TWEETSCAPES – REAL-TIME SONIFICATION OF TWITTER DATA STREAMS FOR RADIO BROADCASTING

Thomas Hermann[1], *Anselm V. Nehls*[2], *Florian Eitel*[3], *Tarik Barri, Marcus Gammel*[4]

[1] Ambient Intelligence Group, CITEC, Bielefeld University, Bielefeld, Germany
[2] HEAVYLISTENING , Berlin, Germany
[3] Freelance programmer, Berlin, Germany
[4] Deutschlandradio Kultur, Berlin, Germany
thermann@techfak.uni-bielefeld.de

ABSTRACT

This paper introduces *tweetscapes*[1], a system that transforms message streams from Twitter in real-time into a soundscape that allows the listener to perceive characteristics of Twitter messages such as their density, origin, impact, or how topics change over time. *Tweetscapes* allows the listener to be in touch with the social platform/medium *Twitter* and to understand its dynamics. We developed *tweetscapes* with and for the Sound Art department of the Germany-wide radio program *Deutschlandradio Kultur* where the sonifications are now broadcasted several times per week for a few minutes since October 2011. The goal was to create a new sense of media awareness and an example of how sound can support monitoring applications differently than mere alarms. This paper introduces the methods, the ideas, the design, the sounds, and it discusses our experiences with, and novel interaction possibilities offered by *tweetscapes*.

Figure 1: Screenshot of the *#tweetscapes* media stream at http://tweetscapes.de (2012-01-18): arrows show replies, #hashtags occur at the location of the tweet.

1. INTRODUCTION

One of the major advantages of sonification is that it enables the communication of information without requiring any visual attention and thus without any interference with a visual task. This makes sonification not only highly attractive for process monitoring tasks, (see [1]), where a process is to be followed while engaged into another primary task, or for information displays for the visually impaired who cannot access any visual information (e.g. see [2]), but also for radio broadcasts where there is simply no visual channel.

Since sonification can convey complex and detailed information, and we live in a decade of steadily growing information spaces, it is astonishing that it is nowadays so rarely used in established radio formats. To our knowledge the first regular use of sonification in a radio program was *broadcasting auditory weather forecasts*[2], a system introduced in [3] that represented many details of the expected weather (e. g. temperature, humidity, precipitation, wind, etc.) and its expected change over time for the next 24 hours in a 12 seconds soundscape, tuned to convey quickly and without the detour via language processing a good impression of how the weather is going to be like. From that project we learned that sonification in radio faces the particular challenge that sound needs to be as self-explanatory as possible and that the sonifications will be heard in many different contexts such as in the car, during work, in noisy environments – which imposes specific constraints on the sonification design.

As partnership and cooperation between Deutschlandradio Kultur and the Ambient Intelligence group at CITEC, we decided to create a new series called *Sonarisations*, where sonifications will be featured within the nationwide radio program *Deutschlandradio Kultur*. The given way of embedding the sonifications into the program – as gap filler between broadcasts and the news – provided some constraints for the selection of the domain as outlined in detail in Section 8. Furthermore we agreed that the tight cooperation of sonification

[1] offical name: '*#tweetscapes*'; we omit the '#' to increase readability
[2] German title: 'Die Wettervorhörsage'

scientists and artists/sound designers would be required.

Tweetscapes is the first and pilot project to establish and kick-off the series of Sonarisations. We conducted a workshop and presentation with support from Sam Auinger and Martin Supper at UdK Berlin (sound studies/acoustic communication) and decided subsequently to follow the second authors' proposal to create a real-time sonification of Twitter traffic. The proposal was then jointly elaborated in tight dialogue between the involved artist and sonification scientist, the process, interesting in itself and discussed in [4], will only be referred to occasionally in this paper. The resulting soundscape aimed to be both aesthetically interesting and useful as a sonification, i.e. key principles for sonification such as reproducibility, precise algorithmic transformation [5] are respected.

Twitter serves as a good example for communication networks where complex interactions between agents have shifted from the real world to the virtual/digital realm; as a whole the network shows an an overall behavior which is difficult to grasp, if at all, from merely looking at few tweets. How do individual messages lead to tweet avalanches which become trending topics? How does the Twitter community respond to events in the worlds, ranging from simple events like the onset of advertisement breaks in the big German TV shows to breaking news? How can sound provide a new level of experience of the digital medium? and how can we best make sonification more widely known and accepted as medium? *Tweetscapes* follows these questions and furthermore showcases an interdisciplinary experiment between media, auditory display research and sound arts. This paper aims at explaining the sonification side as the main focus, but the other aspects will be touched on as well.

We start with a short introduction into the social communication medium Twitter and summarize the key phenomena that we find relevant. This leads us in Section 3 to the goals and design ideas of *tweetscapes*. In Section 5 we introduce the sonification methods stream by stream. For the website, we worked on an audiovisual stream (Section 6) where the synchronization of visual and auditory events helps to better understand the data. Section 7 provides and comments on different *tweetscapes* for typical activity patterns. Finally, we address some practical issues and share our experience when integrating *tweetscapes* to the radio program of Deutschlandradio Kultur.

2. TWITTER — MICRO-BLOGGING DATA STREAMS

Twitter is a social networking service that allows users to send *tweets*: short text messages of up to 140 characters. It has grown since 2006 to a globally known service. Registered users can follow the tweets of other users and thus become 'followers'. Topics are set by using *hashtags* which are simply words prefixed with the # symbol. Instead of watching the posts of users they follow, users can also query the Twitter stream for specific keywords and thus use Twitter as a news filter. According to wikipedia, Twitter has 140 million users[3]. The amount of information per day is incredible and difficult to understand as a whole from the microscopic views that the standard interfaces offer.

3. TWEETSCAPES: GOALS AND DESIGN IDEA

Tweetscapes follows several goals on different levels: from the perspective of sonification research, the goal was to make the idea of sonification more publicly known by integrating it into the regular radio program. From the perspective of radio makers, it should be aesthetically interesting and surprising, and touch a subject that is of public and cultural interest. The real-time sonification of Twitter traffic was a topic that is compatible with these different goals.

The key design idea is to create a soundscape that involves several sound streams, similar to the sound- (or land-) scapes that surround us in real environments. They typically have a foreground, middle- and background. Likewise,

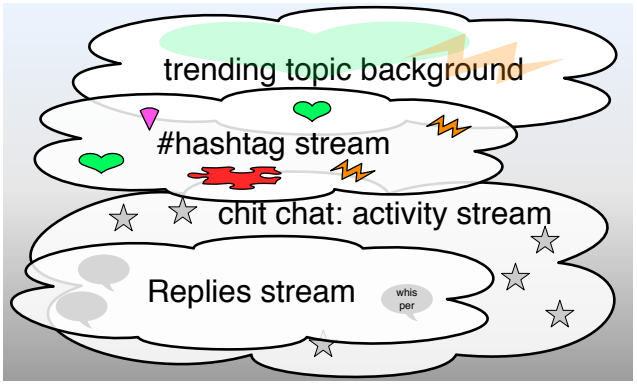

Figure 2: Sonic streams of *#tweetscapes*: salient hashtag events dominate a multi-stream background with activity, replies and topic streams.

tweetscapes represents the Twitter activity in several auditory streams: (a) *chit chat* is a stream where all tweets occur that are neither replies nor have hashtags, (b) *replies* is a stream of sonic events for public tweets exchanged by users, (c) *hashtag events* form the acoustic foreground stream where distinct topics become clearly audible, and finally (d) a *topic stream* makes the three most trending topics continuously perceivable as a background stream.

Apart from (d), all streams consist of individual sound events which are caused by tweets and thus are a true 1:1

[3] on March 21, 2012, see Section on growth on http://en.wikipedia.org/wiki/Twitter

representation of the event-like communication in Twitter. (a), (b), and (c) provide different filters or views. The sound events are chosen from a huge library of with different sonic material of large variation (as explained below) so that the overall sonic shape becomes acoustically rich. Finally the overall activity is estimated by a some features such as the average frequency of tweets. This parameter is used to influence the sound on many levels, such as sound effects, global parameters and post-processing.

4. PRE-PROCESSING OF TWITTER DATA

Twitter can be accessed through the Streaming-API[4]. This returns all tweets matching the particular query parameters in real-time. It is possible to filter by user names, keywords or location. The query is transmitted by a parameter in the HTTP[5] request. Twitter doesn't terminate the connection but sends new, matching tweets in real-time.

Unfortunately, there is no way to query tweets according a particular language. A series of tests showed that only 0.33 tweets per second are labeled with location information. This issue is solved by logging tweets with a very generic search query over a long time period. Based on this data the word frequency is analyzed regarding words from German users and non-German users. This results in a word list filled with words which are mostly used in German language.

Due to performance issues Twitter limits the rate of results on highly general search queries. To cope with these limits and collect nonetheless as complete as possible the German Twitter traffic, we created a restricted word list. The challenge is to filter these and suppress as good as possible the non-German tweets that may appear since words on the list are identical with words in other languages. This is also taken into account with the word list selection.

Every transmitted tweet is encoded as JSON[6] and contains approximately 54 parameters[7], which are related to the tweet or its sender. These characteristics are filtered, processed and enhanced as follows:

The location is important, especially for the visualization (see Section 6, below). If no location is set in the tweet the program takes a guess of coordinates based on location settings in the user preferences. If this is not successful a random position on the German map is created and cached based on user ID for a short time period so that repeated tweets from that user appear at the same location.

The hashtags need particular attention: A counter is incremented for every occurring hashtag h. The relative occurrence estimates the current popularity of the hashtag. The value is updated every 10 s by $m_h = \lambda m_h + (1 - \lambda) N_h$, where N_h is the number of occurrences over the past 10 seconds. We set λ of this 'leaky integrator' to get a half-life value of 5 minutes. This results in a dynamic ranking of all incoming hashtags. A ranking of the top 20 popular keywords is continuously extracted and sent to the visualization and sonification modules.

Additionally many more characteristics are processed, starting from simple metrics such as the number of followers of a user (followers count), retweet count of a tweet or character count of the tweeted message towards more complex parameters such as the time difference between a tweet and

feature	description/ type
realtimestamp	absolute time of tweet (by *tweetscapes*) float, ms since 2011-08-01
created_at	absolute time of tweet (by Twitter) float, ms since 2011-08-01
is_a_reply	flag if tweet is a reply to another integer (0/1)
RT_count	number of retweets integer, upper limit 100
text	chars of tweet text integer
user_followers_count	followers of User integer
user_statuses_count	count of tweets from User integer
RT_created_at	seconds since retweeted status integer, sec (default 0)
RT_statuses_count	count of tweets from retweeted user integer (default 0)
RT_followers_count	count of followers from retweeted user integer (default 0)
weekday	current weekday integer (mon=0)
sec_since_midnight	seconds since midnight integer, sec
mood	mood of tweet (guessed by emoticons) integer
question	number of question marks integer
longitude	longitude of Tweet float (default random)
latitude	latitude of Tweet float (default random)
tophashtag	best rated hashtag used in Tweet string (default " ")
relative_rating	best rated hashtag / current top hashtag float (default 0)
tweet_id	ID of this tweet string
RT_tweet_id	ID of retweeted tweet string (default " ")

Table 1: Extracted Features that characterize tweets in *#tweetscapes*.

[4]https://dev.twitter.com/docs/streaming-api/
[5]http://tools.ietf.org/html/rfc2616
[6]http://www.json.org/
[7]https://dev.twitter.com/docs/api/1/

later retweets, or 'mood detection' via the emoticons contained in tweets. Table 1 gives a complete overview of all extracted features. Finally, and very relevant for *tweetscapes*, statuses and users are filtered by a blacklist and identical tweets are blocked in a given time period to avoid spam.

These preprocessing results in 20 features which are encoded in Open Sound Control[8] (OSC) messages sent to the visualization and sonification modules. To allow multiple applications to access the data (debug, visuals, sound, logging, ...) the stream is not sent – as designed by OSC – over UDP but using a TCP connection. This enables the encapsulation into multiple servers and a clean interface between the different parts. The OSC processing applications usually require UDP packets, so a reliable proxy is used for parsing OSC packets out of the TCP stream and to translate them back to UDP.

5. SONIFICATION METHODS FOR THE *TWEETSCAPES* SOUND STREAMS

We will now discuss the sound streams and explain why and how the tweet features control the parameters of the sound events. Please navigate to http://tweetscapes.de[9] to familiarize yourself via the real-time stream with the sounds. As outlined in Section 3, the sonification contains four sound streams which we introduce next.

5.1. The chit-chat stream

As tweets are events, the most straightforward idea is to take a 1:1 manifestation of tweets as sound events. This resembles the Geiger counter that represents individual radioactive events as sound grains. Likewise a direct event sonification creates perceptual units on a higher level, such as the perception of momentary density and its change, of rhythms and waves. Beyond that, with the event sounds conveying details of the tweets, temporal patterns emerge that may become auditory gestalts. Our first attempt for such a granular texture of event actually used chirped sine tones to create a soundscape similar to literally twittering birds. Two sound examples are provided at our website.[10] Obviously, the bird sounds fill the sound space quite intensively. For that reason we considered other timbre spaces. We finally decided to use highly transient, non-pitched, short sounds. As sound source material we chose 8 sample sets of each 20 sounds from the area of communication, including single typewriter events, computer keystrokes, morse keys and relay clicks. Instead of modulating or manipulating features of single sounds, we decided to start from *ordered set of sounds*, (e.g. keystroke recordings at increasing force) and select the sample to be used according to the tweet's feature value. In this way, we automatically encode a data feature as a coherent auditory unit. For instance, the sample selection is driven by the number of followers of the tweet writer. Since tweets have obviously a higher impact depending on that feature, this 'impact' becomes literally perceivable as keystroke impact, which manifests in correlated level, brightness, complexity, duration etc. Technically this method can be regarded as a *parameterized auditory icon* [6] approach, yet the parameterization is here not achieved by a complex synthesis but via a table look-up. The term 'Sound Font' can be used for this battery of ordered samples.

In a nutshell the mapping[11] to sonic features is:

- impact (couples attack, level, timbre, etc., achieved via data-driven sample selection in ordered sample set [0,19]) ← user_followers_count.
- stereo panning [left, right] ← longitude [eastern, western edge of Germany], i.e. stereo position is as if the listener would be located in the center of Germany.
- reverberation [dry, wet] ← distance [0, 1000 km] from the center of Germany
- delay time decreases, and delay feedback increases with increasing RT_count, so that retweets can be recognized by their echo effect.
- sample set selection [complex, tiny] ← global activity [low, high], i. e. during lower activity the higher sparseness allows the program to select more complex sounds.

The algorithm is prepared to work with N-channel audio systems so that beyond a stereo panning also the latitude is properly mapped. Sound example S3 demonstrates chit-chat events for two single tweets, one near east, the second far away in the south. S4 contains two retweets, the first with RT_count $= 30$, the second with > 100. The spatial drift represents the spatial difference between the original tweet and the retweet location. Sound example S5 contains 5 selected chit-chat events with increasing impact (i. e. user_follower_count) Finally, sound example S6 is a typical chit-chat stream for German Twitter traffic.

5.2. The Replies sound stream

Replies are part of the public conversation at Twitter, but they are usually directed at a specific person. They should stand out of the chit-chat stream and have their own character and timbre so that listeners can perceive the ratio of non-replies tweets to replies from their occurrence frequency. A good metaphor is that of whispering. Similarly to the sound font

[8] http://opensoundcontrol.org/

[9] english version at http://tweetscapes.de/?lang=en

[10] http://techfak.uni-bielefeld.de/ags/ami/publications/HNEBG2012-TRT

[11] reported as sound parameter [min., max.] ← data feature [min., max.], using a linear mapping if not otherwise stated.

approach for chit-chat, here some longer samples of whispering are used where the whispering style gets more and more excited and faster with time. The length of a reply in characters is then mapped to the onset in this buffer to extract a snippet of appropriate whispering density that is further processed to deliver the reply sound event. Thereby longer replies sound more excited and faster without becoming unnatural.

Technically this method can again be regarded as a parameterized auditory icon mapping, but different from the approach in the chit-chat stream with discrete events in a sound font, we here realize a continuous selection process. While the actual psychophysical judgments of excitement may not increase strictly monotonously due to gaps and the details in the recorded whispering, the general trend will be dominating. The additional mappings are:

- sample file selection ← mood estimation, from :-) via :-| to :-(and nr. of '!' in the tweet.
- position in sample (degree of excitement) [begin, end] ← length of the tweet [0, 140 characters]
- The position and reverberation is consistent with the mappings for chit-chat events explained above.

Sound examples S7 contains a number of replies with increasing excitement (length of text). S8 contains a sequence of replies with average text length and different mood. They sound all neutral in space as they are the versions before any further post-processing.

5.3. The Hashtag sound stream

Hashtags are the parts of the tweets which we consider as relevant for judging the topics. Since hashtags can be freely invented by any user, it is impossible to set up a catalogue of possible strings and organize them in any meaningful way. As the sonification needs to create a sound in real-time without any intervention and reviewing by an editor, the sound needs to be synthesized from the string alone. Certainly, the first thought is to use any sort of text-to-speech system, or, to save time and avoid cluttering, to compress these spoken words just as *spearcons* do [7]. However, this would turn the sonification into a very verbal soundscape and possibly it would fail to convey what the Twitter dialogue is about. Thus we selected a more abstract way of encoding hashtags into sound-tags, oriented along two principles: (a) whenever a hashtag reoccurs, it has to be sonified by the identical sound as the previous one, (b) the hashtags cover a huge variety of sound events, just as words cover a huge variety of topics. Practically, we solve the problem by computing a hash which is reproducible for any hashtag string, with low risk that different strings result in the same hash value. We then use this hash to determine (i) a sound file in an extensible sample library with sounds from all areas of life, and (ii) details such as what snippet is extracted from the file and how it is distorted so that we obtain a very specific sound event for that hashtag. There is no easy way of generating a steady mapping between strings and sounds, so the hashtag #icad may sound very different from #icad2012. There is no underlying semantic analysis or categorization of words into classes such as economy, leisure, etc. Such extensions may be considered for specific continuations of the project.

Specifically the hashtag sound events are processed further using the following mappings:

- granular synthesis (sample, trigger rate, grain duration, etc.) ← hash(hashtag)
- sonority (how pitched vs. noise-like, via sample selection) ← ratio of consonants to hashtag length [0,1]
- delay, reverb, panning ← are consistent with chit-chat mappings.
- duration of hashtag events increase with decreasing global average activity (tweets per minute)

Perceptually, hashtag events stand out and appear as if in the foreground. Their unpredictability results in an element of surprise and should make listening to *tweetscapes* interesting even if there is no explicit interest to listen to it as a sonification. On longer and frequent listening to *tweetscapes*, users may remember and recognize certain sounds, such as #google, or #ff (short for #followfriday) on Fridays. Thematic changes are typically so slow that it is difficult to perceive them in continuous listening, but when listening to *tweetscapes* on different times or days, qualitative changes can be heard.

Sound example S8 and S9 are the hashtags for #papst (pope) and #piraten (a political party in Germany)[12]. Note that 'piraten' has more vocals and is somewhat more resonant. An example Tweetscape with these hashtags is discussed later on.

5.4. The Dominant Topics sound stream

As explained in Section 4, a ranking of hashtag frequencies is computed with a leaky integrator with 5 minute half-life. The technique to condense event streams into more complex events that represent aggregate properties was introduced in [8] and coined *Auditory Information buckets*. The idea is that a bucket collects information incrementally and flushes a more complex sound once the bucket is full. Here we take inspiration from this tipping bucket idea to define analogue structures that gather information about the dominance of topics. Only the three most filled collectors are selected for further sonification. Instead of a complex event localized

[12] as of Oct 2011, the algorithm has been refined meanwhile

in time we here create a continuous background sound that represents the hashtag sound as a stationary soundscape, so that the acoustic space is soaked with the idea of that topic. Certainly, this topic sound is the same as the corresponding hashtag, but using granular synthesis looped into a stationary pattern. Sound examples S10 and S11 present the corresponding topic sounds for the hashtags #papst and #piraten discussed in the previous section. To avoid a permanent overfilling of the sonic space with these topic sounds for the first 3 ranked topics, they are furthermore only added when they exceed a certain frequency (resp. counter value). The detailed mappings are:

- stereo panning [left, center, right] ← rank [2, 1, 3]
- level ← frequency counter [$f_{\min}$, $f_{\max}$], $-\infty$ below a threshold $f_{\min}$

5.5. Putting streams together

It is a difficult design task to tune all parameters and source sounds so that the individual streams work together as a coherent soundscape. Here particular effort was invested by the second author. The *tweetscapes* were first tuned according to our observation that the number of tweets rarely exceeded 5 per second using our filters. Sound example S12 is an example tweetscape with these data. However, modifications on the data interface to better capture the full German Twitter traffic led to an increase of the data volume per minute. In consequence a retuning was necessary since the soundscape became too densely filled. Sound examples S13 and S14 are two different versions for this more dense Twitter traffic. The solution to better cope with the available sonic space in time was to use the global activity (as already introduced above) to select the complexity and duration of events. This leads to less intrusive sounds once the intensity increases, as can be heard in sound example S14. From a sonification standpoint this procedure is debatable, since it breaks with the persistence of information. If we assume, however, that the main information lies in the level, frequency, echoes, reverberations and location, and we know that the density-driven selection process is reliable and reproducible we may simply adapt our listening habits and understand the soundscape correctly.

As a further extension we had considered including short verbal utterances that simply 'speak' a hashtag from time to time, at least on of the dominant topics. However, the speech synthesis lacked sufficient quality and robustness, given that hashtags are not necessarily words that can be spoken (e.g. #ff or #s21). So we canceled this path, yet it would probably be something valuable to consider for special application, such as for instance if visually impaired users showed an interest in using *tweetscapes*.

6. TWEETSCAPES VISUALIZATION

A frequent question that came from listeners who were first confronted with *tweetscapes* was 'what do the sounds actually mean?', 'what topics are discussed right now?'. We made clear that this is beyond the scope of the sonification and information we actively decided not to give. For the website at tweetscapes.de, fortunately the visual composer and 4th author Tarik Barri joined the team and created a real-time visual display (using his Versum [9]) that allows much better to connect the hashtag sounds with a particular meaning. The visual display shows the frontier line of Germany on a black background and dynamically creates colored light flashes at the location of the tweet. Furthermore, if it is a tweet with hashtag(s), the strings appear as text next to the light point. The synchronization of light and sound has two effects: (a) sound draws the attention to visual events, and (b) the textual display allows users to build up an association between hashtag sounds and their meaning. A particular feature is that replies to another tweet creates a visual arrow between the locations. This allows users to see how interconnected the Twitter space is.

7. TWEETSCAPES EXAMPLE SOUNDSCAPES

In this section we present three selected tweetscapes. The videos S15, S16, S17 are all captured from the live stream. S15 is a typical everyday activity. S16 represents a tweetscape at night – this is a much less populated soundscape. Finally S17 is a Tweetscape at a specific event. More detailed explanation will be given on the website with the sound examples. Our general experience is that the visual part is quite absorbing and draws the attention very much. So we recommend listening to the tweetscapes also with closed eyes, to investigate whether you can differentiate the situations by listening, or recognize or identify repeated topics.

8. EMBEDDING TWEETSCAPES INTO THE RADIO PROGRAMME

Tweetscapes was tailored to a particular role within Deutschlandradio Kultur's radio drama, documentary and sound art program: In this department, productions rarely match the precise length of their respective slots. The resulting time gaps are usually filled with generic music to be faded out when the news come in. In order to artistically shape this gap, Deutschlandradio Kultur's former sound art editor Götz Naleppa introduced a special format in 1998: 'Das Geräusch der Monats' (the noise of the month) were 5 minute sound art compositions designed to be faded in and out at any given time. This format was replaced by the Sonarisations in October 2011.

The piloting *tweetscapes* project meets the challenges of this particular slot in many ways:

- since *tweetscapes* taps into a live data stream, it can be faded in and out at ease
- *Tweetscapes* presents an artistic take on a topic of general interest
- the elaborate sound design makes *tweetscapes* equally accessible as a musical composition for a larger public and as a carrier of relevant information for experts

Embedding *tweetscapes* into the structures of Deutschlandradio Kultur required a number of thorough preparations. First, the concept needed to be communicated within the hierarchy and different departments concerned. The risks of real-time rendition, with unpredictable sound output needed to be tackled, both in terms of reliability (i. e. what if the synthesis fails?) and quality (i. e. what if Twitter traffic develops so that the tweetscape is unacceptable?). Furthermore, the embedding demanded significant technical infrastructure, from setting up a dedicated computer with the high security standards inside the intranet of the broadcasting station, to procedures to backup and access for maintenance.

Once these steps were taken, the integration into daily use required the production of programs explaining the purpose and idea of the project, as well as the setup of a project website, the edition of short texts for moderators to read before tweetscapes are played, etc.

Finally, the relaunch of tweetscapes.de with the audiovisual live stream challenged the means of a public broadcaster in terms of supporting online projects. However, the website and visualization have proven perfectly complementary to the sound stream, offering greater transparency and accessibility for a wide range of users.

9. DISCUSSION

With *tweetscapes*, we have – for the first time – established sonification into the regular program of a national broadcasting station. This project allowed us to learn many lessons on many levels. One level is the interdisciplinary communication: drawing together radio professionals, sonification researchers and artists/composers proved to be highly beneficial both for the involved persons that appreciated the different views and for the project since it offered to go beyond typical paths that probably would have been taken if not the mutual negotiations helped us to find a view 'in between' the poles. Our take is that it is definitively worth the effort.

The second level is the one of sonification for public media: we were surprised by the huge interest from media and press to report about *tweetscapes*, in fact the project launch event was highly visible due to press releases from DPA and even made it to several nation-wide newspapers. The reception of the project, however, showed a wide range of comments, from 'useful' / 'nice artwork' to 'waste of time'. Only few recognized *tweetscapes* as an example of sonification and understood the idea behind it, which is the general idea to represent complex information reliably by using non-speech sound. They related to *tweetscapes* more as 'making music from Twitter'. Mostly the question arose 'What is the practical use of listening to *tweetscapes*?' Indeed, the practical use is very limited – it is the *idea* that we here wished to transport. Understanding the Twitter space as such by listening is a new experience and that may or may not be inspiring for the listener. When getting in contact with public media, apparently there is the need and tendency to break complex ideas down into the most basic and raw concepts that anybody can connect with. This led to headlines such as 'turning Twitter into music', a phrase where sonification researchers will probably disagree.

9.1. Interactive participatory radio-making

On another level we see the potential of *tweetscapes* to establish something really new in radio broadcasting: the ability that radio listeners can via *tweetscapes* participate and influence the radio broadcast in real-time. This may on first sight only appear to be a neat gimmick, yet on second sight, it may allow completely new forms of radio shows. For instance, imagine that the moderator can ask the audience what they find most interesting to focus on – the radio listeners in turn tweet their opinion using pre-determined hashtags, and they can experience in real-time the distribution and frequency of opinions of others. The moderator can then use this information to refine or adapt the program or to select the next questions in an interview, etc. *Tweetscapes* thus provides not only a new 'unconventional view' on Twitter, it opens and suggests new forms of interactions in radio culture.

10. CONCLUSION

We have introduced *tweetscapes*, a real-time sonification system that allows users to become aware of Twitter traffic by listening. We have reported the goals, design ideas, methods, sonification streams, and played concrete examples for the various elements in *tweetscapes*. The multi-stream event-based sonification uses established parameter-mapping techniques and less frequently used ideas such as sound fonts and continuous sample selection for parameterized auditory icons. We explained how *tweetscapes* has been integrated into the regular program of Deutschlandradio Kultur and we have shown an audio-visual extension (live stream) which is featured on the project website. Finally we outlined some new ideas of how *tweetscapes* could in future inspire new forms of participatory interactive radio. *Tweetscapes* is the pilot project for the continued series 'Sonarisations' that

aims at making sonification publicly known by featuring its possibilities in a nation-wide radio program.

11. ACKNOWLEDGMENT

We thank Deutschlandradio Kultur who enabled the realization of *#tweetscapes*. We thank the German Research Foundation (DFG) and the Center of Excellence 277 Cognitive Interaction Technology (CITEC) that enabled this work within the German Excellence Initiative. We thank Sam Auinger, Holger Schulze, Martin Supper and Georg Spehr for early discussions that lead to *#tweetscapes*. We thank the staff at Twitter for their help, namely Katie Jacobs Stanton, Jason Costa und Carolina Janssen.

12. REFERENCES

[1] P. Vickers, "Sonification for process monitoring," in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. G. Neuhoff, Eds. Berlin, Germany: Logos Publishing House, 2011, ch. 18, pp. 455–491. [Online]. Available: http://sonification.de/handbook/chapters/chapter18/

[2] A. D. N. Edwards, "Auditory display in assistive technology," in *The Sonification Handbook*, T. Hermann, A. Hunt, and J. G. Neuhoff, Eds. Berlin, Germany: Logos Publishing House, 2011, ch. 17, pp. 431–453. [Online]. Available: http://sonification.de/handbook/chapters/chapter17/

[3] T. Hermann, J. M. Drees, and H. Ritter, "Broadcasting auditory weather reports – a pilot project," in *Proceedings of the International Conference on Auditory Display (ICAD 2003)*, E. Brazil and B. Shinn-Cunningham, Eds., International Community for Auditory Display (ICAD). Boston, MA, USA: Boston University Publications Production Department, 07 2003, pp. 208–211.

[4] H. Schulze, "Sonarisationen. ein projekt künstlerischer forschung des deutschlandaradio kultur berlin," in *Das geschulte Ohr*, ser. Sound Studies. Bielefeld, Germany: transcript Verlag, 2012, vol. 4, pp. 283–298.

[5] T. Hermann, "Taxonomy and definitions for sonification and auditory display," in *Proc. 14th Int. Conf. Auditory Display (ICAD 2008)*, B. Katz, Ed., ICAD. Paris, France: ICAD, 06 2008.

[6] W. W. Gaver, "Using and creating auditory icons," in *Auditory Display*, G. Kramer, Ed., ICAD. Reading, MA: Addison-Wesley, 1994, pp. 417–446.

[7] B. N. Walker, A. Nance, and J. Lindsay, "Spearcons: speech-based earcons improve navigation performance in auditory menus," in *Proc. Int. Conf. Auditory Display (ICAD 2006)*, T. S. et al., Ed., ICAD. London, UK: Department of Computer Science, QMC, University of London, 2006, pp. 63–68.

[8] T. Hermann, M. H. Hansen, and H. Ritter, "Sonification of markov-chain monte carlo simulations," in *Proceedings of 7th International Conference on Auditory Display*, J. Hiipakka, N. Zacharov, and T. Takala, Eds., ICAD. Helsinki University of Technology: Laboratory of Acoustics and Audio Signal Processing and the Telecommunications Software and Multimedia Laboratory, 07 2001, pp. 208–216.

[9] T. Barri, "Versum: audiovisual composing in 3D," in *Proc. 15th Int. Conf. Auditory Display (ICAD 2009)*, Copenhagen, Denmark, 06 2009.