EARCONS IN MOTION - DEFINING LANGUAGE FOR AN INTELLIGENT MOBILE DEVICE

Sami Ronkainen

Nokia Mobile Phones
Applications and User Centric Technologies Laboratory
P.O.Box 50, FIN-90571 Oulu, Finland
`sami.ronkainen@nokia.com`

ABSTRACT

This paper describes a case study of creating a musical language consisting of earcon-type sounds for an intelligent mobile device. The focus in sound design has been on machine-initiated interaction cases defining a way for the device to both catch the user's attention and at the same time give initial information about what caused the device to react. Psychological implications regarding alertiveness of the sounds have been considered in the design.

1. INTRODUCTION

In the future mobile devices will provide more and more intelligent services, many based on location and surroundings of the user as discussed e.g. by Brown *et al* [1]. Also the field of intelligent user interfaces (UI) and agent technologies is developing. As a consequence there will be an increasing number of usage situations where the mobile device will actively decide to react e.g. to some service being available at the user's current location or to some predefined event, such as a change in a stock price.

The current PC technology is mainly used in a desktop environment which enables the designer of a user interface safely assume that the user's sight can be utilized to notify him of interesting events. But in the field of mobile devices this is no longer true. It is more likely that the user's sight is indeed *not* available since the mobile device is worn on the user. It is only after catching the user's attention e.g. by using auditory or tactile stimulus that the display can be utilized. Usually auditory stimulus is used because it allows the device to be also away from the user when an interesting event occurs.

There are several known problems related to conveying information with sound. The sound must be distinctive enough so that the user knows which event it represents. But if there are many occasions when the device alerts the user using different sounds then the user may have difficulty to make out what caused the alert. Consequently, he then has to pick up the device every time to look at the display - which ruins the whole point of using more than one alert sound.

Another question that has been discussed in the academic world is if speech should be used instead of nonspeech sounds. In some cases speech certainly would provide some benefits but in practice it should be used sparingly. The problems of using speech have been described e.g. by Brewster [2] . Especially the fact that conveying complex information using speech requires using long sentences creates a high risk of

causing irritation in the user. Usually an easy way to make interaction sounds more pleasant is to keep them short.

In mobile devices speech also causes technical problems such as getting the quality good enough with very limited memory resources. Also supporting many different languages in a single device is problematic which is a logistical problem for devices that are manufactured in high volumes.

In general the aesthetic and user acceptance issues were left out of the research scope although they were recognized as extremely important in a consumer product of this kind.

2. USERS AND USAGE SITUATIONS

The research described in this paper was carried out in a project aiming at creating a demonstrator of an intelligent mobile device. The device created was capable of utilizing location dependent information and local ad hoc networking capabilities combined with methods e.g. for electrical payment. There were several assumptions also about the surroundings of the user, e.g. that there would be intelligent buildings and devices that can communicate with the device. The device had a small loudspeaker and capability of playing sampled sounds.

A typical usage situation would be that the user would walk e.g. by a vending machine that would introduce itself to the mobile device. The device would then alert the user that if he wishes he can buy a soda using the electrical payment capabilities of the device.

Related to the kind of events described above it was also necessary to notify the user if there is some default action which will be taken by the device if the user does not react. In the vending machine example, though, the default action would obviously be no action at all.

Another usage situation would be that when the user walks past the mail room at his office building and the smart room knows that there is mail for the user, the device could play out an alert. In this case there would be no reaction required from the user's part - it would just be a notification that afterwards would act as a reminder. In this case there is nothing urgent happening so the sound should be less alerting than in the previous example.

The device included also a combined set of sensors for enabling context awareness. This was used e.g. for automatically changing between UI modes in the device. In Nokia's current phones the modes have been named as "UI Profiles". They affect different settings in the phone, e.g. loudness of alert sounds and screening of incoming phone calls.

For example when detecting periodical movement and hearing typical street noises the device would turn itself into a

"street profile" including a loud ringing tone, no screening for incoming calls and other such settings. The fact that e.g. screening of incoming calls differs from profile to profile caused a requirement that the user should always be aware of which UI profile is currently active.

Another requirement related to UI profiles was that in some profiles such as on the street the sounds should reflect the high ambient noise level while in others such as in the office they would have to be softer but still audible and recognizable.

Yet another thing to be considered in the design was the huge already existing group of users and their habits. The project was a research project not aiming at a product but still it was seen pointless to spend lots of effort on designing something that in practice would be changed by the users anyway. For example in Europe a huge business has developed around allowing people to download their favourite tunes to be used as call alert tones. Therefore only little attention was paid to designing sounds e.g. for incoming phone calls.

This existing user base and de facto standards was also seen as a benefit. For example the default text message (Short Messaging Service, SMS) sound in all Nokia handsets consists of a pattern of two short notes followed by a pause then playing the same two notes again. This sound is the default one to be played whenever a user receives a text message. It has been used in all Nokia handsets today and has therefore become very widely recognized among user groups who use text messaging. In the research it was studied, though, if some minor aspect such as the intonation could be changed without affecting the recognizability of the sound.

3. DESIGN FRAMEWORK

The design was started by going through several use cases trying to come up with all different events that sounds should be used for. Soon it became evident that the different events could be classified into relatively few classes, each requiring different actions from the user. The identified classes are listed in table 1.

Event type	Explanation
Question, "look-at-me"	The *device* needs the user to look at the screen as it is displaying a query. For example asking if the user would like to buy a soda
Notification, "by-the-way"	Something not requiring the user's immediate attention. For example indication of new mail in the mail room
Reminder	The *user* himself has set up e.g. a calendar notification or a location based notification
Alert	Some *external party* is trying to catch the user's attention. For example a phone call or a text message.
Feedback	The device already has the user's attention and is playing sounds as feedback for the user's actions.

Table 1. *Event classes and their descriptions*

The next step was to consider the semantics of the musical language of the device. The language itself was decided to be kept as simple as possible. No sentence-like structures would be created. The messages would be merely of the kind "Look at me!" (or one of the other event classes) with the details of the

sounds telling more details of the event, if the user is capable of recognizing them.

It was obvious that inside one sound class all the sounds should share one or more qualities. After having decided what the qualities were it could be decided how the quality would refer to the class it represents. At this stage also psychological aspects were considered. For example it has been shown by Haas *et al* that high pitch and fast rhythm associate with high urgency of an alarm [3]. Therefore for the urgent events these qualities should be utilized in the sound describing the event.

It has been studied before that the rhythm and the timbre are the most distinctive qualities of sounds [4]. The two most important factors in the design was decided to be the currently active UI profile and the event class a sound represents. Differentiating between events inside one class was not seen quite as important since in many occasions e.g. in the 'Question' class the user would have to pick the device up anyway to see what the 'question' is. But also in that class some questions would be such that defaulting to no action will suffice. A good design would allow the user to notice also those situations.

It was decided that timbre would be used for indicating the active profile. This was because it was assumed that it would be easier to control things like perceived urgency of an event by changing the rhythm than by changing the timbre. A certain event in any profile should have the same urgency so rhythm was selected to describe the event classes. At this phase it was assumed that rhythm would be the dominant attribute affecting the alertiveness of the sound – this assumption proved wrong later in the user tests!

For profiles with a requirement for penetrating sounds - e.g. the street profile - timbres with rich harmonic characteristics were used. The sensors in the device could give direction about the sound level to be used but for technical limitations e.g. in loudspeaker design it was decided that also sound design should reflect the noise level.

All sounds (with the possible exception of very short sounds such as a key click) in an active UI profile shared the same timbre which would be different for each profile. It was assumed that people would be familiar with the concept of a UI profile and would notice that each profile has its own timbre.

Profile changes were decided to be indicated by *morphing* two sounds together. The concept of audio morphing has been described e.g. by Slaney *et al* [5]. In practice the morphing and editing features of an off-the-shelf sample editing software called Metasynth [6] were used. A sound requesting permission for a profile change would start playing with the timbre for the previously active profile smoothly transforming into the timbre for the new profile to be activated. If the user denies the profile change, the device would play a feedback sound with the sound morphing from the new timbre back into the old one.

As mentioned, in some UI profiles it was required that the sounds should be as undisturbing as possible. An example would be a profile to be used while in a meeting. The requirement was met by using soft percussive sounds such as a wood block as the timbre in those profiles. As a consequence the melodic information in the earcons was almost lost as percussive sounds hardly have any pitch at all. However, since the rhythm indicated the event class this was not believed to be a downside. In fact it was seen as a benefit since it has also been shown that melodies covering a large pitch range are perceived as more urgent than ones in a narrow pitch range [7]. In a meeting a user presumably has little time to concentrate on the mobile device anyway so too much alertiveness was to be

avoided. Another benefit in using a wood block sound was the opportunity to utilize the resemblance to an everyday sound by combining it with a rhythm resembling knocking on wood - utilizing ideas from both earcon and auditory icon paradigms.

This design also scales down well to devices where e.g. a piezoelectric buzzer is used for playing the sounds where little or no changes in timbre are possible for technical reasons.

3.1. Event Classes and Sounds

The relatively fast, three-note rhythm resembling a familiar knock on a door or a tap on one's shoulder was used for the 'Question' class of sounds. The fast pattern was believed to also have a bit alarming quality which suited the idea that the device is indeed trying to catch the user's attention.

The events in the 'Notification' class were regarded as not so urgent. Therefore a much slower rhythm with was selected for that class. It was also decided that four-note patterns were to be used for this class in order to create more difference to the 'Question' class.

The 'Reminder' class differs from the previous classes in that events in it are more important than in the previous ones. In the current implementations there have been used repeating patterns which won't stop until the user reacts. A relatively fast, waltz-type 3/4 rhythm pattern was chosen for this class.

Based on current experience it was known that users personalize their mobile phone sounds in the 'Alert' class quite commonly today. Therefore the sounds in this class were left mostly 'as is'. For testing purposes the SMS sound was changed so that it would use the timbre of the currently active profile, and the intonation of the sound would be a rising one.

The 'Feedback' class was left to lesser attention because the goal in this research was mostly to study the sounds where the device initiates the interaction and tries to catch the user's attention.

The melody was then decided to be used for conveying more detailed information. For example in the 'Question' class a monotonic melody or one going downwards was planned to be used when the default action is no action at all. This was to mimic the intonation of a neutral sentence in speech. A rising melody was used for sounds where the default action was that the device would actually change something such as the active profile. This was to mimic the intonation of a question sentence. Again the idea of having real world like characteristics in an earcon type sound was utilized.

4. USER TESTING

The objective in the user tests was to study if the sound design created would actually work in real usage context as planned. Special attention was paid to whether the users would be able to tell between different event classes or not.

Two separate user tests were arranged. One was a WWW-based test where the users had to group different sounds together. It was then analyzed what characteristics in the sounds were most important in classification. This test was done entirely unsupervised, every test person using their own computers to do the test.

The other test was a moderated user test where test persons were invited to participate in the test. In both tests all the test persons were Nokia personnel.

4.1. WWW-test

The WWW-based test consisted of a WWW page where the users had a table of 13 sounds, each sound followed by a row of seven radio buttons. The users were instructed to create sound groups by selecting the same button for each sound they feel should "go together".

The sounds in the test were built using two different timbres, three different rhythms (fast 3-note, extremely slow 2-by-2 note and a relatively slow 4-note waltz-type) and a rising or lowering intonation in the melody. All combinations of these parameters were used, creating 12 different sounds. The 13th sound used was one decidedly quite different from any other sound. This was to check that the test users had filled out the form correctly and to avoid pranksters.

The differences in computer setups people were using were tried to be compensated by selecting the timbres different enough so that even small laptop loudspeakers would suffice. The first timbre was a bell sound while the other one was completely synthetic. The test was piloted by two test users who commented that the timbres were different enough.

The sounds in the test were presented in random order in the WWW page. Because of the number of parameters it was chosen that at maximum six groups could be created. The seventh group was reserved for the 13th, "odd sound out". By selecting the number of groups this way the users were forced to select which were the most important characteristics for them when creating the groups. In the WWW form there was also a text field where the users were asked to tell in their own words, what they had based the grouping on. There was also a selection for telling if they had felt that there had been enough possible groups available, or if they would have liked to divide the sounds into even more detailed groupings. At the bottom of the page the users were also asked to evaluate their own musicality on a scale from 1 to 7.

After completing the page the form was handled by a PHP script that evaluated what the users had based their grouping on. The different grouping factors were "rhythm", "timbre", "intonation", or any combination of two of these. For example, if all the sounds in a group would have both a common timbre and e.g. a rising intonation, the script would classify as "timbre and intonation" to be the grouping factor for this group. Obviously there were also groups where the grouping factor was ambiguous.

It was found out that most of the users (67.5%) had used the same parameters for classifying each one of their groups. The number of grouping factors used by the test users can be seen in table 2.
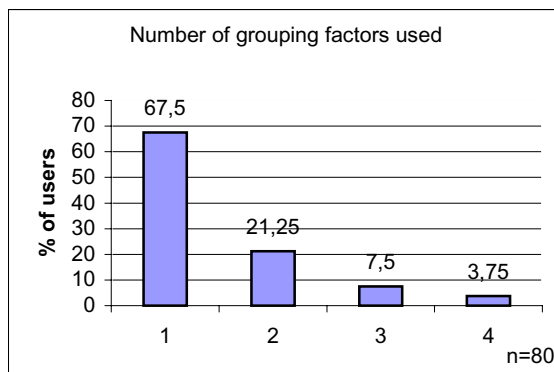

Table 2. *Number of grouping factors, all users*

The different sound parameters that people feel should belong together can be seen in table 3. In the table all groups have been considered, i.e. one user may have set up different groups according to different factors.
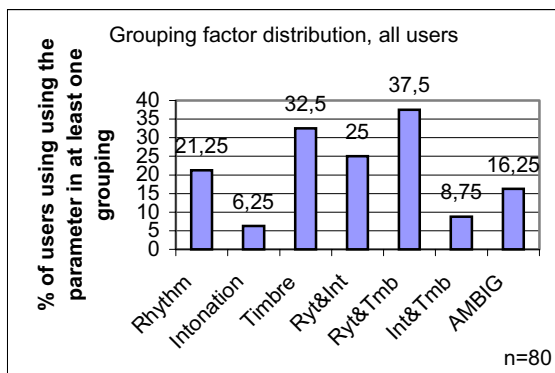

Table 3. *Grouping factor distribution, all users*

The interesting result in table 3 is that rhythm alone does not seem to be as strong a differentiating factor as was assumed in the original design. This result becomes even more obvious if only those users who had done their groupings in a consistent manner. This can be seen in table 4.

In table 4 it can be seen that rhythm alone is not at all an important grouping factor. This result was visible also when comparing the groupings with the opinions users gave in their own words. In many occasions people would say things like "beat" or "tempo" as the factor they based their grouping on, but from the grouping they did it could be seen that they still had often put sounds with the same rhythm but different timbre or intonation into different groups.

Only two out of 80 respondents in the test said that there had not been enough groups available for them. So, it seems that almost all people are somewhat tolerant to differences in sounds belonging together.

As a conclusion from the WWW-based test it can be said that most people were able to create consistent sound groups. The most important grouping were:

- Rhythm combined with timbre
- Timbre alone
- Rhythm combined with intonation

Especially in table 3 these three seem to have almost equal importance. So, for some people the intonation played no role. Two sounds with the same timbre and rhythm would belong together, two with differences in either would not.

For an another group of people timbre was the only dominant parameter, quite oppositely to the third group.

For the third group of people, the timbre was not important at all; two sounds with the same rhythm and same intonation would be classified together regardless of the timbre. Two sounds with the same timbre but differences either in rhythm or intonation would not belong together.

The third group of people was a surprising finding, since the intonation was assumed to be something only few users would actually pay attention to. It was suspected that users in the third group were more musical than those in the other groups. But at least based on the users' own evaluations this

was not true. The group where there were most people regarding themselves above average in their musical abilities was the one where rhythm combined with timbre was the most important differentiating factor. (65% against 57% in the group "rhythm combined with intonation" and 56% in "timbre only")

The least important differentiating factors according to the WWW-test were melody alone (2%) or combined with timbre (2%). Very few people would group two sounds with different timbres or different rhythms as belonging together based on their similar intonation only.
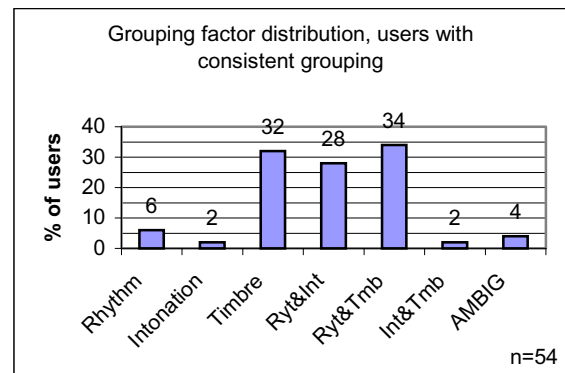

Table 4. *Grouping factors, users with consistent grouping*

4.2. Live Testing

Parallel to the WWW-based testing also a "live" test with invited test persons was arranged.

4.2.1. Test Setup

The test consisted by a short briefing session where the different types of events were presented together with sounds related to them. The sounds presented, however, had different timbres from the ones used in the actual test. This was done to avoid too big a learning effect before the actual test.

In the actual test the user walked in an office environment with the test moderator leading the way. Another test operator walked behind them with a laptop computer and a paper prototype of the UI working as a simulation. On the way there were 13 places where different events took place. When an event occurred, the laptop played a sound relating to that event. The users were then asked to guess what the event or at least the type of event that originated the sound could be. They were also asked about how important they thought the event would be and how they would react to the sound.

The different events and the sounds related to them are presented in table 5. All the 4-note sounds were decidedly much slower than the 3-note ones. The purpose was to make them sound calm and peaceful, as opposed to the relatively fast "knock-knock-knock" rhythm in all the 3-note sounds.

It should be pointed out that the two events related to printers were different. The first one was merely pointing out that the user's print job was finished and ready when he walked by the printer. The second one was a question asking if the user would like to print the SMS he had received.

Another thing to be noticed was that an important message such as a change in stock prices was indicated by exactly the same sound as an unimportant one; a question "Would you like to buy coffee" near any coffee machine. This was done to

provocate the users to comment on the perceived importance in the sound compared to the importance of the different events.

Event		Timbre	Melody
1	Profile change 1	1->2	Rising 3-note
2	Near printer 1	2	Monotonic 4-note
3	Receive stock information 1	2	Monotonic 3-note
4	Near coffee machine 1	2	Monotonic 3-note
5	Profile change 2	2->3	Rising 3-note
6	Receive SMS	3	Rising 2-by-2 note
7	Near fire alarm switch	3	Monotonic 4-note
8	Receive stock information 2	3	Monotonic 3-note
9	Near coffee machine 2	3	Monotonic 3-note
10	Profile change 3	3->4	Rising 3-note
11	Near coffee machine 3	4	Monotonic 3-note
12	Calendar alert	4	Long waltz-type
13	Near printer 2	4	Rising 3-note

Table 5. *Events and their sounds*

At the end of the test some of the sounds in the test were played to the users again, to test short-term learning issues. They also filled out a post-test questionnaire asking their opinions about the events and sounds, and possible suggestions they might have to improve the sounds.

5. FINDINGS

5.1. Sounds and Real-life Sources

Whenever there was something in the environment that the sound could relate to, it was usually the first guess by the users. For example the coffee machine alerts and profile changes were quickly learned by most users. Also the second printer alert was guessed to relate to the printer by most users, but no one pointed out that the alert used was entirely different from the first printer alert.

However, there was also evidence that learning had some effect. For example the users learned that a 3-note sound related to a coffee machine if one was nearby – but the last coffee machine alert was an exception. The last timbre used was so different from the others that people didn't think the sound could be related to the coffee machine.

5.2. Rhythm

Rhythm alone was not distinctive enough. As mentioned, the third coffee machine alert wasn't recognized nearly as easily as the previous ones even though it had exactly the same familiar knock-knock-knock rhythm. Some users even mentioned they recognized the rhythm but still didn't relate the sound to the coffee machine. Instead they would have expected a bright sound just like in the previous coffee machine alerts.

Another example is that no one recognized the SMS sound as having the same rhythm as in their everyday use. Some users did guess "SMS" as the probable origin - but when asked if they noticed that the rhythm was the same, no one admitted recognizing it at a conscious level.

5.3. Timbre

As mentioned above, the timbre had a crucial effect separating the last coffee machine alert from the others even though that was not intended. So radical changes in the timbre has an effect on what the sound is associated to.

On the other hand, the first timbre used (a string) and the second one (a bell) were not noticeably different by all users. Some users questioned e.g. if the second stock alert was the same sound as the first one. Also in the post-test questionnaire some people complained that there should be more differences in the timbres used. So the effect of timbre must be seriously considered in the future designs.

5.4. Rhythm and Timbre Together

The variation in perceived urgency was big in all sound types. The sounds that were clearly perceived as urgent were the 3-note sounds played with the bell timbre. The bell sound had a fast attack phase and rich high harmonics so it is in accordance with Haas *et al* [3]. But for example the SMS and fire alarm sounds' alertiveness was divided among the users. They were played with a bell timbre but the rhythms were slow.

5.5. Big Sounds and Small Sounds

A surprising finding was that length of sound is also an important differentiating factor. This finding was evident both in WWW-based and live tests. Several users commented things like "duration" as their grouping factor in the WWW-test. In the results these showed as "rhythm" (or rhythm combined with some other parameters). In the live tests, many users reacted to long sounds (such as the calendar alert) by saying that the device must have something important to say since it alerts with a "big" sound. Also loud and fast sounds were often referred to as "big" sounds. Especially the coffee machine alerts were commented as being too "big" for such an unimportant event.

As a conclusion it can be said that a slow tempo does relate to less important messages – but if it leads to a long total duration of the sound, the effect is likely to be lost.

5.6. Intuitiveness

One principle in the original design was that the three-note sound should relate to a familiar knocking type rhythm. The test also supports this theory to some extent. The three-note rhythm was clearly more learnable than the other ones and some users indeed did indicate that they recognized the rhythm. But on the other hand this rhythm was also the most frequent one in the test so its learnability must partly be based on that.

Several test users also said that they would have preferred more auditory icon type sounds, although obviously the term "auditory icon" was not mentioned. People liked the last calendar note since it reminded them of jungle drums.

Obviously in a short test of this kind people are not able to learn the sounds and the requirements for more intuitive sounds can be based on that. But in all, intuitiveness should be emphasized in the sound design when possible.

5.7. Language in General

More testing would be required to find out long-term learning effects, but based on this test it can be said that there was no clear indication that people would have understood the differences between the different event classes. For each sound there were some people that said they would pick the device up and some that said they wouldn't. The perceived urgency did usually work for the 3-note sounds. People would more often pick the phone up when hearing one of those, compared e.g. to the slow 4-note sounds.

Some evidence of grouping did occur, though. In the post-test questionnaire the slow 4-note sounds seldom got guessed as anything to do with location, but the fast 3-note knocking type sounds did. The exact meanings of the sounds were generally not remembered.

The timbre effects were a bit lost due to the fact that none of the users seemed to notice that the timbre changes related to profile changes. This interesting finding may also be related to the fact that the profile changes were a bit artificial – for example there was no clear explanation why the last profile used more quiet sounds than the previous ones.

5.8. Live Testing vs. Laboratory Testing

In general it can be said that the users' performance in the live test situation differed from their performance in the laboratory doing the WWW-based test. In the WWW-test many people were able to differentiate sound groups based on rather subtle differences, such as the 'rhythm combined with intonation'.

In the live usage (and test) situations things happen unexpectedly. People have no time to prepare themselves to concentrate on what they are about to hear. People also have other things in mind than testing the sounds. This was especially obvious in one case where a user – a very musical one – was first presented a 3-note sound which was repeated a minute or two later. However, right before the second time the sound was played he noticed a couple of his friends nearby and was clearly uncomfortable with his friends seeing him in the test. The result was that he felt the second sound was more alerting than the first one – even though they were exactly the same sound. There were also other occasions where some users did not recognize two sounds to be the same even though they heard them only minutes apart.

5.9. Problems and Future Issues

There were some problems that need to be addressed in the future tests of this kind. In the WWW-test it was found that e.g. sample playback only worked with some browser versions. This obviously lowered the answer rate.

In the live test situation it was sometimes difficult to say how important or alertive the users really considered the sounds to be. Often they just threw guesses of the origin of the sound. When asked how they would react, their answers reflected their guess what the sound was about. In the future the importance levels of different events could e.g. be graded by the users before going to the actual test.

Disturbances in the environment also affected the users' perception on the sounds. In future tests these effects should be normalized somehow. An interesting topic would be to study the "raw" perception of sounds, i.e. after each sound played the user would first be asked to simply describe what he just heard.

Also the long-term learning effects should be studied further in the future.

6. CONCLUSIONS AND FUTURE WORK

In this paper a case study for designing a musical language for an intelligent mobile device has been presented. The goal was to define a model based on which a learnable and usable set of sounds for machine-initiated interaction can be designed. Two separate user tests were made to test the design.

In the end it was concluded that the design did not work as intended. However, several new findings were found to improve the design in the future.

It was evident in the test that the perceived importance of a sound is a combination of all sound parameters; rhythm, timbre, intonation – and length of sound. It is difficult to parametrize sounds so that the perceived importance could be easily predicted when changing one parameter only.

Intuitiveness and familiar metaphors should be emphasized where possible. It should be considered which everyday sound characteristics could be utilized in earcon design.

If location based events are created, it should be taken into consideration that people tend to easily affiliate alerts with things in their surroundings that they consider as potential causes for the alert. Therefore it should be considered that sounds describing location based events would be clearly separated from those describing e.g. messaging or time based events.

7. REFERENCES

[1] P.J. Brown, J.D. Bovey and X. Chen, "Context-aware Applications: from the laboratory to the market place", *IEEE Personal Communications*, Vol. 4(5), pp.58-64, 1997.

[2] S. Brewster, "Providing a Structured Method for Integrating Non-Speech Audio into Human-Computer Interfaces", Ph.D Thesis, University of York, 1994.

[3] E.C. Haas, J.Edworthy, "Designing urgency into auditory warnings using pitch, speed and loudness", *Computing & Control Engineering Journal* Volume 7 4, pp. 193-198, 1996

[4] M. Blattner, D.Sumikawa and R.Greenberg, "Earcons and icons: Their structure and common design principles", *Human Computer Interaction, 4(1)*, pp.11-44, 1989.

[5] M.Slaney, M. Covell and B. Lassiter, "Automatic Audio Morphing", *IEEE Conference on Acoustics, Speech and Signal Processing* ICASSP-96 Conference Proceedings, Volume 2, pp. 1001-1004, 1996

[6] Anonymous. "Metasynth Presentation.", http://www.uisoftware.com/PAGES/ms_presentation.html, 1998

[7] J. Edworthy, S. Loxley and I. Dennis, "Improving auditory warning design: Relationships between warning sound parameters and perceived urgency", *Human Factors* Volume 33(2), pp. 205-231, 1991