Ubiquitous Collaborative Annotations of Mobile Maps: *how* and *why* people might want to share geographical notes

Mauro Cherubini, Fabrice Hong, Pierre Dillenbourg
École Polytechnique Fédérale de Lausanne
Station 1, CH-1015
Lausanne, Switzerland

{mauro.cherubini, fabrice.hong, pierre.dillenbourg}@epfl.ch

Fabien Girardin
Universitat Pompeu Fabra
Passeig de Circumval·ació, 8
08003 Barcelona, Spain
fabien.girardin@upf.edu

ABSTRACT
We investigated the production of collaborative annotations of maps with mobile devices, looking for the reasons why people might want to produce these notes and how they might do so. We conducted a field trial with 21 participants over 3 months in the center of Geneva, Switzerland. We collected fewer messages than expected and we attributed the reasons for this result with a lack of social conventions surrounding this activity and to the overwhelming effect of the map on participants which caused them to refrain from engaging in content-driven discussions and explorations.

Keywords
Computer Supported Collaborative Learning/Work (CSCL/W), Forum discussions, Informal Learning, Location Based Annotations/Services, Spatial Cognition, Short Messaging Service (SMS).

1. INTRODUCTION
Over the last decade, there have been a growing number of projects on connecting information to geographical positions, such as GeoNotes [7], a prototype developed at SICS in Sweden, ActiveCampus [3], developed at the University of California San Diego, E-graffiti, developed at Cornell University [1], and UrbanTapestries, developed by a non-profit company in the United Kingdom [4]. All these prototypes allow users to express opinions, preferences, recommendations, questions, and jokes, connected to specific places. More recently, commercial companies have launched similar services that received extensive media coverage [5, 6]. Services like `http://dismoiou.fr/` or `http://socialight.com` offer this kind of map plus content mash-ups. However, little research has been done on the user experience of these systems in the real-world and urban environment. Many questions regarding the production and the consumption of information in such systems are still open. In this paper we aim to understand further why people might want to share geographical notes with each other and how they might do so. We investigated the following questions: what are the reasons for sharing pointers to places? Is geographical position enriching the content of the messages? What modalities are most commonly used to publish or to retrieve these notes?

Answering these questions is difficult because research should be conducted in an absence of established practices. In fact, these location-based applications support new types of mobile experiences. A possible solution was proposed by Crabtree [2]: it consists in deploying new technologies in the wild and treating them as "breaching experiments", allowing these new technologies to provoke practices and reveal contingencies between activities and technological interactions.

In this spirit, we deployed a field trial of a mobile location-based annotation system that we called STAMPS. In this experiment that took place in Geneva during summer 2006, 21 participants used the system during three months, producing only 150 messages. The specific design of the application allowed us to record accurately the users' interactions with the system to a high level of detail. We analyzed the collected messages, the locations where they were produced, and the logs of the interactions. In addition, we administered a post-trial questionnaire to the participants.

Results revealed that the interest in this form of communication was high but participation was reduced by several technological and social factors that we will discuss in detail. The analysis of the logs also revealed the overwhelming effect of the messages-on-map on the reading patterns of the participants. This finding suggests that in such an interface the referential context provided by the map overrides the conversational context provided by the messages.

IWCES'07, November 4, 2007, Sanibel Island, Florida, USA.

2. BACKGROUND

The design rationale of GeoNotes [7] was to endorse an open information space, where users could connect notes with space. Each note was then categorized according to the room corresponding to the position where it was generated. A specific design choice made it impossible to read/write notes from a remote position, otherwise the connection between the note and its spatial context would be endangered. The evaluation conducted with 80 students from a university community showed that in general people used the system for chatting with three main aims: object chat, related to an object or physical aspect of the locale; situation chat, related less to physical object and more to ongoing activities and situations in which several users took part and talk-to-me chat, an urge to chat with others independently from time and place. The results showed that the triggers for authoring notes were not primarily physical objects or infrastructure, but rather the ongoing social activities and situations in that physical space.

A similar setup was used in the ActiveCampus project [3], where the researchers chose to create a viral community, because as they stated, for project sustainability they had to increase the application value, which increases with the number of users. One of their most interesting findings concerned the analysis of the actual locations of the sender and receiver of a message. The application logged continuously the position of the participating students. The analysis showed that for 473 out of 539 pairs the distance when messaging was shorter than their average distance. In short, relative location as a context seems to matter in community-oriented computing.

The importance of a 'critical mass' of messages/users was noticed in the evaluation of the E-Graffiti project [1]. The lessons learned included difficulties with a misleading conceptual model. In short, the designers expected an asynchronous use of the tool, but students used it mainly for synchronous communication, similar to that seen in GeoNotes. Authors also noticed a certain lack of use due to the reliance on explicit user input: "the fewer people using the system the fewer notes people will contribute and the less value other people will get out of the system by reading those notes". Finally, the authors highlighted the need for a highly relevant contextual focus: as this kind of technology is not part of daily life, users did not really think about information in terms of location, didn't know what notes to write, and didn't really have anything to share with others at a location.

The three projects described above were developed in the closed context of university campuses. In contrast, UrbanTapestries [4] was designed for use by the general population, aiming at sharing pointers about the city the participants were living in. The authors used the more specific term: "public authoring", meaning the process of mapping and sharing local knowledge and experience. During two field trials conducted between 2004 and 2005, they discerned a series of general feelings and trends about the process and relevance of public authoring to everyday life. One of the key issues was the interaction time: people expressed their need for quick and simple interactions while on the move, versus a more rich interaction when they are at work or at home. Authors also noted how the purpose of public authoring was seen as being about sociability. Participants saw the application as a new way of engaging in conversations about places that are fragmented and happen over time as well as in space.

While these studies carried out an extensive evaluation of the use of geographical messaging systems in a real context, they were either limited to university campuses or they lacked a detailed analysis of logs providing information on the mobile application and the context of its use. Therefore, this paper focuses on specific sessions of use of the system through an accurate analysis of the interaction logs. We were interested on testing the pervasiveness of geographical annotations about a city space, an area that is richer in terms of details that can potentially attract annotations.

3. METHOD

During the summer of 2006 we organized a field trial of geographical messaging in Geneva, Switzerland. Participants were asked to use STAMPS, an application for mobile phone developed in our laboratory (see section 3.5).

3.1 Participants

We recruited participants through leaflets posted in shopping malls and universities. We also posted the call for participation on our university blog. 21 people volunteered to participate to the experiment from three different contexts. Fifteen volunteers were students from two different academic research groups in Geneva. Four other participants were journalists from a local newspaper. The last two participants worked for a pharmaceutical company in Geneva. They were all native French speakers.

Each subject was offered the reimbursement of the connection costs. Also, we informed each group that the most active participant would have received a prize of 100 Swiss Francs (~ 69 Euro).

3.2 Apparatus

Each participant received a mobile phone (maker: *Nokia*, model: *3320*) with the installed STAMPS application. As we did not want them to use a distinct mobile for their personal communications, we encouraged them to put their SIM card inside the provided phone and use only that device for the period of the trial. That maximized the chances of having the application readily available when the need for it arose.

3.3 Procedure

As we wanted to leave the user free to use the system in any way they considered useful, we did not give structured instruction of what to do with STAMPS. Instead, we only offered a list of generic scenarios of use. The phones were delivered during the first days of June. Each participant received full support for installing the SIM card in the mobile phone, explanations on how to copy contacts and a tutorial on how to use STAMPS. During the period we offered a phone help desk to troubleshoot incurring problems. Participants had at their disposal a web site where they could find a Frequently Asked Question guide. The trial lasted three months.

3.4 Measures

The application generated fine-grained logs of the user's interaction with the system. Each action in the application was recorded and associated with a timestamp and extra details customized to every kind of action. For instance, a 'zoom' command in the application was associated also the coordinates of the resulting portion of map that was displayed after the command execution. A 'read' command would have been associated with the unique identifier of the message retrieved, etc.

3.5 STAMPS

STAMPS is an application for Symbian Series 60 smartphones. It combines two main functions: it allows the user to visualize the maps of the place where s/he is located and annotate these maps. Maps are streamed from GoogleMaps. Users can move around the map, and zoom in and out to the available levels of details (part (b) of figure 1). To annotate, the user can locate a specific point on a map and associate a message to it. Once posted, a square is shown on the chosen position (part (a) of figure 1). Other commands allow the retrieval of all the messages displayed in the area currently visualized (part (c) of figure 1), the search for specific content, and the filtering of content based on a temporal criterion (part (d) of figure 1). This last option was designed to allow the user to avoid the cluttering of the display with message landmarks (see for instance part (b) of figure 1).

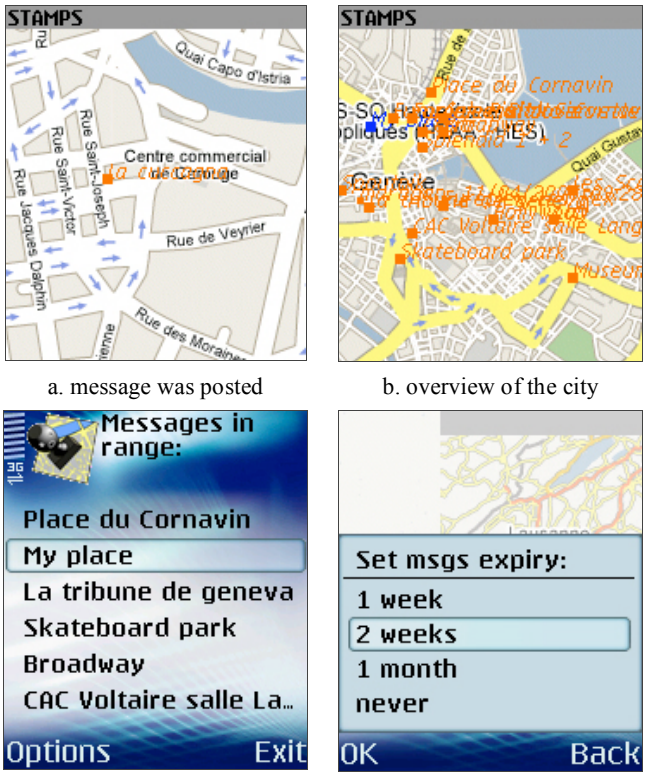

a. message was posted b. overview of the city

c. messages in range d. messages expire time

Figure 1. Display-captures of STAMPS

4. RESULTS

We categorized participants into three groups based on their usage patterns: (**P**) 'the passives' (5 participants), those who logged in the system only once and that did not produce any contribution; (**C**) 'the curious' (7 participants), those who participated briefly in the activities posting one or two messages and logging in average 5 times; and finally (**A**) 'the adopters' (9 participants), those who logged into the system frequently, often leaving their application running for a long time, produced most of the messages and engaged in many conversations. Table 1 shows some quantitative data of the analyzed dataset.

We focused our analysis on the last two groups, looking for differences in their login sessions, their consumption and production style in the system. We then categorized the produced messages. Finally, we report results from the post-trial survey deployed to understand the overall experience and the reasons of use.

Table 1. General statistics for three months of system usage

group	pseudonym user	# messages	# answers	# sessions	av.duration session (sec.)	# searches
A	*Cyril*	9	7	61	12386	7
C	*Edwin*	1	0	6	179	0
C	*Cperroud*	0	0	27	175	4
A	*Faril*	11	1	56	24236	1
A	*Yakari*	4	1	17	378	1
P	*Vinch*	1	0	1	665	0
C	*aldomanus*	3	0	5	222	0
P	*faril*	0	0	1	24	0
C	*Jack*	0	0	9	190	0
A	*sid*	65	2	232	282	2
A	*martigan*	5	0	26	251	3
A	*bawawa*	2	0	20	329	0
A	*Rodellar*	9	0	9	478	0
C	*Rebus*	4	1	5	1535	1
C	*schmoggi*	0	0	5	258	0
A	*icon*	5	4	15	326	4
P	*Julie*	0	0	1	17	0
C	*Amapelli*	0	0	26	2210	0
A	*Neuneu*	8	1	34	11772	1
P	*barrault*	0	0	2	91	0
P	*nigelsh*	0	0	1	8093	0

4.1 Login sessions

In average, the adopters logged in 52 times and their sessions lasted 93 minutes, while the curious logged in 11 times, and their sessions lasted 11 minutes.

Looking at the single login sessions we could observe how the 'curious' used the system mostly with a 'browsing' attitude, moving around rapidly through the tiles of the map, zooming in and out and rarely taking time to read messages and to perform searches in the database. On the contrary, 'adopters' allowed time between each action for the application to correctly load the tiles. They choose with care the regions to explore with a few clicks and moves and finally posted messages and read available contributions. Figure 2 offers a comparison between these different attitudes. Circles represent zooming in or out the map, while triangles are moves in the four cardinal directions.

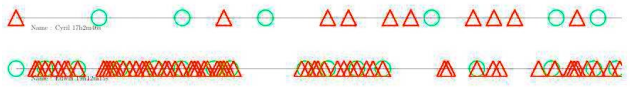

Figure 2. Timeline comparison of login sessions. An 'adopter' on the top and a 'curious' on the bottom (triangle: move up/down/left/right; circle: zoom in/out)

Additionally we looked at the hours of connection of each participant. We found that the whole population used STAMPS early in the morning, immediately during lunch break and after 5 pm. Figure 3 shows usage patterns for the 'adopters' and the 'curious' and a cumulative representation of the login hours for all the participants.

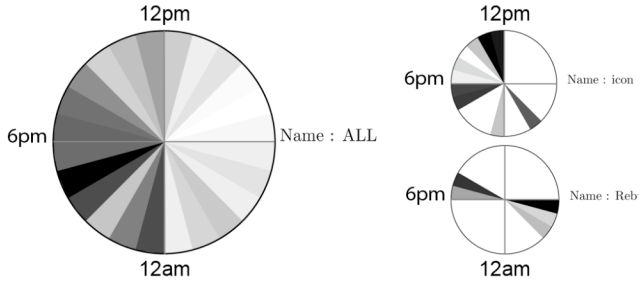

Figure 3. Left, cumulative representation of the login hours for all the participants. Right-top, login hours for an 'adopter'. Right-bottom, login hours for a 'curious'

4.2 Consumption style

On average, 'adopters' ran two queries in the database during the three months of the field trial, while the 'curious' did not run any query.

We looked at how participants retrieved the messages left by other participants and we found that the majority of users accessed content using the map navigation to isolate a region of interest first and then reading all the content available in that particular area. In a minority of situations, users were retrieving content through specific keywords: matching messages returned by the search engine might have been 'attached' to distant geographical locations. Nonetheless, they were retrieved and read systematically. Figure 4 shows the logs for these two behaviors, respectively. Rectangles represent a search by keyword, while squares represents read actions.

Figure 4. Timeline comparison of reading styles. Retrieval by content on the top, and retrieval by position on the bottom

4.3 Production style

Each time a message was posted, we logged the GSM network cell identifier to which the mobile was connected at the moment of submission. This is a unique number that distinguishes each antenna worldwide. In a densely urbanized environment like that of Geneva, the radius of each cell ranges between 100 and 500 meters (~ 109–547 yards). This was used as a rough indication of the position of the emitter of a message. We found that all the messages in the database were posted from 49 different antennas. A more detailed analysis revealed that participants used one to four different antennas to post their messages (see figure 5).

We then calculated the average distance of the messages posted using the same GSM antenna, under the assumption that if these messages were concerning events or items in the region covered by the antenna, then their distance should not exceed two times the radius of a GSM cell, namely 100 meters in a densely urbanized area. We found that most of the message posted exceeded 500 meters of average distance, revealing an attitude of the users to publish content 'attached' to locations far from their actual position (see figure 6).

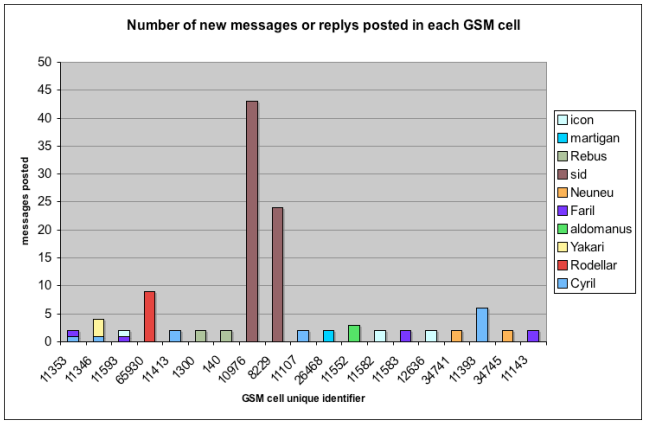

Figure 5. Frequency of message posted by GSM antenna

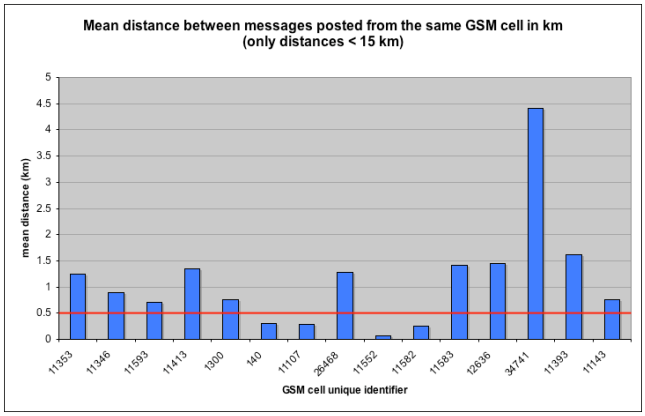

Figure 6. Average distance of the messages posted while connected to the same GSM antenna

4.4 Content of the messages

In this section we report examples of messages produced during the trial. The messages have been translated as they were originally written in French.

We categorized the messages posted during the trial using a coding scheme aimed at distinguishing the content of the contributions. As there are no widely accepted coding schemas for geo-localized messages, we arbitrarily defined our own following a discussion with other researchers building similar types of spatial annotation software [8]. We finally used five categories. The main category is the tips/assistance/warnings, notes offering useful information for the reader at a particular location (e.g., "*Beautiful view of the lake from the bridge*"). We subdivided this category in personal notes (TPP) and general use notes (TPG), depending on whether the message was intended to the group of friends (e.g., "*This is the place where I work with Rork*") or to all the users of the system (e.g., "*Best pizzeria of the city*"). Messages in this category did not have a particular temporal validity. On the contrary, messages in the events category (EV) lost their value after a certain temporal window (e.g., "*In the afternoon the tram went off the rails. People are still working on it*"). We classed in this group advertisements for items on sale (e.g., "*I am selling my bike, 150 CHF*"), invitations to parties (e.g., "*The faculty fiesta is there tonight. Hope I'll have fun*"), concerts (e.g., "*Concert Saturday, Sand over skara, make it pink and honey for petzi. 21h*

Piment rouge, 10 CHF"), other entertainment events (e.g., "*Improvisation matches, from the 2 to the 11 of November 2006 www.impro.ch*").

A different category describes 'spatialized' requests (RQ): people looking for a particular good or service in a specific spot of the city (e.g., "*A friend is looking for a roommate for 6 months. 700 francs/months. Call me!*"). Finally we used a separate category for tests and messages that could not be coded with the above categories (NA) (e.g., "*Nice to meet you*"). Table 2 resumes the frequency of each category on the messages posted in the system during the trial. The last column of the table reports the average number of characters of the messages in each of the 5 categories.

Table 2. Number of messages produced for each category

Category	# of msgs	Av. chars
Tips/Assistance/Warnings Personal (TPP)	16	49
Tips/Assistance/Warning General (TPG)	102	47
Events/Announcements/Ads (EV)	29	48
Requests (RQ)	6	57
None of the above (NA)	9	15
TOTAL	**162**	**47**

4.5 Questionnaire analysis

At the end of the field trial, all the participants received a short questionnaire consisting of five items. Six participants completed the questionnaire.

4.5.1 In which situations do you think using messages that refer to specific locations in space, like those of STAMPS, might be useful?

All the respondents answered that geo-located messages can be useful in situations where at least two people want to communicate content that is related to a physical location. The situations most frequently suggested were: showing local recommendations or the history of a place; revealing personal footprints to friends; information about happenings like concerts, highway traffic, places left in a certain parking lot; ubiquitous games.

One participant highlighted the fact that in order for the system to be useful its use should be related to an existing scenario, like a school trip or a work assignment. He specifically said that it might be interesting for learners to have information in context, which may lead to enriching discussions about buildings/monuments, etc.

4.5.2 How does STAMPS compare to other messaging systems like SMS or Newsgroups? Why/When would you use a SMS, an email or a news post instead of posting in STAMPS?

All the respondents noted how an SMS or messages on a forum are independent of location. They are one-to-one or one-to-many asynchronous conversations like messages on STAMPS but they do not relate specifically to space. Respondents described STAMPS as a communication tool in between SMS and newsgroup messages in terms of time-resources necessary to produce a message. While a news post is often long and detailed as directed to a wide community, an SMS can be quite informal, easy to compose and directed to a single person. STAMPS requires an extra effort to locate the specific points of anchorage on the map.

4.5.3 Do you remember any episode when you found something interesting/useful in the messages collected in STAMPS?

One participant, originally from Bern, said that she found interesting tourist advices of Geneva. She found indications on the locations of the beaches on the lakeside, public baths, and nice sightseeing spots. All the other respondents answered that they could not remember any particular episode where the application was useful. They all wished the application could have contained advices in particular situations like a traffic jam or a strike, but it was not the case.

4.5.4 Have you ever written an article for Wikipedia or similar community-driven sites? Could you tell the story? Even if you did not, how would it compare to posting message in STAMPS?

Most of the respondents did not have experiences in writing contributions on a wiki. Two participants mentioned differences in the user interface: STAMPS does not have a reviewing process like Wikipedia. Also, it is easier to produce written contributions with a computer keyboard than with a mobile keypad. Finally, a participant mentioned that the lack of automatic positioning introduced mistakes in the actual location of the places to which the messages referred.

4.5.5 What is the biggest limitation of STAMPS? Why did you stop using it? In your opinion, why it did not work?

Three participants mentioned that they did not have enough location-based interests in common with the community of users that were testing STAMPS. Another point that was mentioned was the lack of richness in the database: participants mostly already knew the information that was posted in the system.

One of the participants mentioned that messages should be produced on a desktop computer, while retrieval is fine with a mobile phone. He also mentioned that new messages should prompt the user to explore the position where they were attached and not the other way around.

5. DISCUSSION

The analysis reported in this paper shows that sharing location-based annotations with a mobile device is an emerging practice for which there are no established social conventions. We derived this conclusion from the answers given to our first survey question, in which participants listed a wide number of situations where the application might results useful. Many of these, like the social navigation of the city, are widely observable social practices (e.g., the number of cars parked in front of a restaurant, as well as the waiting line before a theatre indicate the places popularity). However, we found little evidence of these practices performed in STAMPS. The reasons of why this was not the case are difficult to backtrack. Nevertheless, questionnaire respondents explicitly stated that one of the reasons the application was uninteresting was because it was not available to their entire social network, or somewhat similarly, that participants did not have enough location-based interests in common. Many participants, especially the 'passives', did not see the utility of the service because it lacked content and perhaps usability.

We did not offer a structured scenario to follow and interestingly, participants produced, most of the time, messages aimed at the

whole population and not to a specific user. This finding is contrary to that of Persson and Fagergerg [7], who found that notes were targeted at other participants for social communication. STAMPS was perceived more as a person-to-community asynchronous information tool than a person-to-person communication application. This result goes against Burrell and Gay's report on the use of the E-graffiti platform [1]. They found mainly a synchronous use of the messaging system.

However, our system was missing a notification service for new messages. Thus, emitters could not have the certitude that a message sent to a specific person was read. We derived that, in order for STAMPS to function as a chat, recipients need to be notified promptly of incoming messages.

Another important point is that we observed an overwhelming effect of the map in the user interaction with the application. Participants were 'attracted' by the map and spent most of the time just browsing tiles instead of looking for content. We derived this argument from the small number of content-driven queries to the system and from the small number of discussions engaged during the trial. STAMPS followed a map-first interaction paradigm, in the sense that the map was used to route the user toward the content and not vice versa. While this 'map-first' approach supports users just browsing for interesting things nearby well, the mechanism is burdensome for those looking for specific and new content without a particular location in mind.

Finally, annotating maps is clearly a leisure activity that was mostly performed during breaks or during commuting time. Contrary to our expectations, participants produced most of their annotations while being in the same locations and not while being in the actual locations referred to in the content of the messages as originally hypothesized. This led us to think that the content posted in the system was somehow familiar to the authors of the posts and not discovered at their physical location. This finding was contrary to that of Griswold et al. [3], who found that local context was very important for the content of the messages.

We imputed the differences of our results to the other studies of location-based annotations on a campus setting [1, 3, 7] to the different levels of geographical scale. While a university campus can be identified with a specific social group, like that of the students who inhabit it, a city space is 'impersonal' as being used by a multitude of different groups. Therefore, we registered a majority of messages targeted to a generic community. Even if participants knew that their messages were mainly visible to their peers, this did not help to see STAMPS as a place for interpersonal communications.

As we felt that a more specific scenario given to the users could play an important role in the obtained results, we are currently running a second field trial with students from the urban planning course in our university. Students are asked to walk in a city and collect 'impressions' on the surroundings. These informal observations are then formalized in their course report.

To conclude, more research is needed to understand whether ubiquitous collaborative annotations of maps are useful and how best to support this activity. This study should serve as a cautionary tale to researchers who are trying to build such systems. The social characteristics of the annotation activity can be as significant as the design of the interaction mechanism with the user in its success or failure.

6. ACKNOWLEDGMENTS

This work was funded by the Kaleidoscope Network of Excellence, European Grant n° 507838 (OFES 03.0558).

We would like to thank all the participants for taking time to use the application, deal with its bugs and sharing their thoughts on location-based annotations.

Finally, we would like to thank several colleagues that have helped with the deployment of the trial: Giles Lane, Siddarth Jain, Florence Colomb, Shuja Parvez, Nicolas Nova, Cyril Rebetez, Jürgen Scheible, and Regine Buschauer.

7. REFERENCES

[1] Burrell, J., and Gay, G. K. E-graffiti: evaluating real-word use of a context-aware system. *Interacting with Computers*, 14 (2002), 301–312.

[2] Crabtree, A. Design in the absence of practice: breaching experiments. In *DIS '04: Proceedings of the 2004 conference on Designing interactive systems* (New York, NY, USA, 2004), ACM Press, pp. 59–68.

[3] Griswold, W. G., Shanahan, P., Brown, S. W., and Boyer, R. T. ActiveCampus: Experiments in community-oriented ubiquitous computing. *IEEE Computer 37*, 10 (2003).

[4] Lane, G., Thelwall, S., Angus, A., Peckett, V., and West, N. Urban tapestries: Public authoring, place and mobility. Project final report, Proboscis, UK, London, UK, 2005.

[5] Meyer, T. A chacun son guide de la ville sur son mobile. *Le Temps, Switzerland* (Saturday 26th of June 2004), p. 16.

[6] Morgan, R. Neighborhood report: New york up close; post-its for passers-by, 13th of November 2005.

[7] Persson, P., and Fagerberg, P. Geonotes: a real-use study of a public location-aware community system. Technical Report T2002:27, SICS, Sweden, 2002.

[8] Tester, J. *Why people will geo-annotate physical space*. Blogpost. Retrieved September 2007. http://future.iftf.org/2004/10/why_people_will.html