

Ubicomp to the Masses: A Large-scale Study of Two Tangible Interfaces for Download

Enrico Costanza
University of Southampton
Southampton, UK
e.costanza@ieee.org

Matteo Giaccone
WeLaika
Torino, Italy
matteo.giaccone@acmos.net

Olivier Küng
EPFL
Lausanne, Switzerland
olivier.kueng@epfl.ch

Simon Shelley
University of York
York, UK
simonbshelley@yahoo.co.uk

Jeffrey Huang
EPFL
Lausanne, Switzerland
jeffrey.huang@epfl.ch

ABSTRACT

Tangible user interfaces have been promoted and discussed in the Ubicomp and HCI communities for 15 years. In TUIs physical objects are used for the control and representation of digital information, similarly to how icons are used in graphical user interfaces for the same purpose. Most reported TUI systems have the nature of research prototypes, available in laboratories or museums. This paper reports an attempt to understand the impact of TUIs in users' everyday environments through 2 low-cost simple set-up tangible interfaces for music that can be freely downloaded from a web-site. The systems are based on computer vision, printed paper and audio output. A few hundreds of users downloaded them and played with them. We logged users interaction with the interfaces and analysed content posted by them on our own and other web sites to observe and evaluate how they relate to such novel systems, taking measures to protect their privacy. Both the interaction logs and the users' comments indicate that the tangible interfaces were accepted as *normal*: they were perceived just as interfaces to make music rather than esoteric systems. Its time to bring Ubicomp technology to the masses!

Author Keywords

TUI, musical instrument, massive user observation, user generated content, d-touch, drum machine, sequencer

ACM Classification Keywords

H.5.2 User Interface: Evaluation/methodology

General Terms

Design, Experimentation, Human Factors.

INTRODUCTION

Tangible user interfaces (TUIs) have been promoted and discussed in the Ubicomp and HCI communities for 15 years, with journals special issues [10] and a specific conference, TEI¹, dedicated to them. In TUIs physical objects are used for the control and representation of digital information [8, 17, 11], similarly to how icons are used in graphical user interfaces for the same purpose. Proposed areas of application span from education [20, 7], to creative expression [12, 13] to urban design [18].

Most reported TUI systems have the nature of research prototypes, available in laboratories, or interactive installations on display in museums. This is partially due to expensive or complex set-ups, involving, for example, custom-built electronic sensing systems [13] or retro-projected surfaces [12]. As a consequence, most accounts of TUI usage to date have been restricted to controlled settings [20] or short-term interaction [7]. The suitability of TUIs for application in users' everyday environments, such as homes, offices and schools, as well as people's reaction beyond the initial stage, remain open questions.

At the same time, many in the Ubicomp community have stressed the importance of long-term deployments and study of new systems "in the wild" [15, 3] in order to understand the impact of design and implementation choices on real usage. A recent trend in HCI is raising the attention towards Web 2.0 and online communities as resources to gather research data. Examples include the use of existing user-generated content (UGC) related to commercial interactive products for their evaluation [2] and making GUI research prototypes freely available to online users to study them [19].

We argue for this approach to be extended further, making tangible and embodied interfaces available through the Internet to remotely study their adoption. As a first example, in this paper we report an attempt to understand the impact of TUIs in users everyday environments through an Internet-based observation of volunteer users. We developed *Audio d-touch*, a collection of low-cost tangible interfaces for musical composition and performance, and made

©ACM, 2010. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in Proceedings of the Ubicomp 2010 conference, <http://doi.acm.org/10.1145/nnnnnn.nnnnnn>

¹<http://www.tei-conf.org>

it freely and fully available for *download* from our website. Audio d-touch requires only a regular computer (Mac or PC) equipped with a webcam, an audio card and a printer. The physical interactive objects can be folded out of ordinary paper, with visual markers printed on it, or made by simply gluing printed marker labels onto existing objects. The webcam, hanging from a desk lamp or other improvised stand, tracks the interface items from above using the d-touch marker recognition system [5]: the objects position is mapped to control music synthesis parameters. We logged interaction with the interfaces and analysed content posted by users on our own and other web sites to observe and evaluate how people relate to such novel systems. As detailed below, beside predictable enthusiasm for the novelty factor of the interfaces, most users' comments focused on specific aspects related to the musical nature of the applications and their integration within existing ecologies of software music tools, indicating that the tangible interfaces were accepted as normal.

The next section provides an overview of related work, followed by a description of the Audio d-touch applications and their simple installation. The on-line promotion, the remote observation of the system and the analysis of user generated content are then reported, followed by a discussion and concluding remarks.

RELATED WORK

In the past fifteen years a number of researchers have worked on tangible user interfaces. Fitzmaurice, Buxton, and Ishii explored the possibilities of "graspable user interfaces" in 1995 [8], then Ishii and Ullmer introduced the term of tangible user interface in 1997 [11]. Based on this work a new user interaction model has been developed, in contraposition with the usual GUI model. TUIs represent physically the internal model of the software and, at the same time, they are the input and the output device [17]. Based on the idea that people have developed skills to manipulate physical objects, TUIs try to exploit this base knowledge to help the user and to build a better and seamless interaction.

The Audio d-touch system follows the paradigm of tangible user interfaces, in particular in the field devoted to musical applications. The musical TUIs are in an already established research and commercial field, as proven by Reactable and Audiopad [12, 13]. These instruments generate sounds through the interaction of the different blocks and their motions. Unlike these two well-known systems, Audio d-touch enters the category of *sequencer*, which is based purely on loops and repetition of patterns. A large number of other amateur and professional built musical TUIs² fit into this category with one of the most popular being BeatBearing [1]. Other musical systems include touch based instruments on table-tops such as the Microsoft surface or the popular Jeff Han's prototype [9]. The main difference with Audio d-touch lies in its low-cost and easy *do it yourself* properties. Instead of just watching prototypes in labs, museums or expensive recording studios, users can try it at their home. The other main difference with these systems is that the Audio

²<http://modin.yuri.at/tangibles>



Figure 1. Setup of the d-touch system with the webcam, the active surface, the paper blocks and the speakers.

d-touch applications are developed as a tool to remotely observe and track people usage of TUIs in their own settings, enabling the first large scale test of musical TUIs outside a laboratory.

Thanks to the availability of the popular ARToolkit, and other toolkits, as a library for Adobe Flash³, recently a number of augmented reality demos started to be available on the Web. Users can print ARToolkit markers, point a webcam to them and see them on their monitor augmented with 3D objects. However, many of these projects seem purely targeted at creating a *wow-effect* to promote commercial products and services⁴, rather than exploring new interfaces for actual applications. Moreover, no report was found of studies assessing the users interaction with these systems. In contrast our attempt is to use the Web to distribute complete applications that can be controlled using a tangible user interface, and study how real users interact with them. An earlier and less detailed description of this work [4] was presented at alt.chi 2010 (juried and non-archival track of the ACM CHI conference).

AUDIO D-TOUCH DRUM MACHINE AND SEQUENCER

Audio d-touch is a small collection of applications for real-time musical composition and performance with tangible user interfaces. It currently includes a drum machine and a sampling sequencer. Both are controlled by physically arranging a set of graspable interactive blocks (this kind of objects are sometimes named in the TUI literature as *bricks*, *tokens*, or *phicons*) on a flat surface, such as a table-top. We refer to the surface where the interaction takes place as the

³<http://www.libspark.org/wiki/saqoosha/FLARToolKit/en>

⁴See for example: <http://www.megabaile.com/> or <http://www.bmw.co.uk/z43d> or <http://www.mini.de/webcam>

interactive surface (or simply as the *board*). The blocks are tracked through an inexpensive webcam connected to a standard consumer-grade personal computer and their position, relative to the surface, controls the parameters of the audio synthesis applications.

Both the blocks and the interactive surface are marked with printed pieces of paper, containing graphic symbols used both to convey interaction cues to the users and to make easier to track them through the webcam. The d-touch marker recognition system [5] is used for the recognition and tracking. Four d-touch markers are placed at the corners of the interactive surface, for calibration, so that the interface behaviour is tolerant to small camera movements. One marker is also attached to each block, making it possible to define multiple types and to track in real-time the blocks x- and y-position as well as the rotation angle with respect to the surface. For ease of construction the interactive surface is normally defined by an A4 piece of paper, which defines a block size of approximately 2.5 by 3 cm. However, the system can be scaled to practically any size, as long as interactive surface and blocks are scaled by the same factor and everything is in the field of view of the camera. An example of the setup can be seen on Figure 1.

From the audio point of view, both applications are loop-based sequencers – similar in principle to many commercial electronic instruments. The Audio d-touch instruments allow users to create sequences of sounds arranged in periodic loops. Physical objects in the interface represent sounds that get reproduced when hit by an invisible virtual cursor that periodically scans the interactive surface from left to right.

Spatial Mapping

The interfaces of both Audio d-touch instruments were designed to be simple and easy to learn. Sounds and other *objects* (e.g. *loop start*, *loop end*) are represented by the graspable blocks, while *actions* (e.g. *play*, *record*, ...) are represented by flat areas of the interface, we call these *active areas*. The relative position of a block inside an active area can determine variations of how the action is applied to the object. This description is based on Schneiderman's action-object paradigm [16], which states that any user interface can be described as a set of actions and a set of objects. We think that in the context of tangible user interfaces associating objects to graspable items and actions to flat elements (or at least clearly separating the representation of objects from the representation of actions) can reduce ambiguity. In other words, the Audio d-touch interfaces are defined through spatial mapping: they are operated by placing the blocks in specific positions. The following paragraph describe how each application fits in this general description.

The d-touch **drum machine** was designed to be simple and playful and allow enjoyment with little effort. The object-blocks represent *drum hits* and the entire interactive surface is covered by a *play* active area: when blocks are placed on it the corresponding drum hits are played. The relative




Figure 2. d-touch drum machine. The x-position corresponds to the time instant within the loop, the y-position determines the type of drum.

position of the blocks within the play area determines the playback characteristics: the x-position corresponds to the time instant within the loop when the sound is played, while the y-position determines the type of drum – the active area is uniformly divided into 11 rows, each representing a different drum type, as illustrated in Figure 2. Two types of blocks are available, one corresponds to a normal drum hit, while the other corresponds to a louder hit. This can be useful to place accents within the drum patterns. In the drum machine the rotation of the blocks is not mapped.

The d-touch **sequencer** is a more flexible and advanced application. In this case the interactive surface contains 3 active areas: a *record* area, a *store pattern* area and a larger *play* active area divided into two identical tracks, as illustrated in Figure 3. Eighteen block types act as *sound container* objects, when one of them is placed in the play area its x-position determines the playback instant within the loop (as in the drum machine), while the y-position relative to the track determines the playback volume. The rotation of the block is mapped to the playback speed, coloured triangles on the blocks indicate the forward playback direction. The rotation angle is mapped to playback speed using a quadratic function, to place more *inertia* on the normal speed, and avoid small angles to accidentally cause sound distortion. When a block is placed in the record area the live input to the sound-card gets recorded “on it”, replacing any prior content. Similarly when a block is placed in the store area the content of the playback areas gets assigned to it, allowing the recursive construction of complex patterns. More than one block per type can be placed in the play area, all blocks of the same type will represent the same audio clip, but the sound will be reproduced according to each individual position. The size of the recording and storing areas are such that they can only contain one block at the time. The sequencer also includes 3 control blocks, representing the loop start, the loop end and the loop duration. The control blocks have different geometries from the sound blocks to make them easily distinguishable and convey their different role.



Figure 3. The d-touch sequencer. The x-position corresponds to the time instant within the loop, y-position relative to the track determines the playback volume. The two areas on the top right allow the recording of new sounds and to store the content of the playback area.

A Printable User Interface

Low cost and easy set-up were factors driving the design of the project from its start. The system requires no soldering or other electronics construction, only consumer-grade hardware that could be found at local computer stores: a standard personal computer running Windows or OS X equipped with an audio card and a webcam, or other camera (e.g. a consumer DV camera). The piece of paper covering the interactive surface can be printed using any inexpensive ink-jet printer, and the same is true for the labels on the interactive objects. The interactive objects can be built as small foldable boxes made of cardboard, or simply by attaching the labels to any small objects available, such as nuts, small chocolate bars, candies, toy bricks... To position the camera so that it correctly observes the interactive surface it can be attached to a desk lamp, or standard tripod, or even to a custom made stand built out of cardboard, or other inexpensive material.

To facilitate the distribution of the project we designed a proper kit in the form of cut-along-the-lines graphics distributed in the PDF format. These include parts to create a cardboard camera stand, parts to create the interactive objects and a board to define the interactive surface. In order to create the interactive objects, the user can either use printed parts to create cardboard boxes or labels that can be attached to existing small object. The purpose of this kit was mainly to help the project promotional communication – we wanted to convey that Audio d-touch can be simply and precisely reproduced following straightforward instructions and easy-to-find materials.

TUIS TO THE MASSES

Thanks to the portability and low-cost nature of the prototypes, since the very beginning, the Audio d-touch applications have informally been tested by a number of musicians and non-musicians across a wide variety of situations [6]: from laboratory open houses to letting our friends try them from our laptops. We generally received enthusiastic reactions, people were intrigued by the system, for its unusual aspects but also for its playfulness, simplicity and low cost

nature. Musicians particularly appreciated using an electronic instrument without a video monitor and compared it to audio effects pedals. The d-touch sequencer was used in few concerts by the professional composer and cellist G. Solima⁵ starting from 2006. A video of the 2003 prototype of the drum machine was posted on Youtube at the end of 2006 and without any promotion received about 22000 views in about 2 years. We also received a number of emails from people inquiring about the instruments and being interested in trying them.

This informal positive feedback, demonstrating that the instruments were potentially appealing to a wide audience, encouraged us to try and arrange a large scale observation. The low cost nature of the system was key in deciding to make it available for download from the d-touch.org website. The decision was also influenced by the general popularity of Reactable [12]. The goal of making a tangible user interface freely and fully available for download and being able to remotely observe its usage entailed a number of non-trivial technical challenges. First and foremost, it has been necessary to evolve the system beyond demo, creating a prototype able to run on a variety of machines and configurations. The process required an amount of testing and debugging more intensive than what is normally required for proof-of-concept applications which are expected to run only on few selected computers. Details such as simple installation and launching also had to be addressed. We deemed it important to make the system available on both Windows and OS X platforms, given the somehow different user demographics associated to each of them.

To observe how users interact with Audio d-touch, and in turn understand how they relate to its type of interfaces we used remote logging. We decided to collect the minimum amount of information that would let us reconstruct the interaction while limiting the invasion of the privacy of our users. The coordinates of all markers detected by the system, frame by frame, and (only for the sequencer) the audio clips recorded by the users were transmitted over Internet to a server and stored there in a secured database. To be able to identify repeated usage patterns we requested users to register on the website and employ the same username and password when the application is launched. During the registration process background information was collected through a questionnaire, asking users about their age, gender, occupation, whether they play any musical instrument and how frequently and whether they had any knowledge or practice about tangible user interfaces, and if so how. No personal information such as name and geographical provenance was requested. To limit and detect the possibility of having a single account shared by multiple users, we kept track of the users IP address. From the technical point of view, the data was posted from the application via HTTP to our server every 25 frames. All information was marked with a timestamp and stored on a MySQL database. A number of measures were taken to respect the privacy of our voluntary users. The remote logging operated by the system is

⁵A review and videos of the concerts are available on <http://d-touch.org/audio/concerts>




Figure 4. Histogram showing how long each user interacted with the system, in minutes. The box on the top right shows a more details about those who used it for less than 1 hour.

clearly stated in the conditions of use that users have to accept before downloading the software. Users can have their data removed from our storage at any point.

A Diary of the Launch

To distribute the system we created the d-touch.org website. Here we made available the applications for both Windows and OS X, the graphic files for the physical parts of the interfaces and documentation including photographs and videos showing the instruments in action and how to set them up. Initially only the drum machine was released, while the logging system for the sequencer was still being completed. The videos were also posted on YouTube, to facilitate their diffusion.

To promote Audio d-touch on June 28 2009 we started with an entry on the instructables.com website, a “Web 2.0” website dedicated to DIY, and then reached for a more general public. Our entry described how to build the drum machine setup providing a link to our website to download the software and the other required files, following the website suggestions we documented each step with images and provided a video demonstrating the finished working system. The project got featured on the homepage of instructables.com and within one day we received about 1000 visits. In the following days we advertised the project on the Twitter microblogging service and received the attention and the support of colleagues external to our group who helped us advertising the project. A number of blogs reported about the project – few of which we contacted directly – including the very popular engadget.com on July 2. As a result, by July 5 we noticed more than 12000 visits to the website, 671 users registering and 208 trying to launch the application, 112 successfully (more than 1 minute of usage). Somewhat to our surprise, most of the attention from the blogs was focused on the videos we posted on YouTube, rather than the actual system available for download.

In the following 6 weeks, until mid August, we released 4 updates on the drum machine fixing a number of bugs, while we continued working on the sequencer. The website received a total of about 25000 visits, while the project was

featured in 30 or more blogs, including 2 which reported a full hands-on review.

On August 17 we officially launched the sequencer, and a further update of the drum machine. We published a new video showing both interfaces, edited by taking on board some of the suggestions we received about the first video. We promoted the event via Twitter and by sending a formal press release to a number of blogs and printed magazines, including those who featured us earlier. Probably due to the reduced novelty, the communication received less attention from the general public: in 8 days we received about 7800 visits, approximately half of what we saw during the same period after the first launch. In total, from August 17 to December 15, the number of registered users was 1252. During the same time, 389 users tried to launch one of the applications, and 273 were successful (more than 1 minute of usage), resulting in a total of 199 hours of usage.

COLLECTED DATA

The interaction logs reported in this section are relative to the period between August 17 and December 15. Only users who started playing with the system on or after August 17 were selected because bugs in the earlier versions of the software may have resulted in the system behaving incorrectly and producing faulty logs. Feedback, UGC and comments were collected and analysed from the pre-launch (June 28) to September 9.

Registration

All users needed to register on the website and use the login to activate and use the software. Of all the self-reported demographic information that we collected during the registration phase we report the data relative to the 273 users who interacted with the system for more than 1 minute, except 7 who filled up the registration form with arbitrary text. In total 27% of our users were under 20 years of age, 73% under 30 and 94% under 40. Females were very few: only the 2% of our active users.

Familiarity with musical instruments and tangible user interfaces were assessed through 4 open questions. The answers were grouped into broad categories. 83 users had no musical knowledge; 104 reported to be amateur musicians (they can play an instrument but do not exercise daily) and 86 played daily. 132 users did not know TUIs before seeing our project; 132 had seen a TUI before (e.g. on the media) and only 9 had direct experience of TUIs.

Interaction Logs and Reconstructed Videos

The logs were processed to extract the amount of time that each user spent using the interface, shown as a histogram in Fig. 4, and the frequency of usage. Usage was generally infrequent: only 21% of the users interacted with the system over a period longer than 2 days and only the 11% over a period longer than a week. Data was also elaborated with respect to interaction sessions, defined as the time between the launch and termination of the application. The average session duration in minutes, the average number of interaction sessions per user and the average number of blocks

	Sequencer	Drum	Overall
Sess. len. (min)	8.75 (10.12)	6.75 (8.05)	7.34 (8.76)
Sess. per user	3.33 (3.13)	5.45 (6.35)	5.55 (6.34)
Blocks in sess.	1.80 (1.71)	3.06 (3.75)	2.69 (3.33)

Table 1. Average values for the audio d-touch usage sessions, gathered from the interaction logs. In brackets the standard deviation.

employed in each session are reported in Table 1. A *Liliefors test* confirmed that the 3 variables are normally distributed ($p < 0.001$ in all cases). A 2-way ANOVA analysis on the average number of blocks used in each session with respect to the instrument revealed a statistically significant difference ($p < 0.01$, $F = 31.38$): on average 3.08 (SD=3.77) blocks were used with the drum machine, while 2.16 (SD=2.44) with the sequencer. The same type of analysis did not show any significant effect of the instrument on the average session duration nor the average interaction time per user.

The interaction logs collected during the first three weeks after the launch were processed to reconstruct the layout of the blocks visible through the camera on each frame. An example reconstructed frame is shown in Figure 5. Frames belonging to each of the sessions were sequenced into video clips and assembled with the corresponding audio tracks. The audio tracks were generated using modified versions of the audio applications, which took the input from the logs rather than from the camera.

Comments, Feedback and User Generated Content

A considerable amount of informal data related to Audio d-touch was produced and collected. More than 120 emails, more than 330 posts on the d-touch forum, more than 50 blog posts, hundreds comments on the web and more than 220 Twitter posts. We also asked our users to share with us videos and photographs in order to see how they setup and use Audio d-touch. Six videos were posted on YouTube⁶, one photo on Flickr and two photos were sent to us via email; in addition two websites published hands-on reviews of the drum machine and included various photographs of the setup used. The content came from several places in the World, mostly from USA, Germany, UK, Chile, Brazil, Japan and other minor contributions from other places.

ANALYSIS

Logs and Reconstructed Videos

The statistically significant difference in the average number of blocks employed with each instrument reflects their design. In the sequencer users can record their own sounds "on" each block, so it is generally possible to create more involved compositions with fewer blocks. In contrast in the drum machine the sounds associated with each block are immutable and they are very short, possibly leading to the use of more blocks. More in general, this quantitative difference suggests that the users interaction with the two instruments went beyond simply trying a new purposeless gadget, reflecting instead their specific characteristics.

⁶Collected on <http://d-touch.org/audio/community/>




Figure 5. Snapshot of a video reconstructed from a user's log. This pattern produces a techno-like beat. The red square at the bottom right indicates that system is correctly calibrated. Green are normal blocks and blue are "louder" blocks used to put an accent on the beat.

As it was not possible to gather meaningful numerical data out of the reconstructed videos, we analysed the content in a qualitative way. As we have only analysed a subset of the videos, due to the overall length greater than 90 hours, we will provide just trends. A majority of the users we observed managed to explore the interface and to eventually produce a basic rhythm. By exploring we mean moving around blocks on the calibrated active area, and we define a basic rhythm as a pattern of more than four blocks. For this exploration to make sense the system needs to be calibrated, i.e. the four markers on the corners of the active area need to be detected most of the time, which means that the camera has to correctly point to the sheet. Relatively few users tried to go beyond the basic rhythm step and produce advanced rhythms with many blocks: only a minority of the sessions contained more than 10 blocks.

Another phenomenon observed in the reconstructed videos is that, in some sessions, markers are recognised intermittently, which negatively affects the user experience, as the layout of physical blocks in the interface may not correspond to the audio pattern generated by the applications. Given that the algorithm has been massively tested through more than 5 years in a variety of conditions, and that a simple tracking filter is used to limit the effect of temporary recognition failures, we suppose that this issue is related to severely poor lighting conditions. After this observation we realised that we did not provide to our users specific guidelines about lighting setup, and that the feedback that the applications provide about proper lighting and the stability of the marker recognition is difficult to interpret.

Generally users were able to create rhythmic patterns using the applications, except when the recognition was intermittent, or the camera unstable. Under these conditions, users seemed to engage in a *basic exploration* of the interface: they would place one or few blocks on different positions, as to explore what happens.




Figure 6 : A user setup publicly posted on Flickr, Uploaded on July 5, 2009 by “wolf confetti [I ♥ film]” and made available through Creative Commons “Attribution 2.0” License. <http://trunc.it/56vdg>

User Generated Content: Videos and Photos

We analysed the user generated videos and photos to find out how and where real users setup the Audio d-touch instruments.

First of all we observed that some people preferred to use the markers on flat pieces of paper, without folding it to make up 3D graspable blocks (Figure 6). It was interesting to see that sometimes entire uncut strips of markers were used as blocks to construct repetitive patterns. On the opposite end, we observed one user who built blocks out of wood to create a very polished setup. The cardboard stand that we offered for download was never seen in use, instead we found that webcams or DV cameras were mounted using tripods or hung from shelves. In some cases, as visible in Figure 6, the paper interactive board was raised, on books, to bring it closer to the camera.

In all the available material Audio d-touch was set up on desks, and never on other casual surfaces, such as coffee tables or on the floor. The desks used were often rather small and cluttered by different items, so the interactive board was rarely placed straight in front of the user (following the direction defined by the text), instead it seemed positioned just as space allowed. The interfaces were generally setup in small rooms, bedrooms or individual offices, in one instance in a cubicle. The hardware used for the system was varied: laptops and desktop computers, Windows and OS X. The audio output ranged from laptop built-in speakers to professional hi-fi systems.

In 2 instances we observed setups that we found very interesting in that they appropriate the technology in radical ways. In the first one, in a university setting, a user printed a board around 8 times bigger (as she herself explains) than the standard A4, and she affixed it vertically to a magnetic whiteboard. The blocks were made by gluing the printed d-touch labels to office magnets, so that they could be used on

the vertical surface. In the video the user explains that that such big and vertical setup more easily allows collaborative usage, as it is demonstrated by two fellow students filmed while they create a drum sequence.

Finally, in another we witnessed a setup of the drum machine setup without paper or other physical blocks. In this case the webcam is pointed to the computer screen, where a graphic design software (similar to Adobe Illustrator) is used to display the interactive board and icons representing the markers, all operated through the standard computer mouse. This is curious because it subverts completely the original aim of the system!

User Generated Content: Text

To analyse the text generated by our users we adopted an approach inspired by grounded theory [14, 2]. The analysis started by categorising the material at the sentence level through open codes. Initially 50 open codes were used, later grouped in 8 broader categories: physicality of the interface, audio, technical, improve and extend, field trial related, first impressions, personal data sharing, viral sharing. These categories were further organised in two groups: *actual usage feedback* that includes categories related to a real, first-person, experience of Audio d-touch, while *generic comments* refer to categories related to comments that do not appear grounded in real usage, but just based on first impression of the project.

Actual Usage Feedback

The feedback was almost exclusively explicitly directed at us, sent either via email or via the forum on d-touch.org. Exceptions included two hands-on reviews posted on other websites and some comments on Twitter, directed to their Twitter followers.

Physicality of the interface

From sentences related to the tangible interface we gathered that users appreciated the physical representation of sounds: “When I showed my good bands and songwriters the setup, they loved the realness and ability to touch and move something real to make the sound” or “Software doesn’t have to mean virtualizing everything and letting go of physical objects. On the contrary, it can create all sorts of imaginative, new ways of mapping musical ideas to the physical world. And that’s how we wind up with a walnut drum sequencer.”

Audio

A good number of entries were related to the audio nature of the system. As exemplified by the following quote posted on the forum, audio proved to be a playful and engaging domain of application: “Everyone in this house has now put together a radio-worthy beat by pushing little scraps of paper around under a webcam.” Other users suggested that Audio d-touch could be a good tool to teach music in school, or that they built the system for their children to enjoy the easy music generation with them.

Audio was also the main topic of criticism directed to the system. The limits and incompleteness of the music syn-

thesis parts of the applications are probably the main reason why users stopped using Audio d-touch early. Several users requested the drum machine to send MIDI signals or to be able to load custom samples. Especially for musicians these two restrictions limited the usage as it is illustrated from this extract from a post on our forum: *“it’s just a short-time toy because: 1) can’t send midi 2) can’t even load custom samples”* or again from the forum: *“I Could probably live without the Midi Sync if there was a BPM selector in the D-touch program in the video window.”*

Technical

The technical category gather data as bug report, problems with the software or setup and comments about the technical architecture of the software. The high number of interactions in this category can be seen as a great interest in trying and understanding the software by our users. Users who posted this kind of content generally showed to have technical competence.

Improve and extend

The people interested in Do-It-Yourself are generally interested in using their practical ability to enjoy their time. This is a natural source of interest for this kind of tangible interface that you can easily build. Even if intuitively spending time in building the interface maybe a point that discourage people, someone is attracted by this, as we can see from some comment: *“The system is made mostly of paper and cardboard; software and instructions of how to build the system are given on the website, in pure style Do-It-Yourself (probably the most appealing part of the project, from the user perspective).”* *“I spent a happy half an hour cutting out the shapes and putting the little boxes together.”* *“Really enjoyed the glue, lentil and paper time though, reminds me how good it can be to get down and dirty with materials.”* From this kind of comments no one complained about the difficulties of building the setup, instead some users proposed much more complex setup.

Field trial related

Some people complained about the remote data logging as well as the registration phase which was asking for too much data, but generally we received appreciation for this research. Often the fact that we were observing the usage was accepted because the application was free to download. We received questions about the possibility of paying a small amount of money to get the application, without the necessity of Internet connection, because this made it difficult to use the system on stage or in some recording studios.

Generic Comments

Here we analyze those comments that do not appear to relate directly to actual usage of the system. This content appears to be more generically related to emotional and initial reactions that people have when they see the project on the web. Even though this material is less informative than users feedback based on actual usage, we report it because it shows a large general interest for the Audio d-touch system and, more in general, for tangible user interfaces that are freely available for download.

First impressions

We found numerous comments, especially on blogs and on Twitter, which expressed surprise and enthusiasm for the project, a *wow-effect*, for example: *“Sci-Fi??? No more! You can actually build it!”*. We consider these as “first impressions”, representing initial emotional responses to the project, probably expressed before people actually attempted to use Audio d-touch.

Personal data sharing

By social sharing we mean passing personal information between people. This especially happened on our forum where people shared publicly and freely their location, or age, or other data. This is a useful side-effect of sharing the software with the on-line community, because we can assume that people are trying to build a community or simply to share knowledge about d-touch between themselves.

Viral sharing (Twitter)

As described above, Twitter was used as a platform to publicise the launch of Audio d-touch. Many Twitter users posted content related to the project, however, this often included only a link to our website or to other websites featuring our project. A very small part of the Twitter content was about personal feelings or experience about the project, but in any case it was always superficial.

DISCUSSION

At a general level the strategy of deploying our experimental systems through the Web and using the same channel for collecting information from voluntary users proved to be very rewarding. More than 97% of our users were cooperative in sharing with us useful information during the registration process on the website. Several users spontaneously shared with us their experience of the interfaces, documenting it through emails, photographs and videos. It is worth emphasising once more that no reward was offered for content sharing. We believe these results indicate that participants are very interested in experiencing first hand Ubicomp technology, which they normally can only indirectly observe, and that the social incentives typical of the Web 2.0 phenomenon transfer to the kind of experiment we conducted.

The analysis of the reconstructed interaction videos shows that the simple set-up works. Many users managed to get the system to work perfectly with the minimal amounts of instructions provided from our website, demonstrating that the low cost nature of the system was effective for spreading an experimental technology to a large number of users in their own environments. The general concept of interacting by laying out blocks over the board with the camera pointed to it is clear. However, from the reconstructed video analysis, we could see that in some sessions markers are recognised intermittently, even to the extent that calibration could not be performed. This was most likely due to poor lighting conditions and to inadequate camera position and orientation. In relation to this phenomenon, it must be emphasised again that we realised that very little information was provided to our users regarding optimal lighting conditions. Moreover

audio d-touch currently provides very little visual feedback about correct marker detection.

Most users explored the interface and produced basic rhythm, but very few used it regularly with more advanced patterns. This is probably due to the fact that the instrument is seen mostly as a toy, as pointed out in some of the comments: many users did not consider audio d-touch as a real music-making software, rather as a way to explore a new type of interface and to have fun making simple music beats.

The trends found in the interaction logs, as summarised in Table 1, supports that users generally perceived audio d-touch as toy. The more playful drum machine has been used more than the more complex sequencer (65% of usage) and the time spent on Audio d-touch is compressed in less than 2 days for the 79% of users. From the observations of the video sessions it emerged that almost all users explored the functionality of the new system to different degrees. In some cases the exploration was limited to the assessment of the basic functionality of the system: placing one or few blocks on the board (with the system being calibrated or not) users experienced the effect of having them in different positions, without attempting to create rhythmical structures. In other cases, the exploration was more directly targeted at exploring the ability of the system to create complex rhythmic patterns.

Generally the people most interested in d-touch are computer users of medium to advanced level, lovers of DIY and musicians. Everyone had good points for missing features or technical problems related to the application, but no one had specific negative comments about the tangible interface. We have in general very few comments about it, even if it's a completely new interface for common musical instruments. Only a small minority of our users reported, on our registration form, previous experience with tangible interfaces, nevertheless we observed from quantitative and qualitative data, that there is a rapid learning from users who have previous musical knowledge, showing that the interface is almost natural to users.

The user generated content showed several examples of *user appropriation*. Two teachers who wanted to use it in their classes, one to teach music to young children and the other to teach game design, they also changed the Audio d-touch setup to better fit their purpose, even though this involved more complex construction. Others, who wanted to use the d-touch applications for their music band, told us that they wanted to make the board in wood and use heavy blocks. Some users contacted us about building a large scale version of the system to be used in festivals or performances. One user managed to use d-touch even without a printer using the screen as the board, making a *virtual tangible interface*.

The high number of people registering on the website compared to the number of people successfully trying out the interface may be due to several reasons, and it is difficult to make precise guesses; we suspect that incompatibility with specific hardware platforms may be one of them. Audio d-

touch was developed using a number of open source third party libraries. While these libraries are extremely valuable and make the project feasible, they might interact in an unpredictable way on untested hardware configurations.

In summary, we argue that the Web distribution and user adoption of the audio d-touch tangible user interfaces was a success, even though the applications were perceived more as toys than as proper musical instruments. The quantitative data from the logs show that, beside technical problems, a large number of users were successful in interacting with the tangible interfaces and exploring their functionality. Very few comments were made about the user interface itself, and all of them were positive. As discussed above, we interpret this lack of comments as evidence that users found the interface “natural” or “obvious”, despite the fact that very few of them reported having experienced a tangible user interface ever before. Probably this circumstance is partially due to a considerable presence of tangible and multi-touch interfaces in popular media in recent years. Finally, the multiple cases of user appropriation indicate a strong interest and advanced understanding of this technology.

LIMITATIONS AND FUTURE WORK

Because inclusion or exclusion criteria were not defined for the study and participation was completely voluntary, results reported in this paper may be influenced by users representing only one specific population. While questions in the registration form assessed participants' familiarity with music and TUIs, no information was gathered about their engagement with DIY practices, nor how they had learned about Audio d-touch. In the future more questions may be included in the registration form, however, there is a fundamental tension between acquiring more detailed information and discouraging prospective users with too many questions. Additionally, because users invested their time in setting up the system before interacting with it, this additional step may have influenced the engagement dynamics, for example motivating them to play with the system for longer.

Based on the analysis of the reconstructed videos, our top priority for future development is the improvement of the visual feedback that the system provides about its correct setup. We believe that simple changes will allow users to better understand how to setup their the webcam and adjust the lighting for a better user experience. We are looking into mechanisms to foster the development of a community of users and sharing of music compositions.

In the future we are interested in extending audio d-touch and developing new low-cost tangible interfaces to include also visual output. Such extension could open up opportunities for wider domains of application.

CONCLUSION

This paper reported an observation of tangible interfaces in users' everyday environments. Two low-cost, easy to set-up TUIs for musical composition and performance were made freely available for users to download from a website. We observed users through interaction logs and analysed the com-

ments they posted on forums as well as multimedia user-generated content they posted on media sharing websites.

Despite some technical difficulties and missing features at the application level that could enhance the overall experience of the user, audio d-touch received very good responses. Few comments were expressed about the interfaces per-se and all of these were strongly positive; we noticed several examples of user appropriation. These observations show that the time is mature to distribute tangible user interfaces in an inexpensive and democratic way; future efforts should venture in domains other than audio.

At a more general level, using the web to distribute our experimental system and gather feedback about its usage was successful. Participants engaged in assembling the system and they were very cooperative in sharing information about their background knowledge and documenting their experience of audio d-touch. It's time to bring Ubicomp technology to the masses!

REFERENCES

1. P. Bennett and S. O'Modhrain. The beatbearing: a tangible rhythm sequencer. In *Proc. of NordiCHI 2008*, 2008.
2. M. Blythe and P. Cairns. Critical methods and user generated content: the iphone on youtube. In *Proc. of CHI '09*, pages 1467–1476, New York, NY, USA, 2009. ACM.
3. S. Carter and J. Mankoff. Prototypes in the wild: Lessons from three ubicomp systems. *IEEE Pervasive Computing*, 4:51–57, 2005.
4. E. Costanza, M. Giaccone, O. Küng, S. Shelley, and J. Huang. Tangible interfaces for download: Initial observations from users' everyday environments. In *alt.chi*, 2010. <http://trunc.it/6cdkk>.
5. E. Costanza and J. Huang. Designable visual markers. In *Proc. of CHI '09*, pages 1879–1888, New York, NY, USA, 2009. ACM.
6. E. Costanza, S. B. Shelley, and J. Robinson. Introducing audio d-touch: A tangible user interface for music composition and performance. In *Proc. of the 2003 International Conference on Digital Audio EffectsDAFx03*, 2003.
7. T. Drori and M. Rinott. Pixel materiali: a system for creating and understanding pixel animations. In *IDC '07: Proc. of the 6th int. conf. on Interaction design and children*, pages 157–160, New York, NY, USA, 2007. ACM.
8. G. W. Fitzmaurice, H. Ishii, and W. A. S. Buxton. Bricks: laying the foundations for graspable user interfaces. In *Proc. of CHI '95*, pages 442–449, New York, NY, USA, 1995. ACM Press/Addison-Wesley Publishing Co.
9. J. Y. Han. Low-cost multi-touch sensing through frustrated total internal reflection. In *UIST '05: Proc. of the 18th annual ACM symposium on User interface software and technology*, pages 115–118, New York, NY, USA, 2005. ACM.
10. L. E. Holmquist, A. Schmidt, and B. Ullmer. Tangible interfaces in perspective. *Personal and Ubiquitous Computing*, 8(5):291–293, 2004.
11. H. Ishii and B. Ullmer. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proc. of CHI '97*, pages 234–241, New York, NY, USA, 1997. ACM.
12. S. Jordà, G. Geiger, M. Alonso, and M. Kaltenbrunner. The reactable: exploring the synergy between live music performance and tabletop tangible interfaces. In *TEI '07: Proc. of the 1st int. conf. on Tangible and embedded interaction*, pages 139–146, New York, NY, USA, 2007. ACM.
13. J. Patten, B. Recht, and H. Ishii. Audiopad: a tag-based interface for musical performance. In *NIME '02: Proc. of the 2002 conf. on New interfaces for musical expression*, pages 1–6, Singapore, Singapore, 2002. National University of Singapore.
14. J. Preece, Y. Rogers, and H. Sharp. *Interaction Design*. John Wiley & Sons, Inc., New York, NY, USA, 2002.
15. Y. Rogers, K. Connelly, L. Tedesco, W. R. Hazlewood, A. Kurtz, R. E. Hall, J. Hursey, and T. Toscos. Why it's worth the hassle: The value of in-situ studies when designing ubicomp. In *Ubicomp*, pages 336–353, 2007.
16. B. Shneiderman and C. Plaisant. *Designing the User Interface: Strategies for Effective Human-Computer Interaction (4th Edition)*. Addison Wesley, 4 edition, April 2004.
17. B. Ullmer and H. Ishii. Emerging frameworks for tangible user interfaces. *IBM Syst. J.*, 39(3-4):915–931, 2000.
18. J. Underkoffler and H. Ishii. Urp: a luminous-tangible workbench for urban planning and design. In *Proc. of CHI '99*, pages 386–393, New York, NY, USA, 1999. ACM.
19. M. G. Van Kleek, M. Bernstein, K. Panovich, G. G. Vargas, D. R. Karger, and M. Schraefel. Note to self: examining personal information keeping in a lightweight note-taking tool. In *Proc. of CHI '09*, pages 1477–1480, New York, NY, USA, 2009. ACM.
20. D. Xu, J. C. Read, E. Mazzone, and M. Brown. Designing and testing a tangible interface prototype. In *IDC '07: Proc. of the 6th int. conf. on Interaction design and children*, pages 25–28, New York, NY, USA, 2007. ACM.