

Job Alerts in the Wild: Study of Expectations and Effects of Location-based Notifications in an Existing Mobile Crowdsourcing Application

Delphine Reinhardt
University of Bonn and
Fraunhofer FKIE
Bonn, Germany
delphine.reinhardt@cs.uni-
bonn.de

Martin Michalak
Secure Mobile Networking Lab
TU Darmstadt
Darmstadt, Germany

Robert Lokaiczkyk
wer denkt was GmbH
Darmstadt, Germany
lokaiczkyk@appjobber.com

ABSTRACT

Mobile crowdsourcing applications leverage volunteers to collect information using their personal devices. It is hence vital to foster the volunteers' engagement and their contributions to ensure the long-term viability of these applications. A method to reach this goal is to notify participants about new tasks in their physical proximity. Such location-based notifications can however impact the users' experiences, disclose their location, and/or incur additional resource consumption. In this paper, we therefore investigate the potential worthiness of introducing location-based notifications in the wild. Our analysis includes the perspectives of both potential users and campaign managers. To this end, we have conducted two questionnaire-based studies counting 335 participants in total. By doing so, we gain insights about the participants' expectations and the value attributed to invested resources. We further identify significant factors that influence their readiness to activate this new function. Finally, we measure the impact of job alerts and rewards on both the quantity and quality of users' contributions. Consequently, our findings can help campaign administrators in the design of future crowdsensing applications.

CCS Concepts

•**Human-centered computing** → Collaborative and social computing systems and tools; Empirical studies in collaborative and social computing; Empirical studies in ubiquitous and mobile computing; *Ubiquitous and mobile computing systems and tools*; •**Security and privacy** → Economics of security and privacy;

Keywords

Mobile crowdsourcing, location-based notifications, field and user studies

1. INTRODUCTION

The number of worldwide mobile subscriptions has been estimated to reach 7.377 billion at the end of 2016 [12]. This ubiquity and the increasing number of sensors embedded in mobile phones have led to the emergence of new sensing paradigms. Different terminologies, such as crowdsensing and participatory sensing, have been introduced, but all of them share a common denominator: citizens leverage their personal devices to contribute user-generated sensor readings [4, 6]. Based on this concept, a myriad of applications have been proposed whose purposes range from people-centric to environment-centric scenarios [7]. Among them, we herein focus on a mobile crowdsourcing application called *appJobber*¹. This application is detailed in Sec. 3 and essentially proposes different participatory tasks to be fulfilled in exchange of monetary rewards.

However, *appJobber* and all crowd-based applications rely on the so-called network effect, thus highly depending on the active utilization and contributions of their potential users [18, 21]. In presence of a multitude of these applications, it becomes more and more difficult to first attract and then maintain a solid user base. As a result, many applications have not yet crossed the chasm and are still in early adoption phase missing large-scale deployments [4, 17, 22]. A possible explanation is that personal resources invested by users and the corresponding return of investment are not balanced [16, 25]. For example, users spend time to search for new tasks and/or reach tasks bounded to particular locations.

An opportunity to mitigate this problem is to notify users about nearby tasks. Thus, our objectives are two-fold. We aim at reducing the efforts needed by the users to, e.g., search for new tasks and reduce the geographic distance they need to cover in order to reach the task's location, while simultaneously fostering contributions to the crowdsourcing campaign. However, location-based alerts rely on continuously probing the user's position and thus incur additional energy consumption and disclose the users' location to the application server. In this work, we analyze and quantify whether the benefits drawn from these notifications are worth the incurred drawbacks in terms of additional resource consumption and reduced location privacy. Our contributions can be summarized as follows:

¹www.appjobber.de

1. We implement and deploy a location-based notification function called *job alerts* in appJobber. While we acknowledge that location-based notifications are already applied in other applications, we aim at understanding and measuring their effects on a representative sample of users having experienced the application under realistic conditions and contributing to a real-world application not especially designed for the purpose of our work. To the best of our knowledge, we are the first to investigate these aspects under these conditions.
2. We conduct a preliminary questionnaire-based study involving 131 participants with the help of which we verify our design decisions and analyze the participants' expectations about this new functionality.
3. We then conduct a follow-up questionnaire-based study with 204 participants coupled with an analysis of the appJobber database. In this second study, we analyze how our participants value the resources they invest in appJobber. We further investigate potential factors influencing users' decision to activate/deactivate these notifications. Besides, we quantify the impact of the job alerts as well as the value of the rewards on both the quantity and quality of users' contributions. As a result, our findings can help future researchers/developers in designing their crowdsourcing applications.

The remaining of this paper is structured as follows. After summarizing existing work in Sec. 2, we revisit the key features of appJobber in Sec. 3. We detail the settings and results of our preliminary and validating studies in Sec. 4 and 5, respectively. We further discuss the limitations of our results in Sec. 6 and make concluding remarks in Sec. 7.

2. RELATED WORK

Our analysis is divided into two key domains addressed in this paper. We first consider studies analyzing potential factors aiming at fostering users' contributions and their effects on the application utility, before focusing on studies dedicated to the disclosure of location information.

2.1 Fostering User-generated Contributions

Depending on the application scenarios, the motivations to contribute user-generated content can vary [4, 7]. For example, users can be interested in quantifying their personal experience. Alternatively, they can contribute by altruism or be motivated by competition aspects, such as gaining virtual rewards [22]. Monetary incentives can also be introduced to reward contributing participants like in appJobber, the crowdsourcing application considered in this paper. We hence focus on this motivation factor in what follows.

In mobile crowdsensing, several work address the optimization of the reward distribution, such as [17], with regards to the contribution quality, the expected coverage, and the delay between the tasks' publication and fulfillment [4]. In addition to propose new optimization schemes, further works analyze how different monetary-incentive modalities influence the users' contribution rate and their behavior in participatory sensing data collection. It has been shown in [21] that users fulfill more tasks when those are individually paid as compared to a flat rate for multiple tasks.

Moreover, monetary incentives have an impact on the users' behavior, as those are ready to spend additional time in specific locations or visit additional locations when necessary. In the context of crowdsourcing and especially Amazon Mechanical Turk, further studies have demonstrated that the greater the reward for a task, the faster it is fulfilled, but the worse the quality [10, 11, 19].

Our work shares similarities with [1], in which the authors have investigated location-based crowdsourcing and shown that users are prevalently interested in finding tasks that can be fulfilled at or close to their current location. However, their results indicate that users prefer searching for new tasks themselves than getting suggestions from the application. Note that no monetary incentives are given to the participants to fulfill tasks in the case study presented in [1]. Our work builds upon these results and aims at verifying how users assess the advantages gained by location-based notifications vs. the inherent drawback in terms of battery lifetime and location disclosure in practice. Moreover, we aim at quantifying the impact of the introduction of these notifications on the application in presence of monetary incentives.

2.2 Disclosure of Location Information

Existing works on location privacy can be classified according to two main categories. The goal of the first category is to analyze factors influencing users' privacy concerns about revealing location information. For example, the results published in [3, 8, 23] show the impact of gender on the users' privacy concerns; women are more concerned than men about their location privacy. Additionally, the users' privacy concerns also depend on their degree of comprehension of the underlying technology as shown in [15]. The study conducted in [20] further shows that users are aware of potential risks to their privacy and 25% have already regretted to have shared their location at least once. In another study published in [5], we have explored multiple factors that may influence users to contribute privacy-sensitive data to participatory sensing. Consequently, we leverage these studies to investigate whether their findings also apply in our scenario.

The second category of works focuses on assessing the value of location information based on user studies. In [9], the participants indicated that they would be ready to provide their monthly location data for 15 Euros for academic purposes and 29 Euros for commercial purposes in average. The same ratio has been found in [8]: 43 and 86 Euros, respectively. The results further show that these values do not proportionally increase with the collection duration. Both studies are however based on questionnaires and targeted students in computer science. In contrast, a two-month deployment was conducted in [3] as support to the valuation of location privacy. In this study, both academic and commercial values were the same in average: 78 Euros for one month if preliminary anonymized and 118 Euros if not. As compared to the aforementioned studies that focused on continuous location information, the study conducted in [2] focuses on discrete location data. In a real-world deployment, the participants were asked daily in which location they currently were and whether they liked to share it for a coupon in a coffee shop. The coupon's value was randomly chosen between 1 and 20 Euros. In average, the participants in this study were willing to share their home location for

8 Euros, their work location for 5 Euros and the remaining locations for 3 Euros. Note that both the publication of the participants' location data and the associated rewards were fictive.

In comparison with these works, we use an existing crowdsourcing application that is not specifically designed for the purposes of our studies and count experienced users with various profiles. Similarly, we aim at analyzing potential factors that motivate participants to activate/deactivate job alerts and how these participants value the transmission of their current information to the application server. To this end, we carefully design and embed our questions, so that we do not artificially raise the participants' awareness about potential privacy issues. Moreover, we investigate the privacy value by providing a reference in order to mitigate the limitations of open scales observed in previous works.

3. APPJOBBER

As baseline for our contributions, we have leveraged appJobber, which is an application developed and managed by *wer denkt was GmbH*². AppJobber is a crowdsourcing application that serves as an interface between clients and users. Clients define tasks called *jobs* to be fulfilled by potential users called *jobbers*. Note that appJobber is not a research prototype especially implemented for this study, but a fully functioning crowdsourcing application. AppJobber counts more than 300,000 registered jobbers deployed across multiple European countries fulfilling tasks for more than 400 clients, such as TomTom, Sony, and Deutsche Bahn. Using their mobile phones, jobbers can browse and temporarily reserve jobs to be completed. These jobs are classified based on (1) their type, (2) the associated subtask(s), and (3) the reward. The job type is defined by the jobbers' location requirements. Jobs can be location-specific, region-specific, or independent of the current jobbers' location. An example of location-specific jobs is taking a picture of a given speed sign. Depending on the job types, the associated subtasks are *questions*, *pictures*, *waypoints*, and *links*. In the first subtask category, jobbers must enter the answer to a given question and take a predefined picture in the second category. In *waypoint* subtasks, the jobbers' need to follow a particular path which is sent to the application server, while jobbers need to open hyperlinks (e.g., websites or other apps) in *links* subtasks. A job can be composed of one or more subtasks. The results of the subtask(s) are then transmitted to the application server where they are individually verified. Depending on their quality, the results are either validated and the associated reward is paid or discarded.

3.1 Introduction of Job Alerts in AppJobber

AppJobber originally allows jobbers to search and reserve jobs by browsing them on a map. Within the scope of this work, the underlying mobile application has been extended by a location-based notification function referred to as *job alerts*. By implementing it, the objective is to further reduce jobbers' efforts by reducing the time needed to find new jobs and increasing their exposure to jobs in physical proximity, thus reducing the distances to be covered. As a result, we expect an increased job fulfillment rate.

To cater for the jobber's acceptance, the notification service should however have the lowest impact possible on the

user experience. This means that jobbers should not be overwhelmed by notifications and the resource consumption incurred by the location-based service should be minimized. Otherwise, the opt out rate may increase [13, 14]. Simultaneously, the location information should be accurate enough to point out relevant jobs to potential jobbers in proximity.

By taking into account these requirements, we have implemented and evaluated different solutions based on the location services available on iOS phones. Among the existing solutions, we have chosen to apply the *significant-change* option. With this option, a request is sent to the application server only if the distance between two successive locations reaches a given threshold. The application server replies with potential jobs located within a radius of 500 m around the jobber's current location, i.e., approximately reachable within a five minute walking distance. To prevent jobbers from getting annoyed by too many job alerts, we have chosen to limit the number of notifications to one every two hours between 7.00 am and 9.00 pm that corresponds to the core timespan during which jobbers usually complete jobs. As a result, a jobber daily receives at most eight notifications. Using these configurations, a first version of job alerts has been integrated in the iOS version of appJobber and deployed. At the first application start, jobbers have been informed about this new function and could directly disable it. Alternatively, they could later deactivate it in the application menu. Additional information about the job alert functionality is available on the appJobber website, including details about its functioning and its impact on the phone's battery lifetime.

4. PRELIMINARY QUESTIONNAIRE

Simultaneously to the release of the new iOS version including job alerts, we have published a new job including different questions to validate our design decisions detailed in Sec. 3.1. No monetary incentives were provided. The questionnaire was written in German and was available to all registered users in Germany, Austria, and Switzerland having installed the latest release. In total, 131 jobbers answered the questionnaire, including 119 who had activated the job alerts and 12 who had deactivated them. We submitted a specific set of questions to the jobbers depending on the group they belonged to. As a result, we first present and comment on the answers given by the jobbers having activated the location-based notifications, before considering those who have deactivated them.

4.1 Activated Job Alerts

In this questionnaire, we first indicated to the jobbers that we needed their feedback to further improve the application, before submitting them the following questions related to the configuration of job alerts.

4.1.1 Maximal Number of Daily Alerts

We firstly asked the jobbers to indicate an upper bound to the daily number of job alerts they would like to receive. As a result, around 76% of the participants would prefer to receive fewer than 8 alerts. The remaining indicated to be ready to get 8-9 daily notifications. To satisfy most participants, the number of received alerts should however be lower than the limit configured in our implementation. Nevertheless, we expect that participants will receive fewer than 8 daily notifications in practice, as their mobility pattern

²<http://wordpress.werdenktwas.com>




Figure 1: Distribution of the participants’ answers to the question: “Indicate the duration of a walk (in minutes) you would be ready to cover to reach and complete a job” ($n=92$)

may not trigger a new request to the application server every hour. Moreover, four participants indicated that they would like to get as many job alerts as possible, while another participant wishes to have personalized alerts based on the reward values.

4.1.2 Distance to Jobs

We next asked the jobbers to indicate the duration of a walk (in minutes) they would be ready to cover to reach and complete a job. The most selected answer was 5-10 min chosen by 35% of the participants as illustrated in Fig. 1. The participants’ answers illustrate the differences between the jobber profiles and their readiness to invest efforts in fulfilling jobs. By choosing a short duration (and hence a small radius around the jobbers), we logically include most participants. A time duration between 5-10 min hence appears as appropriate. In the free-text field, one participant indicated that (s)he would even be ready to walk 30-45 min to reach a job. Other participants answered in terms of distance: 5 km (one participant), 10 km (two participants), and 20 km (one participant). A participant also indicated his/her personal reward-based conditions: (s)he would walk up to 5 min for a job with a reward lower than 2 Euros, while (s)he would be ready to invest more time for a greater reward.

4.1.3 Resources Consumption

To balance the tradeoff between frequent location updates and energy consumption, we asked the participants to indicate the maximal percentage of battery lifetime per hour they would agree consuming when using appJobber. Around 54% of the participants would like to remain below 4% of hourly battery consumption, while 23% agreed with 9-10%. A participant noted that “[s]he rarely uses this function because the GPS consumes much energy”. We were also interested in knowing how often Wi-Fi is enabled on the participants’ phone. Around 55% have it always enabled, while 16% enable it often, 13% occasionally, 11% rarely, and 6% never activate it.

As expected, the participants’ answers illustrate the differences in jobber profiles, ranging from occasional to frequent jobbers. Particularly, we show that the greater the number of job alerts they are ready to receive:

- The longer they are ready to walk ($r(104)=0.218$, $p<0.050$ with r the Pearson’s coefficient),
- The greater loss in battery lifetime they are ready to accept ($r(108)=0.433$, $p<0.001$), and




Figure 2: Distribution of the participants’ answers to the question: “Which are the reasons why you do not want to use the job alert function?” (multiple choice possible, $n=12$)

- The higher the expected reward ($r(99)=0.296$, $p<0.001$).

A significant positive correlation is also shown between the participants’ answers regarding the distance to jobs and the battery consumption ($r(110)=0.366$, $p<0.001$), the distance to jobs and the expected reward ($r(102)=0.281$, $p<0.001$), as well as the battery consumption and the expected reward ($r(107)=0.274$, $p<0.001$). Consequently, the gained insights do not only allow us verifying our design decisions, but may also be useful for the development of future crowdsourcing applications.

4.2 Deactivated Job Alerts

Next, we analyze the reason(s) why jobbers had deactivated job alerts. Fig. 2 illustrates the reasons selected by the twelve participants of our sample. The main reason selected by these participants is that they do not want to be disturbed by the alerts followed by the additional battery consumption incurred by the localization function. While we intentionally not included the risks to their privacy as a potential reason, three participants indicated this aspect in the miscellaneous field: “I simply do not want that my location is continuously sent”, “Do not want to be continuously localized even if job alerts is deactivated”, and “For me, it is more location monitoring that I do not like”.

5. QUESTIONNAIRE-BASED STUDY AND VALIDATION

In a second step, we build on the results detailed in Sec. 4 and conduct a second questionnaire-based study. Our objectives are three-fold. We first investigate how participants value invested resources when fulfilling jobs in Sec. 5.2, before examining factors potentially influencing the participants’ decisions to activate/deactivate job alerts in Sec. 5.3. We finally quantify the impact of job alerts and the reward value on: (1) the time between the publication of new jobs and their completion, (2) the distance covered to complete jobs, (3) the quality of the job results, and (4) the application utilization. To this end, we formulate different hypotheses compiled in Tab. 1 and verify them in Sec. 5.4.

5.1 Study Settings

In this second study, the questionnaire is available to both Android and iOS jobbers registered in Germany and can be answered as a location-independent job. The questionnaire counts at most 22 questions, The questions depend on

Table 1: Considered hypotheses with regard to the impact of job alerts and reward values on appJobber

| | |
|--------------|---|
| Hypothesis 1 | Job alerts reduce the time between the job publication and its reservation by jobbers |
| Hypothesis 2 | The greater a job reward, the faster it is reserved |
| Hypothesis 3 | Job alerts reduce the distance to jobs |
| Hypothesis 4 | The greater a job reward, the larger the distance covered by the jobbers |
| Hypothesis 5 | Job alerts do not impact the quality of the jobbers' contributions |




Figure 3: Ratio of completed jobs after notification compared to all completed jobs in T_{post} ($Q_2=16\%$)

whether the jobbers have enabled job alerts and are clustered into four categories: (1) job alerts, (2) appJobber experience, (3) knowledge and (4) social network utilization. For answering the questionnaire, jobbers are rewarded by 1 Euro. In the knowledge section, we especially ask the participants several questions to quantify their technical knowledge about location-based technologies as well as their understanding of potential risks to their privacy. The jobbers' answers are completed by an analysis of the appJobber database during three periods. The first period referred to as T_{ante} includes the 11 weeks prior to the release of the job alerts, while the 11 weeks after the release are referred to as T_{post} . T_{year} ends with T_{post} and includes all retrospective data during one year. In T_{ante} and T_{year} , the corresponding databases contain the reward, the times between publication, reservation and completion, the distance to the job (i.e., shortest distance between the jobber's location at the reservation and the job), and the result status (i.e., validated or discarded). In T_{post} , an additional entry indicates whether the jobber has been notified about this job using a job alert. The rewards are classified into three categories: small (1 Euro), medium (1-4 Euros), and big (>4 Euros). As reference, Fig. 3 shows the percentage of completed jobs notified by job alerts as compared to all completed jobs during T_{post} . Note that no artificial jobs have been created for the purpose of our study.

5.1.1 Demographics and Sharing Behavior

In total, 207 jobbers answered our second questionnaire. We discarded the answers of three of them based on their abnormally short fulfillment duration. In what follows unless noted otherwise, we hence consider a sample of 204 jobbers equally balanced between Android and iOS users. Our sample counts 66% of male that corresponds to the gender distribution observed in appJobber. A majority of participants (43%) are between 20 and 29 with 5% over 50 (quartiles $Q_1=24$, $Q_2=29$, $Q_3=36$). Their education level is almost balanced between secondary school (40%), A-levels (30%) and university studies (25%). Among those having studied or studying, the most represented fields are engi-

neering (25%) followed by economics and law (17%), computer science (16%), and social sciences (12%). In our sample, over 73% are employed, while 17% are students. Their working fields include: IT (16%), commerce (12%), manual skills (14%), customer services (10%), transport and logistics (8%), as well as health and social (8%). In addition to their demographics, we were interested in gaining insights about their online sharing behavior. The results show that a large majority (87%) are registered and use online social networks. Among them, 46% indicated to post personal information online at least once a week, while 36% even share it several times a week. Only 2% claim not to intentionally publish any personal data online.

5.1.2 AppJobber Experience

Most of our 204 participants (85%) have already completed at least one job. Moreover, 87% actively look for new jobs by, e.g., consulting the map available in appJobber. To this end, 56% consult appJobber on a daily basis (including 34% who do it several times a day), while 39% do it on a weekly basis. The remaining participants do it less frequently. To still reduce the burden associated to the job search, 80% indicated to be in favor of getting notifications about new jobs. In addition to invest time and resource to look for jobs, the participants are also ready to cover more distance to complete jobs. Only 14% of the participants indicated to have completed 80-100% of their jobs with a maximum detour of 5 min. The remaining 68% indicated a lower percentage of jobs involving such a short detour, thus suggesting that they experience longer detours more often. The results hence confirm the participants' readiness to invest resources to search and complete jobs.

5.2 Reward for Invested Resources

In addition to gain insights about their experience, we are interested in inferring the reward value that appJobbers are expecting in return for their efforts. Instead of directly asking the participants to give a numerical value, we submitted them a scenario based on an existing 1-Euro job, which aims at checking traffic speed limits in a particular area. To this end, a jobber would need to fulfill three subtasks: (1) walk between two waypoints, (2) take two pictures at both waypoints, and (3) manually enter the corresponding speed limit of this section. Note that the location of both waypoints is recorded by the application. We then asked the participants to indicate how many cents of the original reward (i.e., 1 Euro) they would attribute to each subtask with a 5-cent granularity.

The participants' answers are compiled in Fig. 4 and the respective extrema and quartiles are summarized in Tab. 2. A Friedman test shows a significant difference in the part of the reward attributed to each subtask ($\chi^2(2)=114$, $p=0.000$). A post hoc analysis with Wilcoxon signed-rank tests with a Bonferroni correction leads to a significance level of $p<0.017$.




Figure 4: Distribution of the 1-Euro reward partitioned between the three given subtasks ($n=201$)

Table 2: Reward extrema and quartiles attributed to each subtask by the participants ($n=201$)

| | Walking between waypoints | Taking pictures | Entering the answer |
|------------|---------------------------|-----------------|---------------------|
| <i>min</i> | 5 | 5 | 5 |
| Q_1 | 25 | 25 | 10 |
| Q_2 | 40 | 40 | 20 |
| Q_3 | 50 | 50 | 30 |
| <i>max</i> | 90 | 80 | 70 |

As a result, there is no significant difference between the partial reward attributed to both the *walking to waypoints* and *taking pictures* subtasks ($Z=-0.373$, $p>0.050$). In contrast, this difference is significant between the *walking to waypoints* and *entering the answer* subtasks ($Z=-8.33$, $p=0.000$) and the *taking pictures* and *entering the answer* subtasks ($Z=-9.19$, $p=0.000$). This means that participants attribute a similar reward to both the *walking to waypoints* and *taking pictures* subtasks, while the attributed reward is lower for manually entering the answer. The comparable rewards between *walking between waypoints* and *taking pictures* are surprising as we would have expected that the participants would attribute a higher reward to the walking subtask. Indeed, it may require more time and physical efforts than taking pictures. Moreover, it involves the disclosure of the current participants’ locations to the application server. As a result, the participants estimate the value of walking between two waypoints and providing their locations to 40 cents in average—a value significantly below the ones observed in existing surveys detailed in Sec. 2.2.

When further considering the reward attributed to the *walking to waypoints* subtask, a Mann-Whitney-U test shows that this value does not significantly differ depending on the participants’ gender ($U=4120$, $z=-1.04$, $p>0.050$), their knowledge about potential risks to privacy ($H(3)=0.257$, $p>0.050$), and age ($H(4)=9.00$, $p>0.050$). However, we demonstrate that the participants’ education level has a significant influence on the attributed value ($H(3)=19.9$, $p<0.001$). A post-hoc test following the method of Dunn with Bonferroni correction confirms significant differences between participants with the highest education level and those with both the intermediary and lowest levels ($p<0.001$, respectively). Surprisingly, participants with the highest education level indicated a value lower for the reward attributed to the *walking between waypoints* subtask than the remaining participants.




Figure 5: Distribution of the answers of the iOS of our first sample and the Android users of our second sample about the additional reward they are expecting to gain when using job alerts ($n=119$ for iOS users (first sample) and $n=92$ for Android users (second sample))

We finally investigate the additional gain that our participants are expecting when activating the job alert function as compared to their current gain. Fig. 5 compiles the answers of the participants of our preliminary study detailed in Sec. 4, i.e., 119 iOS users, and 92 Android users contributing to this second study. In order to limit the number of questions to the minimum and hence reduce the participants’ overhead, we have decided not to include this question in the second questionnaire dedicated to the iOS users. Instead, we have preferred to focus on their practical experience of job alerts. The results show that 27% of the Android users in our second study and 13% of the iOS users in our first study do not expect any increase of their rewards from location-based notifications. In addition, the iOS users of our first study are significantly expecting more additional gain than the Android users of this study ($U=4016$, $z=-3.36$, $p<0.001$, $r'=-0.226$, with r' the effect size). Overall, the participants’ expectations remain low and the majority of them do not expect to earn more than half of their current reward.

5.3 Activation/Deactivation of Job Alerts

We next analyze the influence of the participants’ demographics on the jobbers’ decision to activate/deactivate job alerts. To this end, we distinguish iOS and Android users, as iOS users had the possibility to test this function in practice and their answers to the questionnaire can be completed by an analysis of the appJobber database. We further analyze the experience of the iOS users with the job alerts.

Gender.

Among the iOS users, a database analysis indicates that more female jobbers have activated the job alert function

as compared to male jobbers. A Chi-Square test demonstrates that the gender difference is statistically significant ($\chi^2(1)=11.9, p<0.001$). This result however differs from our expectations based on [15], which shows that women are more concerned about their location privacy than men. A possible explanation is that female jobbers are less informed about possible implications of job alerts. Indeed, significantly fewer female jobbers have visited the information menu as compared to male jobbers ($\chi^2(1)=6.65, p=0.001$). For the participants having not yet activated the function (i.e., Android and remaining iOS users), the same trend is observable in their answers. However, the gender difference is in both cases not statistically significant ($p>0.050$).

Field of Activity.

For Android users, the participants' field of activity makes a significant difference. Participants working in the area of electronic data processing, information technology, and e-business are significantly less ready to activate the job alert function than others ($U=944, z=2.19, p<0.050, r'=-0.217$). This however does not hold for iOS users where the difference is not statistically significant.

Knowledge about Potential Risks to Privacy.

As expected, the answers of the participants show that Android users with knowledge about potential risks for their privacy are significantly less ready to activate it as compared to others ($U=691, z=-2.32, p<0.050, r'=-0.229$). This result is however not valid in the group of iOS users.

In contrast, the participants' age, education level, education background, knowledge about location-based technologies, and contribution in online social media do not significantly impact their readiness to activate the job alert function in both Android and iOS groups.

In summary, we observe the same trends between Android and iOS users about factors potentially influencing their decision to activate or deactivate jobs. However, the impact of the gender, the field of activity, and the knowledge about potential risks to privacy can only be statistically validated for Android users. The fact that iOS users have been able to experience job alerts in practice may explain this difference.

5.3.1 Job Alerts Experience

We now focus on the iOS users having experienced job alerts in practice. In this group, 70% have enabled the function, while 14% have temporarily activated it. In contrast, 7% have directly deactivated it and 7% did not know about this function. As shown in Fig. 3, around 50% of the iOS participants have already received at least one job alert as shown by an analysis of the database. Among them, the number of received alerts ranges between 1 and 83 ($Q_1=2, Q_2=7, Q_3=10$). Figs. 6 and 7 show the advantages and drawbacks perceived by the participants having already experienced job alerts, respectively. Overall, the participants' answers are in line with our expectations and the results of the preliminary study detailed in Sec. 4. It however seems difficult to fully satisfy the jobbers with regards to the number of received job alerts. In our preliminary study, only a minority of participants would be ready to get more than 8 alerts per day. While the participants of our second study have actually experienced fewer alerts, 46.5% indicated that they did not receive enough alerts. Even if the number of alerts is directly determined by both the job and jobber lo-




Figure 6: Distribution of the participants' answers to the question: "Why do you like job alerts?" (multiple choice possible, $n=86$)




Figure 7: Distribution of the participants' answers to the question: "Why do you dislike job alerts?" (multiple choice possible, $n=86$)




Figure 8: Distribution of participants' answers about the reasons why they have disabled job alerts (multiple choice possible, $n=17$)

ocations, an additional option should be introduced so that jobbers can tailor the maximal number of alerts to their preferences. As shown in Fig. 8, the most cited reason why iOS have not enabled job alerts is that they do not want to be disturbed followed by their impact on the battery lifetime. This is in line with the answers given in our previous study. As compared to iOS users, Android users have not tested the function, but 90% indicated their potential readiness to enable this function, while 8% would deactivate it.

5.4 Impact on the AppJobber Application

We further investigate the impact of the job alerts and/or the value of the rewards from the application perspective. To this end, we consider the hypotheses formulated in Tab. 1 and we focus on the participating iOS users having already completed a job after having received a job alert, i.e., 30% of all iOS users ($n=26$). Moreover, we take into consideration their answers to our questionnaire and/or the database entries during T_{ante} and/or T_{post} (i.e., before/after the in-

roduction of the job alerts). We complete these results by an analysis of all jobbers' entries during one-year T_{year} .

Hypothesis 1 *Job alerts reduce the time between the job publication and its reservation by jobbers*

Both the participants' answers and an analysis of the appJobber database do however not support this hypothesis. Indeed, a majority (54%) believes that job alerts do not lead to shorter reservation time, while 38% believe the opposite. The remaining are undecided. We further compare the time to reservation between regular jobs and notified jobs during T_{post} . During this period, the median time is 90 hours for regular jobs and 119 hours for notified ones. A Mann-Whitney-U test confirms that the difference is statistically significant ($U=219192$, $z=3.23$, $p<0.001$, $r'=-0.079$), meaning that the time to reservation for notified jobs is longer than for regular ones. We however observe a significant reduction in the reservation time between T_{ante} and T_{post} ($U=119475$, $z=-2.84$, $p<0.001$, $r'=-0.081$). Since some participants indicated that they are using appJobber more frequently due to the alerts, a potential explanation could be that job alerts have hence contributed to a time reduction for regular jobs. The same effect is however also observable for Android users. As a result, this reduction cannot be attributed to the job alerts. This hypothesis is hence rejected.

Hypothesis 2 *The greater a job reward, the faster it is reserved*

Independently on job alerts, the median time is 521 hours for small rewards, 101 and 87 hours for medium and big rewards, respectively. A Kruskal-Wallis test shows that there is a significant difference in the reservation time depending on the reward categories ($H(2) = 2230$, $p<0.001$). A post-hoc test shows that the difference is significant between all tested pairs of reward categories ($p<0.001$). As a result, the reservation time decreases when the reward increases, thus confirming this second hypothesis.

Hypothesis 3 *Job alerts reduce the distance to jobs*

53% of the participants having completed a notified job agree with this hypothesis. 8% disagree and the remaining are undecided. Again, the participants' answers reflect the results of the database analysis. The median distance covered for notified and regular jobs is 136 m and 235 m, respectively. A Mann-Whitney-U test confirms our third hypothesis ($U=731750$, $z=-5.98$, $p<0.001$, $r'=-0.102$).

Hypothesis 4 *The greater a job reward, the larger the distance covered by the jobbers*

We examine the appJobber database during T_{year} to verify it. The median distance covered for small rewards is 177 m, 463 m for medium rewards, and 1,065 m for big rewards. A Kruskal-Wallis test shows a significant difference between the considered reward categories on the covered distance ($H(2) = 225$, $p<0.001$). Each category pair is significantly different ($p<0.001$), thus validating our fourth hypothesis.

Hypothesis 5 *Job alerts do not impact the quality of the jobbers' contributions*

A database analysis shows that 94% of the results of regular jobs are accepted, i.e., their quality is sufficient to be

exploited, while the acceptance rate is 95% for jobs following a job alert. A Chi-Square test for independence confirms our last hypothesis ($\chi^2(1)=0.705$, $p>0.050$). The quality of the uploaded contributions however depends on the reward categories ($\chi^2(12)=91.120$, $p<0.001$). Indeed, the higher the reward, the lower the acceptance rate.

Additionally, the introduction of job alerts has a positive impact on the appJobber utilization as reported by the participants: 30% of the iOS users indicated in the study that they are using appJobber more often, while 25% reported to also fulfill more jobs.

In summary, we have shown based on the answers of our participants coupled with the database analysis that job alerts significantly reduce the distance to jobs and have no significant influence on the quality of the participants' contributions. Besides, we have confirmed that the reward value impact both the reservation time as well as the quality of the contributions.

6. DISCUSSIONS

We first comment on the observed effects of job alerts and the expected reward for invested resources, before discussing the limitations of our contributions.

6.1 Positive Effects of Job Alerts

As expected, the introduction of job alerts has a positive impact on the user participation. Our participants indicated to both utilize appJobber and complete jobs more often. These results are in line with the results presented in [24]. Moreover, we have verified the impact of reward values based on a existing real-world crowdsourcing applications. We have confirmed that higher rewards decrease the observed delay between the job publication and its execution and also motivate jobbers to cover larger distances. High rewards do however not foster the collection of high-quality results. This is in line with the results obtained for Amazon Mechanical Turk [10, 11, 19]. In our case, a potential explanation is that these jobs may be more difficult to complete. Alternatively, jobbers may be attracted by the rewards, but tend to minimize the efforts necessary to correctly complete these jobs. Moreover, manipulated contributions are observed more frequently with higher rewards. This means that higher rewards are overall beneficial to the application, but require more control.

Our results however dispute the results obtained in [1] showing that users prefer searching manually for new jobs instead of being automatically notified based on their current location. In our study, 92% of our iOS users have activated the job alert function and more than 90% of the Android users claimed to be ready to do it. These numbers hence confirm the interest of the users in getting additional information about jobs to complete through this channel. The difference between these results may be due to (1) the size of the sample (204 in our study vs. 9 in [1]), (2) the evaluation platform used (an existing fully-functioning application in our case vs. a synthetic application created for the study purpose in [1]), as well as (3) the monetary reward provided to the participants (1 Euro in our case). Moreover, the activation rate we have observed is higher than the one usually observed (50%) when considering applications including location-based push notifications according to [24]. This may be due to the advantages perceived by the participants who experienced shorter distance to cover in order

to reach and complete jobs as shown by their answers and confirmed by an analysis of the appJobber database and as compared to regular jobs. Additionally, our participants indicated that they would expect to save time when searching for new jobs.

6.2 Reward for Invested Resources

We have shown in Sec. 5.2 and especially in Tab. 2 that our participants value the subtask *walking between two waypoints* to 0.40 Euros in average when considering a given 1-Euro job. The value of this subtask hence includes the reward for covering the distance between the waypoints and also providing the corresponding locations to the application server. Moreover, some of our participants indicated to expect to gain additional rewards when activating the job alert function. When converting their answers to a monetary value based on their previous rewards, we find that this corresponds to an additional monthly reward of 1.49 Euros in average. We believe that the jobbers' expectations are principally based on their assumption that job alerts may increase their awareness about new jobs and lead to a higher completion rate. Another reason could be that the participants need to invest more resources, such as providing their location information to the application server, and hence expect to be paid more for these tasks. As already highlighted in Sec. 5.2, both our values are lower than those observed in previous studies. In the case of discrete location information, Barak et al found in [2] a value between 2.90 and 8 Euros. In contrast, the monthly collection of continuous location information for commercial purposes was valued 29 Euros in [9], 78 Euros in [3], and 86 Euros in [8]. Again, the differences may be explained by the applied evaluation methods and also by the fact that participants may complete more than one job monthly. Moreover, in existing studies, the participants could indicate any numerical values for the continuous collection of their location information, while we asked our participants to indicate the percentage of the initial 1-Euro reward they would attribute to the location-based subtask. As a result, the participants' answers are given according to a reference, but are obviously bounded by 1 Euro that corresponds to the main reward value used in appJobber. Moreover, we have conducted our questionnaire-based surveys with actual users of an existing application, thus catering for real-world conditions. Since our questionnaires have not focused on privacy aspects, we may also not have artificially increased the participants' privacy awareness, which may have led to higher monetary expectations. Another potential explanation is that the participants better understand why their location is needed as they were able to experience appJobber under real conditions over an extended period of time.

6.3 Limitations

This work is based on two samples of participants recruited within appJobber. Hence, the participants have not installed the application for the purposes of the study, but have been engaged earlier based on their own interests. As a result, our participants have already experienced appJobber under real-world conditions as highlighted in Sec. 5.1.2. This hence contributes to the representativeness of our results. Nevertheless, the deployment of our studies in the wild has introduced additional constraints as compared to studies conducted in controlled environments. Both our question-

naires as well as our implementation should not lead to any user opt-outs. We have therefore designed our studies so that their impact on both the users and the application are reduced to the minimum. For example, we have limited the number of questions and questionnaires submitted to the participants as well as the number of job alerts to maintain a satisfactory user experience without draining the phones' battery. Additionally, we have attempted to remain neutral and conservative in the choice of our questions.

Both our samples do not contain all appJobber users. This might thus result in a bias in our results, as our samples may include users with a larger readiness potential in contributing personal resources as compared to others. We have however shown that our samples contain different jobber profiles.

As in all questionnaire-based studies, most of our results are based on participants' statements and might hence be affected by the usual limitations of such studies, in which the participants' answers and, e.g., their actual behavior might diverge. Nevertheless, we have compared when possible the participants' answers to the corresponding database entries and observed the same trends in both information sources.

7. CONCLUSIONS

We have leveraged the introduction of location-based notifications about nearby jobs in an existing mobile crowdsourcing applications to study the users' expectations, the factors leading to their acceptance, and their impact on the application utility. To this end, we have conducted two questionnaire-based studies involving 335 appJobbers in total. Among the studied factors, our results show significant differences between gender, field of activity, and knowledge of privacy risks about the participants' readiness to activate the job alert function. Against our expectations and the results of previous studies, women are more ready to activate this function as compared to men. The most cited reasons given by the participants for deactivating job alerts is not the potential risks to their privacy, but the incurred reduction in battery lifetime and the potential disturbance caused by impromptu notifications.

From the application perspective, we have shown that job alerts succeed in reducing the distance needed to be covered by the participants, but do not impact the duration between the job publication and its reservation by jobbers. As in other domains, we have demonstrated that the value of the reward statistically influences this duration as well as the quality of the contributed data. In this case, the quality decreases when the reward increases.

The job alert function could be easily improved by allowing users to parametrize the notifications based on their preferences in terms of, e.g., frequency or battery consumption. By doing so, both resources and opportunities to get new jobs could be better balanced and support the different user profiles observed in our studies.

8. ACKNOWLEDGMENTS

Our thanks go to the participants of the user studies and M. Hollick as examiner for M. Michalak.

References

- [1] F. Alt, A. S. Shirazi, A. Schmidt, U. Kramer, and Z. Nawaz. Location-Based Crowdsourcing: Extending Crowdsourcing to the Real World. In *Proceedings of*

the 6th ACM Nordic Conference on Human-Computer Interaction (NordCHI), 2010.

- [2] O. Barak, G. Cohen, A. Gazit, and E. Toch. The Price is Right? Economic Value of Location Sharing. In *Proceedings of the ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication (UbiComp Adjunct)*, 2013.
- [3] A. J. B. Brush, J. Krumm, and J. Scott. Exploring End User Preferences for Location Obfuscation, Location-Based Services, and the Value of Location. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp)*, 2010.
- [4] D. Christin. Privacy in Mobile Participatory Sensing: Current Trends and Future Challenges. *Journal of Systems and Software (JSS)*, 116, 2016.
- [5] D. Christin, C. Büchner, and N. Leibecke. What's the Value of Your Privacy? Exploring Factors That Influence Privacy-sensitive Contributions to Participatory Sensing Applications. In *Proceedings of the IEEE Workshop on Privacy and Anonymity for the Digital Economy (PADE, LCN Workshop)*, 2013.
- [6] D. Christin and M. Hollick. We Must Move – We Will Move: On Mobile Phones as Sensing Platforms. In *Proceedings of the 10th GI/ITG KuVS Fachgespräch „Drahtlose Sensornetze“ (FGSN)*, 2011.
- [7] D. Christin, A. Reinhardt, S. S. Kanhere, and M. Hollick. A Survey on Privacy in Mobile Participatory Sensing Applications. *Journal of Systems and Software (JSS)*, 84(11), 2011.
- [8] D. Cvrcek, M. Kumpost, V. Matyas, and G. Danezis. A Study on the Value of Location Privacy. In *Proceedings of the 5th ACM Workshop on Privacy in Electronic Society (WPES)*, 2006.
- [9] G. Danezis, S. Lewis, and R. Anderson. How Much Is Location Privacy Worth? In *Proceedings of the 4th Workshop on the Economics of Information Security Series (WEIS)*, 2005.
- [10] J. Heer and M. Bostock. Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2010.
- [11] G. Hsieh, R. E. Kraut, and S. E. Hudson. Why Pay?: Exploring How Financial Incentives Are Used for Question & Answer. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI)*, 2010.
- [12] ITU. ICT Facts and Figures 2016. Online: <https://www.itu.int> (accessed in 09.16).
- [13] M. B. Kjærgaard. Location-Based Services on Mobile Phones: Minimizing Power Consumption. *IEEE Pervasive Computing*, 11(1), 2012.
- [14] M. B. Kjærgaard, J. Langdal, T. Godsk, and T. Toftkjær. EnTracked: Energy-Efficient Robust Position Tracking for Mobile Devices. In *Proceedings of the 7th ACM International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2009.
- [15] P. Klasnja, S. Consolvo, T. Choudhury, R. Beckwith, and J. Hightower. Exploring Privacy Concerns About Personal Sensing. In *Proceedings of the 7th IEEE International Conference on Pervasive Computing (Pervasive)*, 2009.
- [16] I. Koutsopoulos. Optimal Incentive-Driven Design of Participatory Sensing Systems. In *Proceedings of the 8th IEEE International Conference on Computer Communications (INFOCOM)*, 2013.
- [17] J.-S. Lee and B. Hoh. Sell Your Experiences: A Market Mechanism Based Incentive for Participatory Sensing. In *Proceedings of 8th IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2010.
- [18] T. Luo and C.-K. Tham. Fairness and Social Welfare in Incentivizing Participatory Sensing. In *Proceedings of the 9th IEEE International Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, 2012.
- [19] W. Mason and D. J. Watts. Financial Incentives and the „Performance of Crowds“. In *Proceedings of the ACM SIGKDD Workshop on Human Computation (HCOMP)*, 2009.
- [20] S. Patil, G. Norcie, A. Kapadia, and A. Lee. “Check Out Where I Am!”: Location-Sharing Motivations, Preferences, and Practices. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems (CHI EA)*, 2012.
- [21] S. Reddy, D. Estrin, M. Hansen, and M. Srivastava. Examining Micro-Payments for Participatory Sensing Data Collections. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp)*, 2010.
- [22] D. Reinhardt and C. Heinig. Survey-based Exploration of Attitudes to Participatory Sensing Tasks in Location-based Gaming Communities. *Pervasive and Mobile Computing (PMC)*, 27, 2016.
- [23] J. Tsai, P. Kelley, L. Cranor, and N. Sadeh. Location-Sharing Technologies: Privacy Risks and Controls. In *Proceedings of the Research Conference on Communication, Information and Internet Policy (TPRC)*, volume 6, 2009.
- [24] Urban Airship. The Good Push Index – How Targeting Boosts Push Messaging Response Rates. Online: <http://urbanairship.com/GPI-targeting> (accessed in 05.2016), 2014.
- [25] D. Yang, G. Xue, X. Fang, and J. Tang. Crowdsourcing to Smartphones: Incentive Mechanism Design for Mobile Phone Sensing. In *Proceedings of the 18th ACM Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2012.