AudioFeeds - A Mobile Auditory Application for Monitoring Online Activities

Tilman Dingler
University of Munich
Amalienstr. 17
80333 Munich, Germany
dingler@cip.ifi.lmu.de

Stephen Brewster
Glasgow Interactive Systems Group
Department of Computing Science
University of Glasgow, G12 8QQ, UK
stephen@dcs.gla.ac.uk

ABSTRACT

User participation has transformed the way news travel the globe. With the rise of the 'Web 2.0' phenomenon [5] users have been empowered with the means of creating and distributing informational items, which we call *social feeds*. Platforms like Twitter[1] and Facebook[2] provide a variety of tools to facilitate real-time communication among people. But social sites are not limited to personal chat; they also provide an effective means for organizing large groups of people in response to catastrophic disasters. Monitoring these feeds can provide time-critical information, but can easily lead to information overload due to the large amount of data being shared.

In this paper we introduce a mobile auditory display application called *AudioFeeds* that allows users to maintain an overview of activities in different social feeds. AudioFeeds runs on a mobile device and enables users to get an overview of their social networks and spot peaks in activity by sonifying social feeds and creating a spatialised soundscape around the user's head. We conducted a user study looking into different aspects of activity monitoring. Results show that our application provides an effective way for monitoring overall activity levels and allows users to identify activity peaks with 86.1% accuracy even when mobile.

Categories and Subject Descriptors

H.5.2 [**User Interfaces**]: Auditory (non-speech) feedback

General Terms

Human Factors, Design

Keywords

Auditory Display, Social Media, Mobile Application

[1]http://www.twitter.com
[2]http://www.facebook.com

MM'10, October 25–29, 2010, Firenze, Italy.

1. INTRODUCTION

Social networks have changed the way people communicate. In an instant, short messages can be sent to a huge audience. Within that audience a piece of information, a link to a video or blogpost can rapidly spread from person to person, almost like an epidemic. People report not only their moods and daily activities, but also share rumours, news, pictures and movie clips. Even commercial news sites such as CNN[3] have discovered microblogging services like Twitter to be fast and effective for not only reaching large audiences, but also for receiving information streams that currently excite the masses. Social networks allow the effective dissemination of large amounts of information in real-time, often times faster than conventional news mechanisms. For example, when the US Airways flight 1549 crashed into the Hudson river in January of 2009, there were hundreds of messages about the crash on Twitter within minutes[4]. A passenger on a ferry that came to rescue survivors took a photo with his cellphone and posted it on Twitter right after the crash happened. The news spread through the net like wildfire as victims and on-lookers posted reports and images as the event unfolded.

Twitter has been widely used during various recent disasters and catastrophes where instant information distribution was required or where crucial information was coming from several sources. During the 2008 California wildfires critical information was aggregated on Twitter using the account 'LATimesFires'[5] which helped to organise and disseminate information. Other examples range from a person getting married and thus causing a lot of chat, up to the volcanic eruption in Iceland when thousands of passengers got stuck at airports all over the world. When information is rapidly shared about events like these, peaks of activity are caused on social platforms.

Twitter alone had 105,779,710 registered users and gained 300,000 new users per day in 2004[6]. There are several other microblogging sites out there like Tumblr[7], Plurk[8] and also established social network platforms such as Facebook provide users with status update functionalities.

Because of the huge amount of information being shared and distributed in real-time, it can be hard for users to keep

[3]http://twitter.com/cnn
[4]http://www.businessinsider.com/2009/1/us-airways-crash-rescue-picture-citizen-jouralism-twitter-at-work
[5]http://twitter.com/LATimesFires
[6]http://mashable.com/2010/04/14/twitter-registered-users
[7]http://www.tumblr.com/
[8]http://www.plurk.com/

Figure 1: Spherical platform placement: Facebook (left), Twitter (front), RSS (right)

Facebook (water)	Twitter (forest)	RSS (abstract)
Inbox Message (splash)	Friend Feed (chirping)	CNN (didgeridoo)
News Feed (bubbles)	Direct Message (crow)	BBC (zither)
Notification (pouring)	Reference (junglefowl)	TechCrunch (wind chime)
Friend Request (drops)	Hashtag (canary)	University News (pan flute)

Table 1: List of message event types that are retrieved and sonified by AudioFeeds

track of it or to make out interesting chunks [4]. But monitoring the general activity levels of those social feeds can be very helpful in cases like catastrophe management or to stay on top of a topic of interest.

To address this overload problem we built an ambient, auditory display that allows users to monitor the social feeds while being occupied or on the go. We use spatialised sounds in order to create a 3D soundscape around the user's head that reacts to the activity levels on platforms like Twitter, Facebook and to RSS updates. In our application, AudioFeeds, we wanted to allow users to maintain an overview of online activity and be able to spot activity peaks through an auditory display that runs on a mobile device.

We chose 12 message event types and designed sound sets for Facebook, Twitter and RSS feeds. These sounds are connected to real-time social feed data and create a spatialised soundscape around the user's head (Fig.1).

To evaluate our system design we conducted a user study to examine soundscape design, message event type allocation and the users' ability to maintain an overview of activities and spot sudden peaks. In this paper we outline our approach, the results of a user study and its implications on the application design.

2. RELATED WORK

The design of AudioFeeds has been inspired by various related projects. Using auditory displays to convey alarms and notifications has been common in mobile devices for some time, but often they are attention grabbing and the challenge remains to develop a rather subtler system, that allows users to remain focused on a primary task but monitor the general state of an application over time.

Sawhney and Schmandt [6] have looked into scalable auditory techniques to provide timely information and built a contextual notification system for wearable audio messaging called Nomadic Radio. AudioFeeds picks up the idea of adaptive notifications and applies it to social feeds and their activity levels.

Another approach for audio notifications has been undertaken by Butz and Jung [2]; they created and examined a method for notifying users with the help of auditory cues in an ambient soundscape. By embedding subtle cues in form of musical instruments and motifs into the auditory environment they delivered notifications to specific users. In our approach, we make use of personalized social feeds and create individual soundscapes for users in a mobile setting.

The design of audio integration in computer systems has been extensively researched. Brewster [1] has looked in great detail into how to enhance interfaces with sound and developed a detailed framework. Garzonis et al. [3] examined mobile service notifications and created design guidelines for auditory cues in terms of intuitiveness, learnability and memorability which we applied to the soundscape design in AudioFeeds.

3. THE AUDIOFEEDS SYSTEM

AudioFeeds is specifically designed for a mobile setting where users might not be able to monitor detailed online activities visually using the screen due to other ongoing tasks, for example walking or driving a car. This implies that notifications from the system should not require the user to fully focus attention in order to gain an overview of the system's activity. Rather than forcing the user to understand each incoming message, AudioFeeds is designed to convey a comprehensive overview of activities in the social feeds and to allow the user to notice specific peaks of interest.

To access individually relevant information streams, AudioFeeds connects to the online platforms Facebook and Twitter, as well as sonifies customizable RSS feeds. If there is an activity spike where a lot of people discuss about an item of interest on Twitter, Facebook or on a relevant news website such as BBC, our application will display this rise in activity to the user by increasing the density of sounds in the soundscape. AudioFeeds presents the events listed in Table 1.

It therefore retrieves updates to those feeds at regular intervals from each platform's servers and adjusts the current soundscape around the user's head according to the incoming event types. Each incoming message is assigned to a platform and a specific message type which determines the displayed sound. AudioFeeds runs on the Apple iPhone and uses OpenAL[9] to spatially position the sound around the user's head. Facebook items are presented on the left of the soundscape, twitter in the centre and RSS on the right. Each platform has an overall sound (Facebook: water scenery, Twitter: forest scenery and RSS: abstract instruments) and within that each event type has a particular sound (see Table 1), giving a total of 12 different sound cues.

The more activity in a platform or for a specific event, the more sound events of the same type will be added to the soundscape making the auditory representation more prominent. Sounds are added sequentially to the soundscape and stay in the system for a certain duration. Thus, users are

[9]http://connect.creativelabs.com/openal

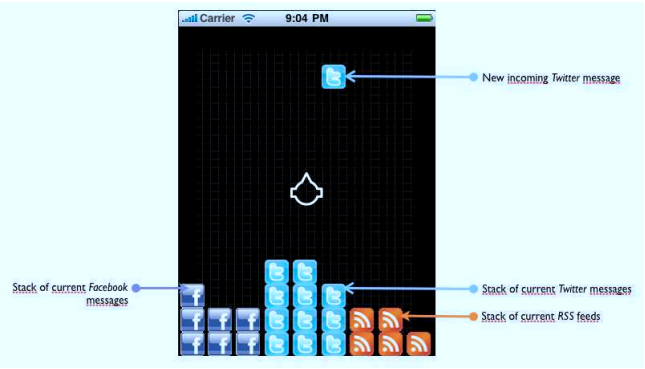

Figure 2: AudioFeeds GUI: Incoming messages are represented by touchable icons that are dropped from the top

presented a constant stream of sonified social feed events that allows them to monitor overall activity and track and identify interesting peaks. For interacting with the soundscape, AudioFeeds provides a graphical user interface where current events are visually presented. It enables users to request detailed information about an event by tapping on the screen (Fig. 2). However, the main focus of the application lies on its auditory display that allows non-visual usage on the go.

Before we came up with the final audio design and the way the audio is processed and affects the soundscape (which is presented here), we conducted a user study to find out about possible masking effects, message-platform allocations and information overload. Based on the results of this the sounds were refined until we had a successful set ready for more formal testing, as reported below.

4. USER STUDY

To evaluate the sound design as well as the application's effectiveness for feed monitoring, we recruited 15 university students as participants (10 men, 5 women between the ages of 20 and 28) who reported normal hearing abilities. There were three stages to our study:

Learnability of Sound Cues

To guarantee that participants understand the presented soundscapes, the first part of the study focused on the learnability of the sounds chosen and ensured participants could identify the sounds. Participants had to assign the 12 sounds to the three platforms and twelve message types correctly before proceeding to the next part of the study. If, after a short training session, they could not respond with $>= 80\%$ accuracy they went through the training again. On average participants needed 2.2 (SD 1.3) training sessions to learn the sounds. After the training, participants were able to assign a sound to a platform with 94.1% accuracy.

Activity Monitoring

AudioFeeds sonifies incoming messages from different social feed platforms. One design requirement was to create a noticeable change in the soundscape to make incoming events perceivable. We examined the conditions under which a sound event was most likely to be noticed and how soundscape density (the number of sounds playing) affected the user's ability to monitor activities before becoming overloaded. Our application's main focus is in conveying a general impression or overview of the current online activity levels. However, we also wanted to see if users could identify particular events and at what point soundscape density made this impossible.

Participants listened to 10 one minute soundscapes which ranged in complexity from very simple (3 sounds playing) to complex (up to 16 sounds playing simultaneously). They were asked to note whenever they detected a new sound event and to assign the corresponding event type, if they knew it.

Peak Recognition

We defined a peak in activity as a rapid sequence of incoming message of one specific event type outnumbering other concurrent sound events by a ratio of at least 2:1. In order to examine how well AudioFeeds facilitated peak recognition, participants listened to 5 different soundscapes that contained various activity peaks. They did this while walking to simulate the effects of using AudioFeeds in a more realistic mobile scenario. We asked participants not to focus on single incoming messages, but rather to maintain an overview over the activity levels of the soundscape presented and to report extraordinary events.

Results

During the Activity Monitoring stage participants listened to 10 soundscapes that differed in the overall quantity of sounds being played and the number of sounds overlapping each other.

Fig. 3 shows that participants' performance was affected by soundscape complexity. The more overlapping activity there was, the more events passed through the system unnoticed. But at the same time, recognition of types of events remained relatively stable as complexity increased. The data confirmed that it was easier to spot the unique occurrence of an event type than counting the total number of messages of that type. Fig. 3 shows this performance difference between the percentage of overall recognised message occurrences (blue bars) and the percentage of recognised unique event types (green bars) that appeared in the soundscape. Even in a highly dense soundscape where many online activities were displayed at the same time, users could still distinguish between different message event types.

In regards to platform recognition, participants were nearly always able to correctly tell which social feed platforms were active (99.8% correct identification across soundscapes). This finding confirms our assumption that spatial positioning of sounds along with different sound themes greatly helps distinguishing platforms and that users can extract overview information from the soundscape.

The results from the Activity Monitoring stage indicate that our system might not convey every single message retrieved from an online platform in detail, but it does give a good overview of the overall activity levels and their meanings, which was our aim.

During the Peak Recognition stage, participants reported peaks of activity they perceived. An event type peak was made up from one particular event type that occurred often. Sometimes participants could not assign the correct event type (e.g. Inbox Message on Facebook), but did notice the

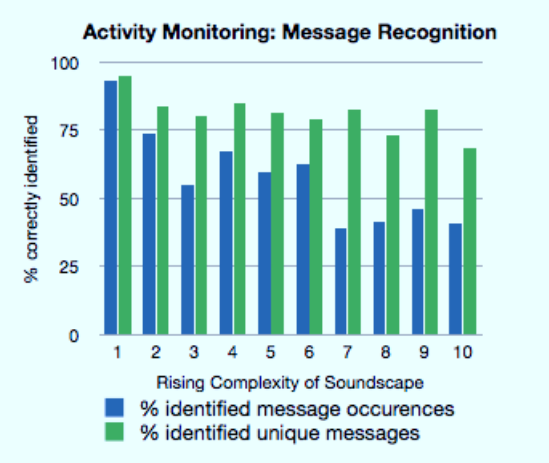

Figure 3: Overall (blue) and unique (green) message identification performance over different levels of sound densities in the soundscape

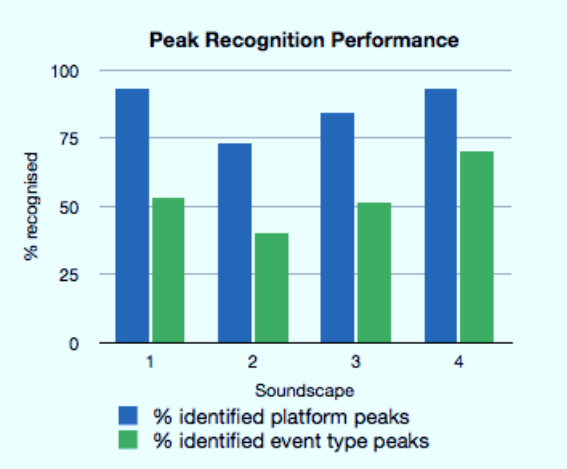

Figure 4: Peak recognition performance across different soundscapes

peak and could assign the corresponding platform (peak on Facebook) (Fig. 4).

The mean unique event type recognition dropped to 71% from the 81% during the Activity Monitoring stage. So the overall ability to identify unique message event types remains fairly good even while walking, which requires attention and concentration. Peaks of activity were noticed with 86.1% accuracy and correctly assigned to the related platforms (Fig. 4). 53.6% of the presented peaks were correctly assigned to the corresponding message event type. One soundscape was completely made up of message events from Twitter and RSS. The absence of any Facebook message was noticed with 86.7% accuracy.

5. DISCUSSION

Results form our user study show that AudioFeeds allows users to maintain an overview of activity levels in social feeds and spot peaks in activity within a feed. When extending the range of these feeds we may run into problems of scale though. There is a vast number of RSS feeds the user may want to stay informed about. Considering the sheer amount of unique feed messages, assigning unique sounds for each feed would be extremely challenging. We could instead categorise RSS feeds and assign sounds to each category (e.g. politics, sports, tech new, etc.).

Our user study has shown us once again how much sound design depends on personal perception and preference. We could support individual requests by providing various sets of sounds. But introducing new sounds increases the problem of potentially conflicting sound types. We are encouraged by the overall positive feedback from participants, especially from those who used Twitter and Facebook on regular basis to monitor friends' activities.

6. CONCLUSION

AudioFeeds is a novel way of monitoring social feeds by applying auditory display techniques and sound design. It allows users to maintain an overview of activity levels in social feeds like Twitter, Facebook and RSS while on the go. We conducted a user study examining performance in message type allocation and peak activity recognition. The results show that AudioFeeds enables users to easily make out interesting social feed activities while maintaining an overview even in complex streams of information. Thus, using AudioFeeds on a mobile device allows users to stay informed about current issues and to spot 'hot topics' while on the go.

7. ACKNOWLEDGEMENTS

This work was done as part of the EPSRC-funded GAIME project (EP/F023405). Equipment and additional funding was provided by Nokia.

8. REFERENCES

[1] S. Brewster. Providing a structured method for integrating non-speech audio into human-computer interfaces. *Department of Computer Science*, 1994.

[2] A. Butz and R. Jung. Seamless user notification in ambient soundscapes. In *Proceedings of the 10th international conference on Intelligent user interfaces*, pages 320–322. ACM, 2005.

[3] S. Garzonis, S. Jones, T. Jay, and E. O'Neill. Auditory icon and earcon mobile service notifications: intuitiveness, learnability, memorability and preference. In *Proceedings of the 27th international conference on Human factors in computing systems*, pages 1513–1522. ACM, 2009.

[4] A. Lavallee. Friends swap twitters, and frustration. new real-time messaging services overwhelm some users with mundane updates from friends. World Wide Web electronic publication, March 2007.

[5] T. O'reilly. What is web 2.0. *Design patterns and business models for the next generation of software*, 30:2005, 2005.

[6] N. Sawhney and C. Schmandt. Nomadic radio: speech and audio interaction for contextual messaging in nomadic environments. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 7(3):383, 2000.