Perceptibility of Mobile Notification Modalities During Multitasking in Smart Environments

Andreas Komninos, Jeries Besharat, Vassilios Stefanis, John Garofalakis
Computer Engineering & Informatics Department
University of Patras
Rio 26504, Greece
[akomninos, besarat, stefanis, garofala]@ceid.upatras.gr

***Abstract*— Smartphone notifications are a major source of interaction with mobile devices. In this paper we address a gap in literature by establishing a foundation that explains the role of modalities with which a notification is delivered on a mobile device. Though previous studies have attempted to address this question, we find that they suffer from significant internal validity problems. As such we conducted an ecologically valid and carefully designed experiment in a controlled environment that simulates a smart home environment. Our work extends into smart environments, by examining a new modality, implemented by pairing the smartphone to a connected lighting system and replicating mobile notifications to smart light bulbs in the user's field of vision. We derive a set of guidelines for choosing notification modalities and set future research directions.**

Keywords*— *Mobile notifications; Modalities; Smart homes; Smart environments

I. Introduction

Notifications on mobile devices are generated by all kinds of applications and services running on our smartphones, and users receive on average more than 60 mobile notifications daily [24]. Notifications are typically tended to within a few minutes, with time taken to dismiss these depending on various context factors [17, 24, 31]. A considerable body of literature deals with the identification of opportune moments in which to notify users of events occurring on their mobile device (e.g. [27]), and the typical behaviour with mobile notifications has been studied recently in a number of key papers (e.g. [32]). An important shortcoming in the available body of literature is lack of research into the perceptibility of mobile notifications with regard to the modality (or combination thereof) with which it is being delivered. Some previous literature attempts to address this gap (e.g. [6, 13, 16]), though the data it reports on comes from field studies, where a number of factors that may affect the reliability of results are not controlled for, reducing from the internal validity of these studies. Addressing the internal validity issue forms the primary goal of this paper. As a secondary goal, this paper aims to add to literature by comparing traditional modalities of delivering mobile notifications on the device, with the concept of extra-device notifications, in the form of ambient smart lighting. The synchronization of mobile notifications across multiple devices has been considered or studied in the past (e.g. [31, 36]) but with the proliferation of domestic connected IoT systems (e.g. Philips Hue, Apple Homekit, GE Link), the question of how domestic appliances may integrate in the user's workflow of managing notifications becomes timely. Connected lighting notifications also extend into the needs of users with disabilities (e.g. hard of hearing) or situations where other modalities are socially inappropriate.

II. Related Work

A. *Notifications on Mobile Devices*

Modern smartphones generate many notifications daily [24], relating to multiple types of events (14 application categories are described in [32]). These notifications are typically dealt with in a short time frame measured in minutes, depending on which application generated it, its perceived importance to the current task, the social relationship between the user and a person relating to the notification, the hour and day of the week, the device that generated them, the user's personality and the current task a user is engaged in [17, 20, 27, 32]. Thus, not all types of notification are important to the users under a given context. It also becomes apparent that the time taken to dismiss a notification does not only depend on the perceptibility of the notification. We can therefore frame the research relating to interaction with notification into three main topics: a) how can the importance of a notification be assessed under context; b) how to determine opportune moments to deliver notifications and thus reduce the impact of interruptions; and c) what modalities are best to use under the given context, in order to efficiently deliver a notification, giving it a high chance to be perceptible while causing minimal disruption. Existing literature has provided useful insights in answering questions a) and b) by extrapolating the notification generation context and user availability through a range of techniques that involve device sensors and analysis of the notification content and generation parameters, to deliver the right notification at the most opportune time [17, 20, 21, 22, 27]. In this paper we are concerned with determining appropriate notification modality.

B. *Modality of notifications on mobile devices*

Typical smartphone notification modalities are visual (including notification icons and the device status LED), auditory (including speech and sound) and haptic (vibration), though sometimes vibration can have unintended audio effects as well (e.g. when a device is vibrating against a hard surface). Users are able to control all three modalities on a modern smartphone like Android, though it is most typical that the users will switch between sound (on/off) and vibration (on/off) ringer mode combinations during the day [4]. Some users may totally silence their devices, however studies such as [16, 24] show that silent mode does not prevent users from becoming aware of the notification events within a reasonable timeframe (this is explained in [4] as the users enter a proactive monitoring state). Users are up to 12 times more likely to immediately attend to a notification if it is delivered with at least one modality [13].

The main distinction between modalities relates to their persistence. Audio tones (with the exception of phone calls) and haptic modalities are momentary, meaning that if the users are unable to respond to the notification immediately, they may never become aware of it, or forget about it [8]. Visual modalities such as the screen display (status bar & lock-screen icons) or device LEDs persist until the notification is dismissed, which helps participants in deciding to react later [13]. When trying to apply the concept of persistence in haptic feedback, researchers in [10] created a device which would constantly vibrate and whose vibrations would increase in frequency and intensity, according to the number of pending notifications active on the device. The concept was found annoying but users were still able to distinguish between the constant "idle" pulse and the more intense vibrations associated with notifications. A further distinction between modalities can be made by thinking of these as private or public, in terms of who can perceive them. An audio notification is public, while a haptic or visual notification is often considered private, though this is not always the case: a blinking device LED can be visible to all who can see the device even from a distance and sometimes a vibration can be heard or felt by others too, e.g. if the device is on a table. The device LED affords users some awareness of which application the notification is coming from through its colour, or importance of the notification depending on the blinking rate [13]. However, device LEDs have the disadvantage of being small and not overly bright, and only a fraction of devices incorporate RGB LEDs. In [7] an attempt was made to discover whether audio modalities could afford the same types of awareness to users and the researchers found that, for distinguishing between application categories, speech was the best performing modality, followed by auditory icons and lastly earcons (typically used in mobile notifications in modern smartphones).

The impact of modality on the perceptibility of notifications has not been widely studied. In [6], the majority of users (65%) were shown to prefer a combination of modes that includes sound. Further field studies report conflicting results based on the current device ringer mode. In [13] it was impossible to discern any statistically significant difference in notification reaction time when comparing across modalities and their combinations, even when asking participants to manually rank their preferred choice of modality (though it was found that users like to associate vibration and sound to important notifications, and that social context plays a role in determining modality choices). In [16], reaction time was found to be lowest with vibrate-only mode, followed by sound-only and sound-vibrate. In [24], reaction time was found to be faster when the phone was on vibrate-only mode, followed by silent mode and normal mode (between which there was no difference). However, there are two main problems with these field studies: Firstly, field studies suffer from inherent internal validity problems which are very pronounced in this case. These studies did not control for a number of everyday behaviors, which might have affected the noticeability (or reaction times) of notifications significantly. For example, if a user left their device in a jacket pocket, or on a desk near other clutter, or in another room e.g. to charge, as reported in [4], then obviously the measurements would be affected. As other literature indicates [17, 20, 27, 32], the current task and social context of the user can strongly affect the measured response times. A clearer demonstration of the internal validity issues from the aforementioned studies comes from [25], where it was discovered that ringer mode (unknown, silent, vibration or sound) is a weak predictor of the attentiveness of a user towards their mobile device with the purpose of noticing a message notification, while other indicators that are not pertinent to the notification itself (e.g. time elapsed since last "screen on" event or "hour of the day") are stronger predictors for attentiveness. It's not clear in this study whether sound and vibration modalities were considered separately or in conjunction with one another, but it highlights that the generalizability of findings from previous field studies is weak, owing possibly to the lack of internal validity.

The second major internal validity issue with these field studies is that they are based on capturing the user's device ringer mode. This is problematic because ringer mode may suppress, but does not add beyond the programmed modality requests (thus will not add a LED illumination, vibration or sound to a notification which is not programmed to have one). When a phone is set on "vibrate only", it doesn't necessarily mean that every application generating a notification will result in a vibration. To infer thus reliable conclusions on how a notification modality influenced response time, a study should capture all types of information (what were the programmed notification modalities, user per-app preferences and what was the current ringer mode at the time of notification).

C. *Notifications with Ambient Lighting Systems*

The use of simple connected household devices to convey information to users (such as ambient lighting and peripheral displays) is a concept that has been discussed under the principles of Calm Computing since the early days of ubiquitous computing (e.g. [11, 12]). With the affordability of connected lighting systems that can interface with smartphones, it is easy to see that a natural synergy for solving the shortcomings of mobile visual notification feedback modalities (notably the device LED) can be achieved. In fact, a synergy of ambient lighting and the smartphone for notifications satisfies most of the criteria for ambient interface design, set out by Gross [8]. We are not aware of any literature that investigates the use of ambient lighting for the delivery of smartphone notifications as a precise mapping of the state of the mobile device LED. However, some previous work exists on ambient notification systems. In [29], it is proposed that ambient information systems may conform to four main design patterns, one being a "Symbolic Sculptural Display", i.e. a system that displays very few pieces of information, usually a single element. A system thus consisting of a single light bulb that replicates mobile notifications can be considered to fall under this category, but the authors do not propose specific ways for designing the function of such systems, other than that the system must support transitions to prevent "change blindness". A thorough survey of existing ALSs that fall into this category can be found in [15], however, there are gaps in all of the surveyed papers: either the systems presented are evaluated in preliminary trials with very few participants (e.g. [2]), or there are, at best, limited comparisons between alternative designs for the perceptibility of conveying simple notifications (e.g. in [23] it is argued that blinking or animated lights should be used, but only blinking vs. static light was actually examined).

In [15], the term "ambient light system (ALS)" is used to describe "a system positioned in the periphery of a person's attention that conveys information using light encodings in a non-distracting way most of the time". The authors propose four general guidelines, one of which states that a light's blinking rate is the most suitable pattern for notification encoding. This guideline is partially supported in [14] via a participatory design process, but the researchers did not experimentally evaluate its effectiveness. Supporting change and state transition in ALSs is demonstrated in [18], where an RGB LED strip reflecting light on the wall behind a computer monitor, gradually changed color from green to red, depending on how much time remained for a user to complete a task with a deadline. This can be seen as a persistent notification system, but has little practical relationship with the majority of spontaneously issued notifications that are typically issued through events in mobile devices and do not contain a temporal dimension. In [30] the smartphone is augmented with additional LEDs able to project light surrounding the device, however this work was presented as a prototype and not evaluated. Other research such as [19, 26], extend the modalities of a smartwatch or a tablet using additional LEDs, but, just as in [30], these extensions are on-device and demand the user's attention is already on the device itself. Hence, they do not conform to the definition of an ALS.

ALS notification systems have privacy implications, as highlighted for example in [9], where participants raised significant concerns. So far, only [34] have investigated the issuing of ambient notifications with the user engaged in a social activity and found that the presence of another person in the room did not affect the acceptability of the notification, regardless of intrusiveness of the modality. However, the participants in the study were real-life couples and the close relationships (trust and familiarity between individuals) might have affected this finding. Finally, as far as the positioning of ALSs is concerned, in [1], it was found that the most suitable location for a general use ALS is a living room or office, while specific room types (e.g. the kitchen) should be reserved for special purpose objects. We can conclude thus that the guidelines for designing ALS notification services are still not definitive but there seems to be some evidence that perceptibility and interpretation can benefit from gradual state transitions and appropriate use of color or blinking patterns.

Summarizing the previous literature, we can derive the following open issues. Firstly, although a range of studies report findings on the impact of modality on the perceptibility of smartphone notifications, the generalizability is limited because they did not control for extrinsic contextual factors. Secondly, research on notification-based ALSs focuses on the use of these systems as stand-alone replacements and not extensions of a smartphone. Where lighting has been investigated as an extension to the modalities available on a device, this has been done by augmenting the device itself, hence negating the definition of an ALS by placing the lighting at the center of the user's attention and not their periphery. There is hence no present understanding of how an ALS can extend the modalities for issuing notifications on a smartphone. Based on the above, our main research questions are (R1) "*How do modality combinations affect the perceptibility of smartphone notifications in a typical use environment?*" and (R2) "*How does the extension of mobile notifications to ambient lighting affect the perceptibility of smartphone notifications, alone or in conjunction with existing smartphone modalities, in a typical use environment?*".

III. Investigating Notification Modalities

A. Experiment set-up

To answer our research questions, we decided to proceed with a laboratory experiment, whose controlled conditions would complement existing studies by focusing on adequate ecological validity and maintaining strong internal validity. As such, our aim was to examine user behaviour in perceiving notifications in a controlled but realistic setting.

Choice of environment: Mobiles are used in a variety of environments and settings. For the purposes of this experiment we chose to simulate an environment such as a home office, where the user might be engaged in multiple simultaneous tasks, hence not constantly paying attention to their mobile device. A study of 693 participants [38] demonstrates that the most probable place for a user's mobile to be found at any time, regardless of environment, is "out on the table or desk" ("right now" 68%, over 24hrs 83%), followed by the front trouser pocket ("right now" 14%, over 24 hrs 64%). Hence, though smartphones are highly mobile, in reality they are mostly placed on flat surfaces near the user. Typical environments that contain such flat surfaces are home or office environments, which are also the natural "habitats" for an ALS [1]. Thus, placing the user and their mobile in a home office setting is only ideal for our experiment which involves an ALS, and reflects a highly ecologically valid scenario. We considered also a set-up where we might position the device in second most common location, i.e. the front trouser pocket of users. However, this would reduce the audibility of sound notifications, depending on the pocket lining material of our users' clothing and the possibility of some female users also wearing thick tights (the experiment was during the winter), thus creating uneven conditions. Additionally, this set-up would remove the visibility of the LED notification, preventing us from examining it. Though there exists no previous literature on where exactly users place their devices on desk, in [5, 38] it is shown that users mostly keep their device within arm's length for easy reach and maximum perceptibility of notifications. We thus selected a smartphone position as per Fig. 1 so as plausibly emulate a user's behaviour (i.e. the device being within easy reach, facing up and within our participants' field of vision, so that the device LED at the bottom of the device can be easily seen). Some smartphone displays light up from sleep when a notification is received, but we excluded this option. Finally, our set-up included an ALS using a single bulb, placed directly behind the participant's monitor, positioned close to the wall (Fig. 1). Hence the participant could not directly see the bulb, but was aware of its state as light reflected on the wall and desk surface behind the monitor.

Activities during experiment: Using multiple activities to overload the user's cognitive processing ability in multiple channels was intentional and essential to our experimental design, as it would help prevent the participants' intentional focus on the mobile device and remove their ability to direct their attentiveness to it, something that would bias their behaviour [4]. Hence, we selected to engage the participants in two parallel tasks, which would overload their vision and hearing channels. The first task was to play a game on a desktop computer screen – for this we selected the well-known Bubble Shooter game, which is very simple to learn and known for its addictiveness. We asked participants to engage in the game, without worrying about high-scores or losing (if they lost, they could start over). At the same time, we played a recording of a basketball game through speakers in the room, at an average volume level of 50db, roughly equivalent to a conversation. Participants were asked to pay attention to the game and take a note every time a particular well-known player's name was mentioned. We developed an application running on a smartphone, which was able to generate notifications using all device modalities including an ALS. Participants were asked to make dismissing the notifications on the smartphone as they noticed them their top priority, and to position the device back to its marked position on the table, as shown in Fig. 1, after dismissing a notification. Notifications were dismissed from the device's lock screen (Fig. 2 right) by swiping on the notification. If a participant did not dismiss a notification before another was issued, both notifications persisted on the screen and the participant was asked to swipe both.

Software, devices and settings: For this experiment, we built a simple Android application that generates notifications to the user's mobile (Fig. 2). These notifications are issued with all the possible combinations including the modalities of sound, device LED, vibration and ambient lighting, with the exception of issuing a notification with none of the above modalities (as the user would have no way of perceiving it). This resulted in 15 different notification modality combinations (Table 1, modality presence in a combination denoted by a dot). The application issues each modality combination twice in the experiment session, resulting in 30 total notifications to the user. Notifications are issued at a regular interval of 30 seconds and the sequence of notifications (with regard to the modality combination) is wholly random. We did not tell our participants that notifications would arrive at regular intervals. This was a restriction due to the task of taking note of the basketball player name mentions, and, as will be explained below, performance in this task could only be equally measured if the task was executed in an equal time period by all participants. We anticipated that because of the random selection of notification modality and the high cognitive load placed by the parallel tasks, participants would not be able to learn and anticipate the notification timing. We will demonstrate in section 3.3.1 that this approach did not lead to learning effects, hence our results are not invalidated by this aspect of the experiment design.

For our experiment, we used a Google Nexus 5 device. For notification sound we selected one of the in-built Android tones ("Tethys"), a simple two-tone short earcon, with the volume level set to 90% of the maximum supported by the device. This was the highest volume setting without the speaker producing "tinny" noises, which could alter participants' perceptions of annoyance of the modality. The device LED color was set to white, blinking at a repeating pattern of 3 seconds ON, 3 seconds OFF. The vibration was set to a pattern of 200ms ON, 200ms OFF, 200ms ON (i.e. two vibrations for each notification). These patterns are similar to popular applications where notifications are an integral part of the user experience (SMS, Gmail, Facebook messenger). In this experiment, we are not concerned with the encoding the importance of a notification but only with the perceptibility of the notification event, hence we chose a white LED color and patterns that are arguably of average intensity so as not to convey priority semantics.

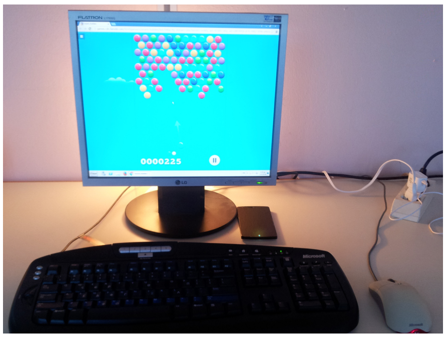

Fig. 1. Participant workstation. A notification is being displayed via device LED and lighting behind the monitor.

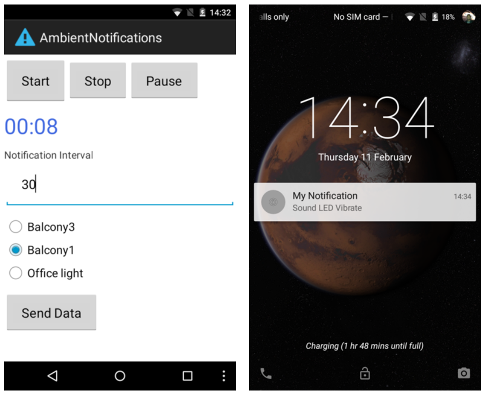

Fig. 2. The notification generator application control interface (left) and locked device screen with an issued notification (right).

TABLE I. NOTIFICATION MODALITY COMBINATIONS

Modality combination	Sound	Device LED	Vibration	Ambient Lighting
A				•
B			•	
C			•	•
D		•		
E		•		•
F		•	•	
G		•	•	•
H	•			
I	•			•
J	•		•	
K	•		•	•
L	•	•		
M	•	•		•
N	•	•	•	
O	•	•	•	•

Finally, issuing notifications vial an ALS, we extended our application to work with Philips HUE, using a single A19 9W bulb, setting the colour of the bulb to white and brightness of

70% (empirically set with 6 colleagues, 3 females, for a comfortable and perceptible level) and with the bulb remaining ON until the user dismisses the notification. Because we are not concerned with the encoding of information in a notification, we used a neutral light colour (white) and did not employ blinking, so as to prevent any assumptions regarding the importance of notifications. Based on [23], we implemented a gradual transition between the "OFF" and "ON" states, to animate the light and attract attention. Finally, the notification text was generic, to exclude any priority semantics (Fig. 2).

B. Experiment

We recruited 25 participants, 7 females. All participants were Computer Science students, and in the age bracket of 18-29 (the questionnaire presented a list of ages brackets and did not require the precise age). Most participants were Android device users (20) but we had also 2 iOS and 3 Windows Mobile users. They were not incentivized for their participation.

Validation of the task and method: Our first concern was to validate the appropriateness of the parallel tasks (i.e. playing the game and listening for the basketball player's name) for occupying the participants' attention. At the start of our experiment, we asked participants to report using 5-point Likert scales both familiarity with the sport (all mentioned often or very often watching or attending basketball games) and the player's name (all reported him to be known, or well known to them). Because the player's name is mentioned a specific number of times during the recording, we needed participants to complete the task in the same timeframe, so as to effectively measure their performance in listening for the player's name. We expected that users would not learn to anticipate the timing of notifications, because of the random selection of notification modalities and the high level of cognitive engagement in the tasks at hand. Indeed, it is common psychology knowledge that human cognitive ability suffers from a processing bottleneck which reduces our ability to perform multiple tasks at the same efficiency as single tasks, so much so that cognitive overload in one or more channels with essential or incidental information, removes our ability to learn new information. Though humans are able to monitor several streams of information for a specified target, if those streams contain a target at the same time, or close together in time, some targets will be unavoidably missed.

Still, we examined the average response times for the first 10 (A), middle 10 (B) and final 10 (C) notifications issued to each participant (Fig. 3). A Friedman test (due to the distribution of the data), failed to reveal any statistically significant differences in these three subsets (χ^2=2.880, p=0.237). Post-hoc Bonferroni corrected pairwise Wilcoxon Signed Rank tests (2-tailed) between the individual sets also did not reveal any statistically significant differences. Since Friedman tests do not allow the computation of effect sizes (Cohen's d), these are presented for the pairwise Wilcoxon signed rank tests (A-B d =0.18, A-C d =0.27, B-C d =0.13). This outcome can mean that either there is no learning effect to be detected, or our study is insufficiently powered to detect one. Common practice where no prior study data is available (as is our case), is to calculate sample sizes according to a general estimate, e.g. Cohen's d=0.5 (a medium size effect), α=0.05 and study power (1-β)=0.80. Based on the above, a power analysis of our study shows that it is sufficiently powered (1-β=0.81) to detect a one-tailed medium size effect (because we are only interested in whether the response time average decreases due to learning effects), at the lower boundaries of Cohen's "medium effects" category (d=0.53). Thus, our sample size can be considered as appropriate.

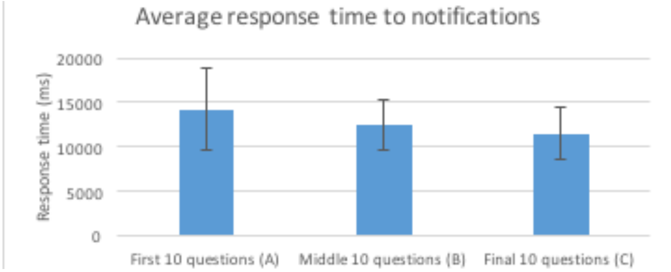

Fig. 3. Perceptibility of notification modalities (error bars at 95% c.i.)

Given the size of the reported effects, our study is not sufficiently powered to determine statistical significance in the discovered small effects, but we can confirm that if any learning effects were indeed present, these were either *very small* or *small*, according to Cohen's d categorisation (the lower thresholds for *very small*, *small* and *medium* effects are 0.1, 0.2 and 0.5 respectively). The long average response time shows that our participants' attention was occupied very heavily by the two parallel tasks. Further corroborating evidence for the heavy cognitive load comes from the fact that out of the total of 19 times that the basketball player's name was mentioned in the 15 minutes of the experiment duration, the average number of instances in which the participants were able to capture the player's name was just 4.36 times (SD=3.2).

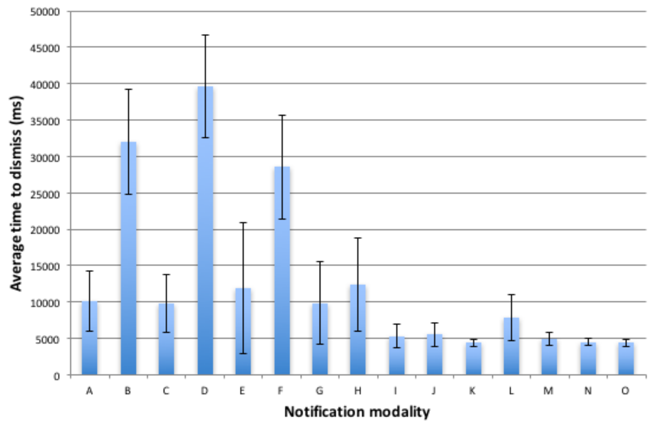

Fig. 4. Notification dismissal time averages (error bars at 95% c.i.).

Experiment results: The results of our experiment are summarized in Fig. 4. The left half depicts all conditions where sound was absent in the notifications, while the right half where sound was present. To examine the results for statistical significance, we used non-parametric tests (Wilcoxon signed rank), due to the distribution of the data. We note that, all other modalities being the same, the presence of ambient lighting has a statistically significant positive impact in reducing the time of dismissal, except in the case where sound is present along with another modality (vibration or device LED) (Table 2). It is also notable that sound has a statistically significant effect on all cases except when the ambient and device LED are displayed,

TABLE II. STATISTICAL SIGNIFICANCE RESULTS ON PAIRWISE CONDITION COMPARISONS FOR THE EFFECT OF MODALITIES IN NOTIFICATIONS

Ambient lighting	**Mod. Combo**	**C - B**	**E - D**	**G - F**	**I - H**	**K - J**	**M - L**	**O - N**
	Z (p-value)	-3.740 (.000)	-3.727 (.000)	-3.484 (.000)	-2.489 (.013)	-.740 (.459)	-1.332 (.183)	-1.251 (.211)
Vibration	**Mod. Combo**	**C - A**	**F - D**	**G - E**	**J - H**	**K - I**	**N - L**	**O - M**
	Z (p-value)	-.202 (.840)	-2.085 (.037)	-.309 (.757)	-2.489 (.013)	-.309 (.757)	-.794 (.427)	-.821 (.412)
Device LED	**Mod. Combo**	**E - A**	**F - B**	**G - C**	**L - H**	**M - I**	**N - J**	**O - K**
	Z (p-value)	-1.251 (.211)	-.040 (.968)	-.821 (.412)	-1.440 (.150)	-.094 (.925)	-.121 (.904)	-.390 (.696)
Sound	**Mod. Combo**	**I - A**	**J - B**	**K - C**	**L - D**	**M - E**	**N - F**	**O - G**
	Z (p-value)	-2.139 (.032)	-4.130 (.000)	-2.516 (.012)	-4.319 (.000)	-.444 (.657)	-4.103 (.000)	-.498 (.619)

and when ambient, device LED and vibration are present. Likewise, vibration seems to have a statistically significant effect on only two cases (when the only other modality is the device LED and when sound and ambient are also present). Finally, the device LED on its own doesn't seem to have any impact on the noticeability of notifications for any of the combinations it participates in. Our results indicate that perceptibility of notifications is strongly affected by the presence of audio cues, followed in efficacy by ambient lighting. While it is not surprising that the device LED did not seem to affect the perceptibility of the notifications, we were surprised to find that vibration did not affect the perceptibility of notifications as it did in previous field studies in [13, 16, 24].

Subjective feedback: We asked participants to rate, post-experiment, how perceptible and annoying each notification modality was. We also asked them if they would prefer a different intensity level for each of the modalities (e.g. sound level, lamp and LED brightness, vibration intensity) and to note their preferences for these questions on 5-point Likert scales. Participants reported high levels of perceptibility for the sound and ambient lighting modalities, while mostly indicating that vibration and the device LED were hard to perceive (Fig. 5 top). These findings support our quantitative analysis. Understandably, the LED provides a rather small visual cue which is also situated at the periphery of the participants' field of view. Despite being persistent and blinking, the LED brightness and size are not enough to draw participants' attention away from their current tasks. With regard to vibration, it can be theorized here that the low perceptibility of this modality is due to the placement of the device on the table and not in direct contact with the user's body. However, although this explanation sounds plausible, it cannot be unreservedly accepted. There is a significant body of literature documenting the experience of "phantom vibrations", with the evidence pointing towards a mental mechanism of manifestation of false perceptions, linked to the context of use and the level of use of the mobile device (e.g. [33]). It is therefore questionable how much of the perceptibility of a vibration is truly dependent on the sensory capability of humans and not affected by mental processes. In [3], for example, it is shown that the perceptibility of haptic feedback is severely impeded under high cognitive workload. Given the cognitive load imposed on our participants and the short duration of the vibrations (as is standard on mobile devices), it is plausible that vibration perceptibility was low, not because the participants weren't able to sense it, but rather, their brain could not process it due to the multitasking conditions.

Modality annoyance was reported higher for the ambient lighting and sound modalities, an expected result, as these were the two modalities that our participants noticed the most (Fig. 5 middle). These two modalities, particularly ambient lighting, provided a sense of urgency to remove the notification for our participants, as reported in post-experiment discussions. Some participants expressed concern on the impact of the longevity of the bulb or the consumption of energy if the bulb was left on for too long, given the fact that the ambient lighting notification continues to remain switched on until the user dismisses the notification. Inversely, the sound caused a sense of urgency, since it was a non-repeating modality and thus some users felt they should dismiss the notification right away and postpone any ongoing tasks, for fear that they might forget about it if left for later.

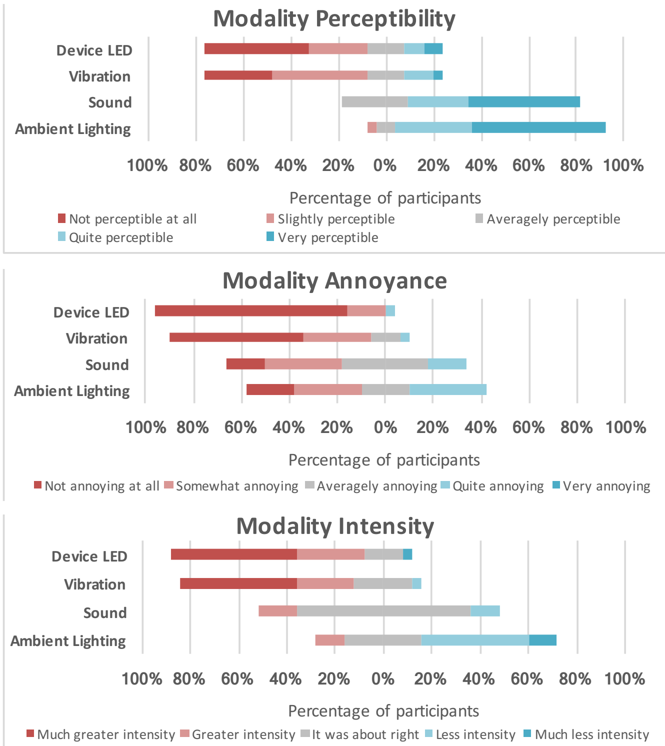

Fig. 5. Response distribution to subjective questions.

Finally, we asked our participants to indicate for each modality, how they might change the intensity with which it was delivered. As can be seen (Fig. 5 bottom), and in line with their previous reports, 20 participants would like to intensify the device LED and 18 the vibration levels. While the latter is not possible on the Android platform, vibration intensity can be changed by prolonging the ON period in the pattern. Extending the ON period or repeating the ON-OFF pattern for longer, might also help with the perceptibility of the haptic feedback when under heavy cognitive load, as suggested by [3]. 18 participants felt that the sound level was about right. Participants were somewhat negative regarding the brightness of the ambient lighting, with 14 participants preferring a lower setting.

IV. Discussion of Results

Our work has clear implications for the design of notification modalities during multitasking. Firstly, we would like to emphasize the importance of having tested the perceptibility of notification modalities in a controlled environment, compared to recent in-the-wild studies. As literature shows, users' response times to notifications are dependent on a wide range of factors. Hence previous studies reporting on the correlation of notification modalities with perceptibility are largely unreliable. We were not able to replicate the users' better response times with vibrations found in [13, 16, 24], placing doubt on the external validity of these studies. While intuitive explanations of some of our results could be given (e.g. that vibrations were not highly perceptible because devices were placed on the desk), literature such as [3, 33] shows that other explanations involving the human cognitive system, not considered in previous studies, may offer better plausibility. The findings and analysis of our quantitative and qualitative data in light of the current literature, lead to some design recommendations for notification modalities when a user is multitasking in a smart environment:

1. Audio and ALS notifications are good for attracting the users' attention. The combination of these modalities with others does not increase the perceptibility of a notification. However, these notification modalities have privacy implications.

2. Vibration and device LED notifications are easy to miss. These modalities are more protective of users' privacy, though under certain circumstances they can be publically perceptible.

3. The device LED is appropriate for private, non-urgent notifications (the user will eventually see it but with a likely long delay). It also serves as a reminder for notifications issued via other modalities (audio, haptic) which the user might have temporarily ignored or missed altogether. However, when an ALS is present, it is best to use the ALS for this role as the LED is the least perceptible modality.

4. Vibrations are appropriate for private, urgent notifications only if the duration of the pattern is long enough to allow it a chance to be perceived. They also encourage "phantom vibration" experiences, especially under certain contexts of use.

V. Conclusion

Our paper offers a comprehensive review of literature on the perceptibility of smartphone notifications and discusses implications from literature about how smartphone notifications can be extended to work with ALS. As can be seen from the literature that we have explored, and from papers in very recent workshops on notifications [28, 35, 37], the topic of modality perceptibility and appropriateness under context is still very much open for research, particularly when considering notifications involving an ALS. Our work does not aim or claim to provide a comprehensive finding to all the open issues in this domain – this is a target far beyond the scope of a single paper. However, it is the first study in this area which attempts to establish a baseline regarding the perceptibility of smartphone and ALS notifications, using a thorough controlled environment trial. Through this, we have aimed to ground the discourse on notification modalities on a more solid basis and to provide some reliable design guidelines, representing one highly ecologically valid scenario of mobile device use in smart environments.

Our results for R1 & R2 show that sound, whether on its own or in combination with other modalities is strongly related to the perceptibility of a notification, when the user is engaged in an attention-demanding task. Additionally, we found that ambient lighting is successful in attracting users' attention during such circumstances, but did not find any evidence to support the use of the device LED, or, in contrast to other researchers, vibration. For the latter, literature suggests that the reasons might be due to the underlying psycho-cognitive mechanism of dealing with haptic stimuli under multitasking. We would have expected the results of these field studies to at least partially confirm our own, since field studies are largely regarded as more externally valid. This contrast in results highlights that previous field studies might have been premature (without considering the role of contextual factors in perceptibility and response time, or the full mechanism behind Android notifications) or too loosely structured to ensure a reasonable degree of internal validity.

Our study has some limitations as can be expected from a laboratory environment. We examined only one contextual setting, where multitasking overloaded participants' visual and audio perception. The experiment should be repeated with the user engaging in other types of task, such as leisurely ones. A further limitation was the placement of the mobile device on the desk. Although literature indicates that this is the most representative and common case in real use of mobile devices, it would be interesting to repeat the experiment with the device placed in contact with the user's body, something that might increase the perceptibility of vibration. This positioning of the device (e.g. in a pocket), might diminish the user's ability to perceive visual and audio cues. Hence it might be worth expanding the notification space to wearable devices (e.g. vibrating smart watch) that form part of the user's device ecology. A final consideration is the use of ALS modalities in a shared space, particularly regarding privacy issues or conflicts, e.g. if used simultaneously as a room's main lighting source.

In general, we can state here that appropriate context awareness and consideration for the users' physiology and cognitive ability has the potential to improve the design of notifications in terms of their modality, increasing their effectiveness and reducing annoyance. An added advantage of careful notification design is the conservation of device battery resources, by avoiding the use of modalities that are ineffective under the given context (especially when being able to off-load notifications to an ALS). There is scope for operating system developers to implement APIs that can be utilized by all third-party developers. As demonstrated by [38], a simple periodic assessment of the device's placement context using integrated sensors can provide cues for selecting notification modalities can either by the application, or managed by the OS itself.

References

[1] F. Alt, A. S. Shirazi, A. Kaiser, K. Pfeuffer, E. Gurkan, A. Schmidt, P. Holleis, and M. Wagner, “Exploring ambient visualizations of context information,” 8th IEEE International Conference on Pervasive Computing and Communications Workshops, 2010.

[2] M. Altosaar, R. Vertegaal, C. Sohn, and D. Cheng, “AuraOrb,” CHI ’06 extended abstracts on Human factors in computing systems, 2006.

[3] A. Chan, K. MacLean, and J. McGrenere, “Learning and Identifying Haptic Icons under Workload,” First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems.

[4] Y. J. Chang and J. C. Tang, “Investigating Mobile Users' Ringer Mode Usage and Attentiveness and Responsiveness to Communication”. In Proc. 17th International Conference on Human-Computer Interaction with Mobile Devices and Services. 2015.

[5] A. K. Dey, K. Wac, D. Ferreira, K. Tassini, J.-H. Hong, and J. Ramos, “Getting closer,” Proc. 13th international conference on Ubiquitous computing - UbiComp ’11, 2011.

[6] J. A. Gallud and R. Tesoriero, “Smartphone Notifications,” Proc. 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI ’15, 2015.

[7] S. Garzonis, C. Bevan, and E. O’Neill, “Mobile service audio notifications,” Proc. 20th Australasian Conference on Computer-Human Interaction Designing for Habitus and Habitat, 2008.

[8] T. Gross, “Ambient interfaces: design challenges and recommendations”. Human computer interaction: theory and practice, 68-72. 2003.

[9] R. Hansson, P. Ljungstrand, and J. Redström, “Subtle and Public Notification Cues for Mobile Devices,” Lecture Notes in Computer Science, pp. 240–246, 2001.

[10] F. Hemmert and G. Joost, “Ambient life: Calm and excited pulsation as a means of life-like permanent tactile status display in mobile phones”. In Proc. of 6th Design & Emotion’08, 2008

[11] A. Kameas, S. Bellis, I. Mavrommati, K. Delaney, M. Colley, and A. Pounds-Cornish, “An architecture that treats everyday objects as communicating tangible components,” Proc. 1st IEEE International Conference on Pervasive Computing and Communications, 2003.

[12] F. Kawsar, K. Fujinami, and T. Nakajima, “Augmenting everyday life with sentient artefacts,” Proc. 2005 joint conference on Smart objects and ambient intelligence innovative context-aware services: usages and technologies - sOc-EUSAI ’05, 2005

[13] A. Mashhadi, A. Mathur, and F. Kawsar, “The myth of subtle notifications,” Proc. 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing Adjunct Publication - UbiComp ’14 Adjunct, 2014.

[14] A. Matviienko, W. Heuten, V. Cobus, H. Müller, J. Fortmann, A. Löcken, S. Boll, M. Rauschenberger, J. Timmermann, and C. Trappe, “Deriving design guidelines for ambient light systems,” Proc. 14th International Conference on Mobile and Ubiquitous Multimedia - MUM ’15, 2015.

[15] A. Matviienko, M. Rauschenberger, V. Cobus, J. Timmermann, J. Fortmann, A. Löcken, H. Müller, C. Trappe, W. Heuten and S. Boll, “Towards New Ambient Light Systems: a Close Look at Existing Encodings of Ambient Light Systems”. IxD&A 26: 10–24. 2015.

[16] A. Mehrotra, V. Pejovic, J. Vermeulen, R. Hendley, and M. Musolesi, “My Phone and Me,” Proc. 2016 CHI Conference on Human Factors in Computing Systems - CHI ’16, 2016.

[17] A. Mehrotra, M. Musolesi, R. Hendley, and V. Pejovic, “Designing content-driven intelligent notification mechanisms for mobile applications,” Proc. 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp ’15, 2015.

[18] H. Müller, A. Kazakova, M. Pielot, W. Heuten, and S. Boll, “Ambient Timer – Unobtrusively Reminding Users of Upcoming Tasks with Ambient Light,” Lecture Notes in Computer Science, pp. 211–228, 2013.

[19] H. Müller, A. Löcken, W. Heuten, and S. Boll, “Sparkle, An Ambient Light Display for Dynamic Off-screen Points of Interest” Proc. 8th Nordic Conference on Human-Computer Interaction Fun, Fast, Foundational - NordiCHI ’14, 2014.

[20] T. Okoshi, J. Ramos, H. Nozaki, J. Nakazawa, A. K. Dey, and H. Tokuda, “Attelia: Reducing user’s cognitive load due to interruptive notifications on smart phones,” Proc. 2015 IEEE International Conference on Pervasive Computing and Communications (PerCom), 2015.

[21] V. Pejovic and M. Musolesi, “InterruptMe,” Proc. 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp ’14 Adjunct, 2014.

[22] V. Pejovic, M. Musolesi, and A. Mehrotra, “Investigating The Role of Task Engagement in Mobile Interruptibility,” Proc. 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI ’15, 2015.

[23] F. Perteneder, E.-M. B. Grossauer, J. Leong, W. Stuerzlinger, and M. Haller, “Glowworms and Fireflies,” Proc. 2016 CHI Conference on Human Factors in Computing Systems - CHI ’16, 2016.

[24] M. Pielot, K. Church, and R. de Oliveira, “An in-situ study of mobile phone notifications,” Proc. 16th international conference on Human-computer interaction with mobile devices & services - MobileHCI ’14, 2014.

[25] M. Pielot, R. de Oliveira, H. Kwak and N. Oliver. 2014. “Didn’t You See My Message?: Predicting Attentiveness to Mobile Instant Messages.” In Proc. SIGCHI Conference on Human Factors in Computing Systems (CHI ’14), 3319–3328. 2014.

[26] H. Pohl, J. Medrek, and M. Rohs, “ScatterWatch,” Proc. 18th International Conference on Human-Computer Interaction with Mobile Devices and Services - MobileHCI ’16, 2016.

[27] B. Poppinga, W. Heuten, and S. Boll, “Sensor-Based Identification of Opportune Moments for Triggering Notifications,” IEEE Pervasive Computing, vol. 13, no. 1, pp. 22–29, 2014.

[28] B. Poppinga, M. Pielot, N. Henze, N. Oliver, K. Church, and A. S. Shirazi, “Smarttention, Please! Intelligent Attention Management on Mobile Devices,” Proc. 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI ’15, 2015.

[29] Z. Pousman and J. Stasko, “A taxonomy of ambient information systems,” Proc. working conference on Advanced visual interfaces - AVI ’06, 2006.

[30] Q. Qin, M. Rohs, and S. Kratz, “Dynamic ambient lighting for mobile devices,” Proc. 24th annual ACM symposium adjunct on User interface software and technology - UIST ’11 Adjunct, 2011.

[31] A. S. Shirazi and N. Henze, “Assessment of Notifications on Smartwatches,” Proc. 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI ’15, 2015.

[32] A. Sahami Shirazi, N. Henze, T. Dingler, M. Pielot, D. Weber, and A. Schmidt, “Large-scale assessment of mobile notifications,” Proc. 32nd annual ACM conference on Human factors in computing systems - CHI ’14, 2014.

[33] V. J. Sauer, S. C. Eimler, S. Maafi, M. Pietrek, and N. C. Krämer, “The phantom in my pocket: Determinants of phantom phone sensations,” Mobile Media & Communication, vol. 3, no. 3, pp. 293–316, Jan. 2015.

[34] M. H. Vastenburg, D. V. Keyson, and H. de Ridder, “Considerate home notification systems: A user study of acceptability of notifications in a living-room laboratory,” International Journal of Human-Computer Studies, vol. 67, no. 9, pp. 814–826, Sep. 2009.

[35] A. Voit, B. Poppinga, D. Weber, M. Böhmer, N. Henze, S. Gehring, T. Okoshi, and V. Pejovic, “UbiTtention,” Proc. 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing Adjunct - UbiComp ’16, 2016.

[36] D. Weber, “Notifications in a multi-device environment”. Diploma Thesis. Institute for Visualization and Interactive Systems, University of Stuttgart, Germany. 2015.

[37] D. Weber, A. S. Shirazi, S. Gehring, N. Henze, B. Poppinga, M. Pielot, and T. Okoshi, “Smarttention, please!, 2nd Workshop on Intelligent Attention Management on Mobile Devices” Proc. 18th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct - MobileHCI ’16, 2016.

[38] J. Wiese, T. S. Saponas, and A. J. B. Brush, “Phoneprioception, Enabling Mobile Phones to Infer Where They Are Kept” Proc. SIGCHI Conference on Human Factors in Computing Systems - CHI ’13, 2013.