rewarded with a 40 EUR Amazon Gift Card for taking part. In addition, we raffled three 100 EUR gift cards amongst them.

All participants indicated that on average they use email and *WhatsApp* several times per day, *Google Talk* once per day, *SMS* several times per week, and *Facebook Messenger* about once per week. Other messengers like *Line*, *TuMe*, or *KakaoTalk* were only reported by few participants, or were used infrequently. When asked to gauge how fast they typically respond to messages, half of the participants indicated that they generally respond to notifications within a few minutes, while the other half indicated that they tend to respond within an hour. Participants also estimated that others expect them to respond within similar time frames, i.e. half within a few minutes, half within an hour.

Procedure

The study was conducted in February and March 2013. Prior to the beginning of the study, the participants installed our logging app, the *Notification Monitor*, distributed as .apk, on their mobile phone. This Android app collected data regarding the notifications received along with details of when and how participants attended to those notifications. The *Notification Monitor* app ran as a background service on the user’s phone. Once installed, the participants were asked to send us their participation ID, which was automatically generated by the app on their behalf. From the point of installation, the notification logging began and continued until the end of the study. The main phase of study commenced once the application was running on all participants’ phones.

This main phase ran for one week, beginning on a Monday and ending on a Sunday. On each day during this period, participants were asked to provide subjective feedback regarding their perceptions and experiences with notifications through the daily diary which was delivered in form of a survey. Every morning, we sent an email to participants with a link to the survey asking them to reflect upon the notifications received on the previous day.

Given past work has shown that many mobile app interactions, in particular with social networking apps [6] and text

³In total, 27 participants started the project, but due to our strict compliance requirements, we discarded the data of 12 of them.

messaging [3], occur most often in the late evening / at nighttime (between 7pm - 1am), we chose to prompt users to fill out the diary every morning of the following day, so that participants could reflect on the whole past day. This conscious design choice turned out to be an important as it allowed us capture periods of communication that take place late at night. To not bias the participants, no statistics of the data collected on previous day was presented. Participants were informed that in order to be eligible for the study incentives, they were required to fill out all 7 diary entries.

The diary included a mix of open-ended and closed questions, of which there were 9 in total, focusing on the 3 key areas:

- *Subjective perceptions* of the volume of received notifications and the time to react to them by category.
- *Attitudes and emotions* felt towards all notifications.
- *Personal stories*, collected via open-text fields, highlighting actual cases encountered by the participant where notifications kept them from doing something they had to, caused feelings of pressure to respond and/or induced feelings of wanting to respond to a notification faster than was possible.

Finally, at the end of the study, participants were debriefed and provided with their compensation. In the following sections, we describe the notification logging and diary portion of the study in more detail.

Quantitative Measures

To understand the quantitative measures, it's important to understand how notifications are implemented in Android OS⁴. When a notification arrives, the phone, depending on its mode, creates a buzz and an audible notification sound. At the same time, a little icon appears in the top left part of the phone screen, the so-called *notification area*. To see more details about the notification, the user can pull down this area and extend it into the *notification drawer*, which is shown in Figure 1 on page 1. In this notification drawer, users are provided with more details about the notifications. For short messages, the whole message can be read there. For longer messages, such as email, the user can read the subject line.

We tracked two events: the reception of a notification, i.e. *Notification Received* and the time between receiving and viewing of a notification, i.e. *Notification Viewed*.

Notification Received

On Android OS, notifications are generated by a number of applications which cannot be considered *true notifications* in the sense of the definition by Iqbal and Bailey [20]. For example, a weather widget that shows the current temperature in the notification area will constantly generate notification events, however, there will be no alert to the user whenever the temperature is updated. To overcome this issue, we created a white-list of applications that deliver true push notifications in the sense of Iqbal and Bailey. We then grouped these

notifications into four categories: *Messengers*, (e.g. WhatsApp, SMS, Line), *Emails* (any email client), *Social network updates* (e.g. Twitter, Facebook, Google+), and *Other* notifications (e.g. system messages, calendar reminders, application updates). For each of these notifications, we store the local time of the notification arrival and the package name of the corresponding application (e.g., com.whatsapp), which serves as a unique identifier for apps on the Google Play Store.

Notification Viewed

There are two ways of viewing notifications. When opening the notification drawer, we consider all pending notifications shown in the drawer as viewed. When the user opens an application, we consider all pending notifications associated with this application as viewed. For the scope of this study, *viewing* is the key event, as in this moment users have clearly become aware of the notification. For each of these *view events*, we store the name of the notification (e.g., com.whatsapp) and the local time at which the message was viewed. How users actually react to notifications, i.e., whether they respond, dismiss it, or ignore it, is out of the scope of this work.

RESULTS

Quantitative Results from Logs

In total, our Notification Monitor application logged 15 person weeks (105 days) of real mobile notification events which equates to 75 weekdays and 30 weekend days. Due to the non-normal distribution of some measurements, we report median rather than mean values (unless otherwise noted) and use non-parametric methods for inferential statistics. Therefore, null-hypothesis testing was performed using the Friedman test and the Wilcoxon rank sum test for matched pairs. Correlations were computed using the Spearman's Rho coefficient and pairwise deletion.

Basic Descriptive Statistics

In the first step, we cleaned the data set by removing instances of misused notifications. Examples are, the Android keyboard that creates a notification event whenever opened, the app market that uses a notification icon to present the percentage of downloaded data – generating 100 events in the process of each update, or a custom battery indicator which used a notification icon to display the battery level. We removed notification events from 64 applications in the process.

The remaining 33 applications were clustered into messengers, email, social networks, and other, as described above. By applying these filters, we ended up with 6854 notifications as defined by Iqbal and Bailey [19]. Thus, each participants received a mean number of 65.3 notifications per day. Table 1 shows the number of notifications received per category.

Cat.	#	Top apps
Msger	3340	WhatsApp, Google Talk, SMS
Email	2210	GMail, K-9 Mail, Android Email
Social	269	TweakDeck, Facebook, Google Plus
Other	1035	Updates, Calendar, Pea Pod

Table 1. Notifications by category.

⁴See <http://developer.android.com/guide/topics/ui/notifiers/notifications.html> for the official description

Since we wanted to study the effect of the notifications that people *actually* receive, we did not require the participants to turn on notifications they typically have disabled. All of the participants received notifications from messengers and other apps. All but two received emails notifications, and 8 received notifications from social networking apps.

5374 (78.4%) of the notifications arrived during a weekday and 1480 (21.6 %) during the weekend, which means that the number of notifications dropped during the weekend. Thus, we split our analysis by weekday and day in weekend to determine if more interesting differences emerge.

Figure 2 shows the average number of notifications received per day for each category, split by weekday versus weekend day. The number of notifications received per day differed significantly by category of notification ($\chi^2(3) = 44.64, p < .001$). Pairwise comparisons using Bonferroni-corrected Mann-Whitney tests showed that our participants received significantly more messages and emails than social network updates and other notifications (all differences significant at $p < .001$).




Figure 2. Median number of notifications received per participant by day, split by weekday versus weekend day

Temporal Patterns

Figure 3 shows at what (local) time notifications arrived. In general, the number of notifications dropped during the night. Most emails arrived during working hours, which two peaks before and after lunch. Messages, in contrast, filled the gaps and arrived primarily during lunchtime and after working hours.

Mostly Viewed Through Notification Drawer

4837 (70.6%) of notifications were first viewed through the notification drawer, the remaining 1966 (28.7%) by directly opening the app. When viewing the notification in the drawer, the participants continued to the app in 1988 (41.1%) of cases, and they did not in 2849 (58.9%) of cases. When any notification was viewed, either through the notification drawer or




Figure 3. Notifications by hour of the day

through the app, on average there were other un-viewed notifications ($Mdn = 1, M = 3.8, SD = 6.0$). This means that our participants typically did not let notifications accumulate.

Notification Viewed Times

Figure 4 shows the average time between the arrival of a notification until it was viewed, for each of the categories of notifications, split by day of the week and day of the weekend. The median time ranges from 3.5 min for messengers on weekends to 27.7 min for email on weekends. The fastest attended notifications were those generated from messengers (6.6 and 3.5 min) and social network applications (3.8 and 7.0 min / weekday and weekend day, respectively).

The box plots indicates that half of the notifications were viewed within a few minutes, and that the majority were attended to within hour (email on weekends being the exception), which confirms the subjective perception of the participants from the recruitment survey. When notifications arrived, the screen had been off in 69.2% of the cases. Given how fast people attended to notifications, this indicates that notifications often triggered interaction with the phone.

The speed at which notifications were viewed on average during the day differed significantly by type ($\chi^2(3) = 9.13, p < .05$). Pairwise comparisons using Bonferroni-corrected Mann-Whitney tests showed that our participants viewed messages significantly faster than emails ($p = .025$).

The Effect of Ringer Mode

One element of mobile phone technology that may or may not effect the interaction with and perceptions of mobile notifications relates to the ringer mode of the mobile phone. Android phones have three different ringer modes: normal mode (sound and vibration), vibration-only, and silent mode (neither sound nor vibration). It might be assumed that users turn the phone to silent mode if they want to avoid interruptions. Of the recorded notifications, 46.2% were received in normal mode, 41.5% in vibration-only mode, and 12.2% in silent mode. This indicates that people frequently disable sound, but rarely disable all alerts. Plotting the distribution, we could not identify any time-related patterns, such as that non-sound modes are prevalent during working hours or night time.

We compared the median time until viewing notifications between the ringer modes and found a significant effect




Figure 4. Median delay between receiving and viewing notifications, split by weekday versus weekend day.

($F(2) = 7.6609, p < .001$). Notifications were viewed significantly faster when the phone was in vibration mode compared to silent mode ($p < .001$) and normal mode ($p < .001$). Notably, the data does not contain any evidence that putting the phone into silent mode – the absence of audio-tactile alerts – leads to slower response times. Thus, it appears that having the phone in silent mode does not appear to help escape the effects of notifications.

Qualitative Results from Diary

In total we collected 97 diary entries over the 1-week period with an average of 6.5 diary entries per participant (Min=4, Max=7, SD=1.1). All diary responses were manually analyzed by two of the authors in an iterative manner to extract emergent themes. This manual analysis took a number of rounds to identify and cluster reoccurring themes.

In order to get further insights into how notifications affected the emotional states of our participants, where possible, we correlated the objective log data from each day with the subjective diary responses for the same day. In our notation, ρ_{obj} indicates correlations with the actual number of notifications received and ρ_{sub} the subjective notification volume, as shown in Figure 5.

Number of Notifications

While in the majority of cases, participants felt as though they received their “usual” amount of notifications per day, we did find differences in their perceptions depending on the type of notification (Figure 5). For example, for messengers and social networks, a significant proportion of participants felt that they received fewer or way fewer notifications than usual over the course of our study (30.5% and 23.7% respectively). Thus, although the number of notifications logged in our study was high (6854 notifications in total, 65.3 notifications per user per day on average), our participants didn’t feel as though they received *more* or *way more* notifications than usual. This implies that participants general perception

of their “usual” number of daily notifications is generally very high.




Figure 5. Subjective perceptions of the number of notifications received compared with the “usual”.

Perceptions of Interruptions

One of the questions explicitly asked in the diary related was how much the majority of their daily notifications interrupted the participants. Based on the responses to that question, we find that mobile messenger applications and emails are the two primary sources of perceptions of interruptions. Participants reported that these apps interrupted them (occasional, frequently, or very frequently) in almost 50% of cases (Figure 6). Taking a closer look at the personal stories surrounding the apps or notifications that cause the most distractions we find that email was *explicitly mentioned* more often than messaging apps (50% vs. 28%).

The amount of emails received during a day was correlated with increased self-reports of negative emotions. Both, subjective and objective email count, lead to higher feelings of being stressed ($\rho_{obj} = 0.356$ and $\rho_{sub} = 0.238$), interrupted ($\rho_{obj} = 0.499$ and $\rho_{sub} = 0.315$), and annoyed (0.412 and $\rho_{sub} = 0.213$). The subjective amount of emails correlated with the feeling of having to deal with a lot of notifications ($\rho_{sub} = 0.374$) and with being overwhelmed ($\rho_{sub} = 0.290$). Additionally, we found a significant correlation between the actual volume of received emails and the number of reported instances where participants felt as though they were kept from doing something ($\rho_{obj} = 0.340$) and pressure to respond ($\rho_{obj} = 0.398$). According to the diary responses, email was mostly associated with work and their impact was generally considered to be interrupting due to their volume, their nature and their unclear urgency, e.g.:

P08: *Email notifications are usually from exchange server, so they are work related. It means if I am not near my computer, I should get to my computer and reply.*

P02: *Email notifications seem to be the most distracting consistently. There are more of them and they are longer.*

P04: *Email notification is most disturbing because the urgency often varies widely.*

Having objectively and subjectively received more notifications from messenger apps during a day correlated with a stronger feeling of having to deal with a lot of messages




Figure 6. Subjective perceptions as to how much the majority of these notifications interrupted from daily tasks

($\rho_{obj} = 0.264$ and $\rho_{sub} = 0.464$) and being overwhelmed ($\rho_{obj} = 0.241$ and $\rho_{sub} = 0.241$). In the diary responses, participants often mentioned messenger apps in this context when being focused on work tasks or when they were busy engaging in other social / leisure activities. In most cases there was an underlying feeling of social pressure, in particular for applications like WhatsApp:

P06: Yesterday, I was working from home and had my phone next to me all day, so I saw all notifications as they came in. Therefore, I was frequently distracted... specially with WhatsApp notifications.

P01: Also WhatsApp was distracting a bit in the evening when I was out at the opera and then for birthday drinks with friends. It was not important to answer, but still a matter of politeness I had to answer in a reasonable time span.

Social network updates also had an impact on emotional states. An increasing subjective volume of social network updates correlated with increased feelings of having to deal with a lot of notifications ($\rho_{sub} = 0.387$), feeling stressed ($\rho_{sub} = 0.224$), interrupted ($\rho_{sub} = 0.221$), and annoyed ($\rho_{sub} = 0.218$). One of the key themes related to social network notifications that emerged from our study is the fact that while social networks are designed to connect people to one another, the many of the notifications generated by social networks are neither directed at the receiver nor are they relevant to the receiver:

P14: Facebook distract me most, because it notifies all the comments that people made in some trending, so a lot of this messages are not for me.

While receiving too many emails and social network updates were related to stress or annoyance, we did not find any of such evidence for messengers. In contrast we found that receiving more messages is significantly correlated with increased feelings of being connected with others ($\rho_{obj} = 0.291$). The same positive correlation was found for the amount of social network updates and feeling connected to others ($\rho_{obj} = 0.205$). Overall this implies that while these services are at times interrupting and overwhelming, the fact

that they help users to keep in touch with friends, family and others appears to also have positive effects.

Perceptions of Response Time

Our log data revealed that most notifications are viewed within a few minutes of reception, and both messenger and social app notifications in particular were viewed very quickly (3.5min - 7min depending on day of week and app type). When we asked our participants about how long it took them to *actually respond to the majority of these notifications*, we found a different pattern. While our participants perceived that they responded to around 30% of messenger notifications within a few minutes and almost 62% of messenger notifications within 1 hour, the perception of responding to email and social app notifications was notably lower (around 30% within an hour), see Figure 7.

Looking to the diary responses we find that social pressure plays a role in how fast users respond to notifications from messenger apps. Specifically we found 3 predominant sources for this pressure, all of which are in line with [7]. First, related to the expectations of others:

P09: WhatsApp, as users don't have option to stop other side of users knowing if you have already read or not [...] people tend to expect immediate answers.

Second, relates to exchanging important, time-critical information, in particular for coordination activities:

P06: I was pending to meet with someone, and was cautious to receiving any WhatsApp from him and be able to reply fast.

P04: I was trying to find out if I should pick up a friend on my way to a meeting.

And finally, the relationship with the receiver seemed to play a role. Given instant messages were mainly used to communicate with loved ones, family, and friends, they seem to generate an increased expectation for faster responses:

P05: In this case, [it] was my girlfriend who asked for some immediate stuff and I should better answer fast if I want to keep her stress less :D

P13: I was talking to my sister. She was sad [and] I wanted to respond immediately to comfort her.

For email and social network updates, in contrast, our data shows that participants deferred attending to these notifications when they started feeling overwhelmed. The more often participants reported that notifications kept them from doing something, the slower they attended to notifications from email ($\rho = 0.255$) and social network updates ($\rho = 0.376$). Participants reported deferring viewing notifications because of other priorities or because of the increased the effort it takes to view and respond to these notifications compared to other communications apps. For example,

P08: Sometimes I take time to respond, because I see it is not urgent and there are other priorities at the moment, so it is my free choice.




Figure 7. Subjective perceptions of the time it took to respond to the “majority” of their notification.

P02: I receive more emails than any other type of notification; and these take more effort to check because you cannot read the message without fully opening the email client. They take longer to read and also require more effort to sort or delete when compared to other messages.

DISCUSSION

Our participants had to deal with a daily mean number of 65 notifications, mostly from communication apps. The majority of the notifications were generated by messengers, such as SMS or WhatsApp, as well as email applications. This implies that mobile notifications are inherently social. When asked about their perceptions of the volume of notifications in the diary, we found that on average our participants largely indicated to have received the “usual” amount of notifications. This implies that our study captured a representative time period of mobile usage and that our participants perception of their “usual” number of daily notifications is generally very high. We found that email notifications were mostly associated with work and were intuitively received more often during work hours, while messenger notifications were mostly associated with personal life and as a result were received during downtimes (lunch breaks) and after work hours.

Notifications were viewed within a few minutes of their reception, with messengers and social network notifications being the fastest attended (3.5-6.6 min and 3.8 -7.0 min respectively), which is inline with the 6 minutes reported by Battestini *et al.* [3] regarding SMS. Notifications were largely pre-screening in the notification drawer. In the diary, the participants indicated that, apart from notifications from the “other” category, they attend to the majority of notifications within an hour. Thus, objectively, participants reacted even faster than subjectively perceived. Overall the fast view times can be attributed to factors like expectations of a response from others, time-critical communication, and the relationship between the sender/receiver (with close-knit relationships appearing to lead to the more feelings of social pressure).

Our analysis revealed a significant relationship between the number of notifications and emotional effects. We found that emails and social networks were correlated with negative emotional responses, in particular to feelings of being

overwhelmed, stressed, interrupted and annoyed. Furthermore, when receiving more emails, participants were also more likely to report experiences where notifications kept them from doing something else or when they felt pressure to respond faster than they were able to. On the other hand, messengers, despite their equally high volume, and social networks to a lesser extent, caused more positive emotional responses. For example we found that receiving more messages is significantly correlated with increased feelings of being connected with others. It’s likely that this again relates to the personal nature of messaging apps.

Our participants could have turned the ringer mode of their phone off, i.e. set their phone to silent mode to escape notifications. However we found that just 12.2% of the recorded notifications were received when the phone was in silent mode. Interestingly, setting the phone to silent mode did not have an effect on how fast our participants viewed the notification. That is, we did not find any evidence that people viewed notifications slower when the phone was in silent mode while they arrived. This could be an indicator that our participants, despite enabling silent mode, still frequently check their phone [24], adhering to the “rule” of immediacy [17]. Thus putting the phone to silent mode does not help to escape the effects of notifications.

In the diary, we found several reports where our participants indicated that they postponed attending to their notifications. Increased feelings of having to deal with a lot of messages correlated with longer delays until emails and social network updates were viewed. Thus, rather than these longer view times being associated with the phone being in silent mode, it appears that at times people make an explicit decision to delay *dealing with notifications*.

During 2013, the number of messages that WhatsApp handled rose from 10B to 56B. Furthermore, younger populations than the sample we studied, particularly teens, have been found to send and receive even more messages per day. Our work represents a snapshot of the situation in March 2013. Given that the volume of communications done via mobile phones is on the rise, as well as the fact that new mobile apps are continually released which rely on notifications, the effects we found in our study might be more pronounced in the future and amongst other user groups.

The presented study has two limitations. First, the data reported in this paper reflects 15 Android users across 7 days. While this is inline with related landmark in-situ studies, which studied between 11 to 20 [10, 13] subjects for 5 to 14 days [22, 20], our findings may not be generalizable to other populations, such as teens or older adults. Initially, we had aimed for more participants, but many volunteers dropped out because they did not comply to filling out the diary every day. As such there are clear opportunities for more research in this domain, focusing on other mobile platforms and other populations over longer time periods.

Second, the notifications we monitored by not only be viewed via the notification centre or the phone apps. For example, on receiving an email notification, our participants might have

checked their emails on their desktop computer rather than on their phone. Or, the user might have glanced at the notification summary that briefly appears in the notification area when the notification arrives. In this case, the delay between reception and viewing of the notification in our data set would be longer than it actually was. Thus, we have to assume that the actual time between receiving and viewing notifications reported in our current work represent a lower bound (i.e. they are likely to be faster than the 3.2 to 22.0 minutes per category), in particular for emails and social network updates.

IMPLICATIONS

Reducing Interruptions

Reducing the number of interruptions by notifications may seem like an intuitive approach, given the large volume of notifications phone users have to deal with every day. For example, Rosenthal *et al.* [27] have shown that phones can be trained to automatically mute according to contextual information. Previous work has shown the benefits of not being exposed to email notifications: information workers checked their emails less opportunistically [20], spent more time on tasks and felt less stressed [22]. They still appreciated the awareness that notifications provide, and all but one participant reverted to the use of notifications [20].

Our work confirms the negative emotions associated with receiving too many emails found in [22]. However, we also found a generally positive attitude towards notifications and even an increased feeling of being connected with others in the case of messaging and social network updates. Given that our participants typically viewed messages and social networks updates within minutes, and given that muting notifications had no effect on those viewing times, we would expect that even if these notifications are reduced, it's likely that mobile users will check their phones more frequently to make sure that no "important" or "urgent" message has been missed. This should be considered when designing systems to reduce the number of interruptions from notifications.

Delivering Notifications at the Right Time

Some prior work has focused on *when* to delivery notifications. Examples include deferring the delivery of emails by a few minutes when the user is busy [18], delivering emails during the transition between two work tasks [20], and delivering phone notifications between physical activities [16] or after finishing a phone call or text message [13].

Our data shows that some participants delayed attending to emails and social network updates when they were busy. On the other hand, notifications from messengers and social networks were typically viewed within minutes, even if they interrupted the receivers. Thus, systems that delay notifications need to consider this level of immediacy, in particular in personal communication. For example, a system that delays the delivery of a message until the end of a meeting may at times violate user expectations.

Communicating Availability

A rich body of past literature has focused on understanding a persons *availability* to allow people to judge whether it is

a good time to engage in a conversation. Proposals for automatically inferring a persons availability in office spaces [15] and desktop instant messengers [2], phone calls [25] and mobile instant messengers [26] have been explored. Other researchers have investigated the use of rich-presence information including motion, location and music to enable users to determine when and how to contact an individual in mobile environments [4].

In our study, we found that participants at times decided to defer viewing or attending to notifications because they were busy with other tasks or priorities. In such cases, we believe that providing information about a person's *availability* is the most promising approach to a better user experience. One of the key challenges in detecting availability to be solved in future work, however, is that in daily life, a person is not strictly available or unavailable. Rather, a person is available for different people or in different situations at different times [4]. Thus, understanding availability is a complex, pervasive problem and one that has not been fully solved by past work, especially for always-connected mobile phone users.

CONCLUSIONS

We found that mobile phone users have to deal with a large volume of notifications, mostly from messengers and email, each day (63.5 on average per day), which was perceived as the usual. Notifications were largely checked within a few minutes of arrival, regardless of whether the phone was in silent mode or not. In particular in the case of personal communication, explanations for these fast reaction times related to high social expectations and the exchange of time-critical information. Increasing numbers of notifications, in particular from email and social networks, correlated with negative emotions, such as stress and feeling overwhelmed. Personal communication, on the other hand, also related to increased feelings of being connected with others.

These findings highlight that strategies are needed to lower negative emotions. Reviewing previously explored approaches, our findings imply that reducing interruptions and deferring notifications may work in a professional context. For a personal context, strategies around communicating (un)availability and managing expectations appear more suited.

While our work sheds first light on the relationship between received notifications and subjective emotional responses, more work is needed to extend our knowledge to other populations. Our findings encourage future work, in particular given that asynchronous communication via mobile phones is still on the rise and becoming a more and more essential form of communication for large groups in our societies.

ACKNOWLEDGEMENTS

We thank the participants for their contributions.

REFERENCES

1. Adamczyk, P. D., and Bailey, B. P. If not now, when?: the effects of interruption at different moments within task execution. In *Proc. CHI '04*, ACM (2004).

2. Avrahami, D., and Hudson, S. E. Responsiveness in instant messaging: predictive models supporting inter-personal communication. In *Proc. CHI '06*, ACM (2006).
3. Battestini, A., Setlur, V., and Sohn, T. A large scale study of text-messaging use. In *Proc. MobileHCI '10*, ACM (2010), 229–238.
4. Bentley, F., and Metcalf, C. J. The use of mobile social presence. *IEEE Pervasive Computing* 8, 4 (Oct. 2009), 35–41.
5. Birnholtz, J., Reynolds, L., Smith, M. E., and Hancock, J. "everyone has to do it: " a joint action approach to managing social inattention. *Comput. Hum. Behav.* 29, 6 (Nov. 2013), 2230–2238.
6. Böhmer, M., Hecht, B., Schöning, J., Krüger, A., and Bauer, G. Falling asleep with angry birds, facebook and kindle: a large scale study on mobile application usage. In *Proc. of MobileHCI '11*, ACM (2011), 47–56.
7. Church, K., and de Oliveira, R. What's up with whatsapp? comparing mobile instant messaging behaviors with traditional sms. In *Proc. MobileHCI '13*, ACM (2013).
8. Cutrell, E., Czerwinski, M., and Horvitz, E. Notification, disruption, and memory: Effects of messaging interruptions on memory and performance. In *Proc. INTERACT '11* (2001).
9. Czerwinski, M., Cutrell, E., and Horvitz, E. Instant messaging and interruption: Influence of task type on performance. In *Proc. OZCHI '00* (2000).
10. Czerwinski, M., Horvitz, E., and Wilhite, S. A diary study of task switching and interruptions. In *Proc. CHI '04*, ACM (2004).
11. de Vries, R. A. J., Lohse, M., Winterboer, A., Groen, F. C., and Evers, V. Combining social strategies and workload: a new design to reduce the negative effects of task interruptions. In *Proc. CHI '13 EA*, ACM (2013).
12. Felt, A. P., Egelman, S., and Wagner, D. I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns. In *Proc. SPSM '12*, ACM (2012).
13. Fischer, J. E., Greenhalgh, C., and Benford, S. Investigating episodes of mobile phone activity as indicators of opportune moments to deliver notifications. In *Proc. MobileHCI '11*, ACM (2011).
14. Fischer, J. E., Yee, N., Bellotti, V., Good, N., Benford, S., and Greenhalgh, C. Effects of content and time of delivery on receptivity to mobile interruptions. In *Proc. MobileHCI '10*, ACM (2010).
15. Fogarty, J., Hudson, S. E., Atkeson, C. G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J. C., and Yang, J. Predicting human interruptibility with sensors. *ACM Trans. Comput.-Hum. Interact.* 12, 1 (Mar 2005), 119–146.
16. Ho, J., and Intille, S. S. Using context-aware computing to reduce the perceived burden of interruptions from mobile devices. In *Proc. CHI '05*, ACM (2005).
17. Holtgraves, T., and Paul, K. Texting versus talking: An exploration in telecommunication language. *Telemat. Inf.* 30, 4 (Nov. 2013), 289–295.
18. Horvitz, E., Apacible, J., and Subramani, M. Balancing awareness and interruption: investigation of notification deferral policies. In *Proc. UM '05*, Springer-Verlag (2005).
19. Iqbal, S. T., and Bailey, B. P. Effects of intelligent notification management on users and their tasks. In *Proc. CHI '08*, ACM (2008).
20. Iqbal, S. T., and Horvitz, E. Notifications and awareness: a field study of alert usage and preferences. In *Proc. CSCW '10*, ACM (2010).
21. Leiva, L., Böhmer, M., Gehring, S., and Krüger, A. Back to the app: the costs of mobile application interruptions. In *Proc. MobileHCI '12*, ACM (2012).
22. Mark, G., Volda, S., and Cardello, A. "a pace not dictated by electrons": an empirical study of work without email. In *Proc. CHI '12*, ACM (2012).
23. Nardi, B. A., Whittaker, S., and Bradner, E. Interaction and outercation: instant messaging in action. In *CSCW '00*, ACM (2000).
24. Oulasvirta, A., Rattenbury, T., Ma, L., and Raita, E. Habits make smartphone use more pervasive. *Personal Ubiquitous Comput.* 16, 1 (Jan 2012), 105–114.
25. Pielot, M. Large-scale evaluation of call availability prediction. In *Proc. UbiComp '14* (2014).
26. Pielot, M., de Oliveira, R., Kwak, H., and Oliver, N. Didn't you see my message? predicting attentiveness in mobile instant messaging. In *Proc CHI '14*, ACM (2014).
27. Rosenthal, S., Dey, A. K., and Veloso, M. Using decision-theoretic experience sampling to build personalized mobile phone interruption models. In *Proc. Pervasive '11*, Springer-Verlag (2011).
28. Sahami Shirazi, A., Henze, N., Pielot, M., Weber, D., and Schmidt, A. Large-scale assessment of mobile notifications. In *Proc. CHI '14*, ACM (2014).
29. Skierkowski, D., and Wood, R. M. To text or not to text? the importance of text messaging among college-aged youth. *Comput. Hum. Behav.* 28, 2 (Mar. 2012), 744–756.
30. Teevan, J., and Hehmeyer, A. Understanding how the projection of availability state impacts the reception incoming communication. In *Proc. CSCW '13*, ACM (2013).
31. Weiser, M. The computer for the 21st century. *Scientific American* 265, 3 (1991), 66–75.
32. Whittaker, S., and Sidner, C. Email overload: exploring personal information management of email. In *Proc. CHI '96*, ACM (1996).