Tangible Bots:
Interaction with Active Tangibles in Tabletop Interfaces

Esben Warming Pedersen & Kasper Hornbæk
Department of Computer Science, University of Copenhagen
DK-2300 Copenhagen S, Denmark
{esbenwp, kash}@diku.dk

ABSTRACT

We present interaction techniques for tangible tabletop interfaces that use active, motorized tangibles, what we call Tangible Bots. Tangible Bots can reflect changes in the digital model and assist users by haptic feedback, by correcting errors, by multi-touch control, and by allowing efficient interaction with multiple tangibles. A first study shows that Tangible Bots are usable for fine-grained manipulation (e.g., rotating tangibles to a particular orientation); for coarse movements, Tangible Bots become useful only when several tangibles are controlled simultaneously. Participants prefer Tangible Bots and find them less taxing than passive, non-motorized tangibles. A second study focuses on usefulness by studying how electronic musicians use Tangible Bots to create music with a tangible tabletop application. We conclude by discussing the further potential of active tangibles, and their relative benefits over passive tangibles and multi-touch.

Author Keywords

Tangible user interfaces, bidirectional interfaces, active tangibles, user evaluation

ACM Classification Keywords

H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems. H.5.2 [Information Interfaces and Presentation]: User Interfaces.

General Terms

Experimentation, Human Factors

INTRODUCTION

Tabletop Tangible Interfaces (TTIs) [23] have emerged as a research topic in the field of tangible computing. TTIs combine the dynamic qualities of digital presentation with the qualities of physical, tangible objects. This powerful combination is illustrated in a TTI system like Urp [24]. In Urp,

the user moves physical models of houses around a top-projected table surface to observe changes in sunlight and shadow. Houses in Urp serve as physical interaction handles and changes made to the spatial configuration of houses are instantly detected and reflected in the digital model. The detection of physical interactions and the digital change that follows are inseparable to users and leave the impression of a unified physical and digital system.

The physical objects in TTIs in general, and also in Urp, are inert. If the digital model changes unrelated to physical input, it would lead to inconsistency between the physical and digital model. For instance, if a house in Urp were assigned a different position in the digital model, its physical position would remain unchanged. The impression of unified input/output would thus be lost. We use the term unidirectional interfaces to categorize TTIs that employ this one-way interaction model and *passive tangibles* as a term for the inert tangibles of such systems. Most TTIs use passive tangibles (e.g., [3, 8, 16, 17, 24]).

In contrast, some researchers have attempted to alleviate inconsistency by creating bidirectional interfaces with *active tangibles*. Active tangibles are capable of moving or otherwise reflecting changes in the digital model. Such tangibles have been used in TTIs to preserve consistency be-


Figure 1. A passive tangible (bottom-left) next to a Tangible Bot (top-right) on a tabletop tangible interface. Tangible Bots are capable of moving and rotating.

tween physical interfaces and virtual online interfaces [19], between remotely coupled TTIs [2, 20], and between TTIs and their underlying digital model [11, 14, 21]. However, little research has explored the potential for new interaction techniques that emerge from using active tangibles.

The present paper makes three contributions to the field of tangible tabletop interfaces. First, we introduce a set of novel interaction techniques for bidirectional interfaces that uses active tangibles. The interaction techniques support haptic feedback for providing physical interaction guidance, introduce techniques for controlling multiple tangibles single-handedly, and allow recording and replaying motions and gestures. Second, we describe an implementation of active tangibles using motorized tangibles, Tangible Bots. Third, we conduct two user studies focusing on the usability and usefulness of active tangibles. Based on the results, we describe the advantages and challenges of using active tangibles on tabletop interfaces.

RELATED WORK

Existing work on bidirectional TTIs can be divided into two groups based on the approach to making the tangibles active. The first group moves tangibles using an electromagnetic array situated under the tabletop surface [2,14,15,25], whereas the second group uses small motorized tangibles that move on the tabletop surface [11,19–21]. We name the two groups electromagnetic and motorized interfaces, respectively.

Brave and colleagues were the first to present an *electromagnetic interface*, PSyBench [2]. PsyBench was constructed from two modified chessboards, which allowed movement on one board to be reproduced on the other board. The 10×8 grid chessboard only supported discrete movement of tangibles into one of the 80 positions. The manner in which tangibles moved was therefore unlike human motion. This issue was addressed in the Actuated Workbench [14]. Pangaro and colleagues were able to place tangibles in between magnets by employing anti-aliasing techniques and activating several electromagnets at once. In PICO [15], the magnetic array was extended to a 16 x 32 grid, resulting in a 30.5 x 61 cm active surface.

An advantage of electromagnetic interfaces is that tangibles do not require power, giving them unlimited operational time. However, electromagnetic interfaces also have drawbacks. Objects tend to move in jerks, resulting in an unnatural and abrupt motion, and the electromagnetic arrays are only capable of moving small lightweight objects. The magnetic arrays are expensive and assembling requires high-level skills in electronics. Moreover, the number of magnets and control boards required to move the tangibles grow with the size of the tabletop, making these interfaces expensive for larger surfaces.

In *motorized interfaces*, small battery-powered tangibles move themselves around the surface. The tangibles are either steered wirelessly by the tabletop computer [19, 21] or move autonomously by following paths or indicators projected visually onto the surface [11, 20]. Typically, two-wheeled custom built robots are used as tangibles, but robots assembled from LEGO's Mindstorm[TM] elements have also been used. The robots used in [21] measure 68×63 mm and have the motor power to move furniture models around the 100×100 cm surface. When working with motorized interfaces, research in the field of Human-Robot Interaction provide several relevant findings to draw upon (i.e., [10,13,26]).

An advantage of motorized interfaces is that existing tabletop systems can be used. Also, pre-assembled robots can be bought off-the-shelf and are getting increasingly cheaper and smaller. Motorized tangibles move and rotate smoothly and are able to repeat human motions very accurately [5]. However, the use of robots as active tangibles introduces several challenges: Motorized tangibles are battery-powered and hence only operational for a limited period of time. Moreover, the wheels of motorized tangibles may limit the freedom of movement for tangibles, as they need to rotate themselves before moving sideways.

Though electromagnetic and motorized interfaces differ in their means of moving tangibles, they address similar research topics. For both groups of interface, the main use of active tangibles has been to allow remote users to collaborate in a physically distributed interaction space. This concept was first explored with PSyBench [2], which allowed changes made to the spatial configuration of tangibles at one interface to be reproduced at a geographically remote user interface. Other systems have used active tangibles to couple physical and virtual models [19]: A TTI was used by engineers to plan the layout of machines in a physical factory model. Remote engineers could follow the interaction in a virtual environment and intervene by dragging the elements in the graphical user interface.

To our knowledge, the research on the use of active tangibles as a way to improve usability by allowing new interaction techniques is limited to [15, 25]. In PICO [15], magnetic attraction and repulsion were used to provide haptic feedback to the users. The haptic feedback was used to guide users as they interact and was found helpful to users when solving spatial optimization problems (e.g., placement of cell-phone towers). In [25], repulsion and attraction were used to allow actuated tangible widgets, called Madgets. Madgets are small physical widgets (i.e., knobs, buttons) that can be placed on the tabletop surface when needed. By activating magnets inside the Madgets, the interface is able to manipulate them and for instance turn a Knob Madget (a physical knob that, when placed on the tabletop, may be turned by both the interface and the user). PICO and Madgets are electromagnetic interfaces. We are not aware of discussions of the general potential of interaction with motorized tangibles.

In sum, interaction techniques for active tangibles seems under-researched and we lack insight into how they may be used to interact with TTIs and what the potential benefits and drawbacks of such interaction might be. Next we therefore outline a first set of interaction techniques for active tangibles. The implementation, usability and usefulness of the techniques form the content of the remainder of the paper.

INTERACTION TECHNIQUES

This section proposes a set of active tangible interaction techniques for bidirectional interfaces. The interaction techniques are divided into four categories: interaction feedback, interaction commands, group interactions, and model-based interactions. This organization is not intended as a taxonomy for active tangible interactions and the list of interaction techniques is not meant to be exhaustive. Our primary objective is to provide a first set of interaction techniques that use active tangibles, with which we can conduct empirical studies. The interaction techniques presented here are designed for small motorized robots that are capable of moving and rotating, and are intended for use on TTIs that have multitouch capabilities.

Interaction feedback

Haptic feedback
Unidirectional TTIs only allow visual and auditory feedback to the user. This makes it difficult for users to interact with multiple tangibles simultaneously while keeping track of their orientation and settings. By controlling the motor, active tangibles may provide users with *haptic feedback* as they interact. For instance, if active tangibles are used for selecting a value on a scale, the motors can create resistance when the user approaches the end points of the scale. The mechanical qualities of panning knobs like the ones found on stereos and mixers can also be simulated by having the tangibles snap to certain angles (see Figure 2). Moreover, repulsion and attraction can be used to guide users as they interact, which could be helpful in systems that constrain the spatial configuration of tangibles (i.e., [8,17]). Attraction and repulsion were explored in PICO [15], where it provided useful guidance to users as they solved spatial optimization problems; the interaction techniques proposed here may be seen as a generalization of those in PICO.


Figure 2. Haptic feedback: The tangibles provide interaction guidance to the user by snapping tangibles to certain angles (left) or by simulating attraction and repulsion (right).

Interaction assistance
Ullmer and Ishii [23] compare TUIs with traditional board games. Both employ physical objects and require the users to be familiar with the rules of the game and the meaning associated with the symbolic artifacts used for playing. With active tangibles it is possible to reposition tangibles that violate syntactic or semantic constraints of the interface (see Figure 3). We propose the term *interaction assistance* to describe this kind of feedback on users' manipulation of tangibles.


Figure 3. Interaction assistance: A user has placed a tangible in an area that does not allow tangibles. The interface provides interaction assistance by moving the misplaced tangible to an acceptable position.

Interaction commands

Indirect interaction
Active tangibles enable users to move and rotate tangibles on the surface without directly moving them by hand. We use the terms *indirect movement* and *indirect rotation* to describe this type of interaction, illustrated in Figure 4. Here, a user draws a movement path on the surface, which the robot will follow when the user stops drawing. Also, it is possible to support more complex commands such as "recall last position" or "align tangibles".


Figure 4. Indirect interaction: A user moves a tangible with indirect movement by drawing a path on the surface (left) and rotates a tangible by tapping (right).

Imitation
With passive tangibles the number of tangibles that can be manipulated simultaneously is limited to the number of hands present at the table. With *rotation imitation* and *movement imitation* the user can record motions and have the interface play them back (Figure 5). In musical systems for live performances [3, 8, 16, 17], as an example of an application, electronic musicians could manipulate effect parameters once and have the interface imitate the manipulation repeatedly.


Figure 5. Imitation: A user moves a tangible in a circular motion and the active tangible imitates the movement.

Group interactions

Active tangibles allow users to group multiple tangibles and control them by interacting with a single member of the group. As illustrated in Figure 6, the interaction commands from the previous section can also be applied to groups.

Here, a user rotates all tangibles to the center position by applying indirect rotation to a single tangible. In addition to the indirect interaction commands, the user can move and rotate a group of tangibles by moving or rotating a member of the group by hand. We use the terms *direct movement* and *direct rotation* for this type of interaction. With group interactions, the maximum number of tangibles that can be manipulated simultaneously is not limited to the total number of hands at the table. In musical TTI systems like mixiTUI [17] and Reactable [8], direct rotation would allow musicians to rotate multiple tangibles in parallel and, for instance, fade all instruments.


Figure 6. Group interactions: Users can control multiple tangibles at once by applying interaction commands to a group.

Model-based interactions

In TTIs, tangibles are used for manipulating the digital model. When using passive tangibles, conflicts between the digital model and the physical model can occur, as changes made to one tangible cannot be reflected by the spatial layout of the other tangibles [7]. As illustrated in Figure 7, active tangibles can be used to ensure physical/digital consistency in TTIs. In the figure, a user rotates a tangible to adjust the temperature in Celcius. The left tangible, which displays temperature in Fahrenheit, rotates proportionally to ensure consistency between the two temperature scales. We use the terms *model-based rotation* and *model-based movement* for these types of interaction.


Figure 7. Model-based interactions: A user rotates a tangible to adjust the temperature in Celsius. The left tangible rotates proportionally to ensure consistency between the two temperature scales.

DESCRIPTION OF HARDWARE

In this section we describe Tangible Bots, the active tangibles used in this paper, and the tabletop setup used in the two studies. Figure 8 shows all system components.

Tabletop

The table used relies on a two camera DSI setup [22]: one camera for tracking of the active tangibles (with reacTIVision [9]) and one for touch tracking (with tbeta[1]). The 80 cm×45 cm surface is illuminated from underneath by a projector running at 1280×720 resolution. A XBee module has been connected to the tabletop computer so as to allow wireless communication with the Tangible Bots.
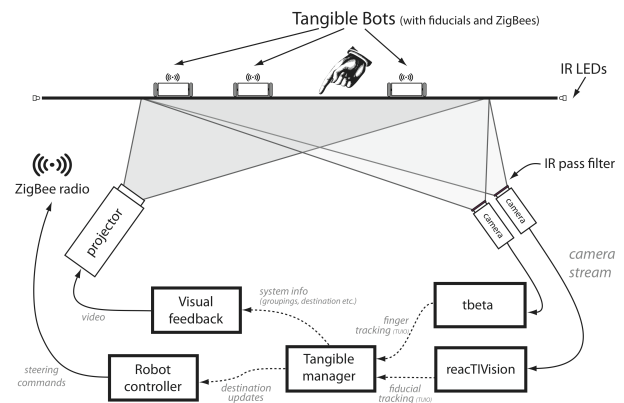
[1]http://ccv.nuigroup.com/

Figure 8. Overview of the components in the system used for the studies.

Tangible Bots

We use small Pololu 3pi robots[2] ($\varnothing$=9.5 cm) as active tangibles. The robots are equipped with a programmable Atmel ATmega328P micro controller, and have had a XBee module added for wireless communication. In order to allow tracking of the robots, a reacTIVision fiducial marker has been attached to the bottom of the robot. Tangible Bots use 4xAAA rechargeable batteries and can stay powered for approximately three hours with normal use. Tangible Bots move at a speed of 24.5 cm/s and rotate at a speed of 220 deg/s.

Tracking and movement computations are performed by the tabletop computer; the Tangible Bots hold no positional information. The computer continuously tracks the orientation and speed of each Tangible Bot, determines the proper action and communicates this action to the robots one at a time. The tabletop computer steers the Tangible Bots by sending movement frames, which are executed by the tangibles. Each movement frame contains three values: (1) speed value for right motor, (2) speed value for left motor, and (3) movement time in milliseconds. The movement time helps avoid unwanted autonomous movement in case of package loss or package delay.

IMPLEMENTATION OF INTERACTION TECHNIQUES

To evaluate the interaction techniques proposed earlier, we designed an application that allowed users to interact with six Tangible Bots. Next, we describe how the interaction techniques have been implemented.

Haptic feedback

If a user turns a tangible below its minimum or maximum setting (indicated by the arc in Figure 9), the tangible will activate its wheels to oppose the motion of the user (interaction guidance).

Interaction commands

Users rotate the tangibles indirectly by tapping a point on the arc projected around tangibles or by dragging their finger

[2]http://www.pololu.com/

Figure 9. The application used in Study #1. Here, a user moves two tangibles using direct movement. The line and the circle show the movement path and destination of the Tangible Bot.


Figure 10. The four tasks in Study #1. In movement tasks participants had to move all tangibles from one position to another, preserving the initial spatial configuration (DM) or to align the tangibles on a straight line (IM). In rotation tasks participants had to add/subtract a value to/from the values shown above the tangibles (DR) or rotate all tangibles to the same value (IR). Target settings are shown for DR and IR as small red circles.

along the arc. To align the tangibles, users select a group of tangibles and tap a dotted line in the middle of the screen. Imitation recording is activated and deactivated by tapping twice with two fingers below a tangible. After having tapped to activate imitation recording, the movements and rotations are recorded and subsequently replayed.

Group interactions
We have implemented three gestures for forming groups, based on recommendations from Micire et al. [13]: A user can group two tangibles by placing a finger below one tangible and double tap below another. Moreover, users can group multiple tangibles by lassoing them with a finger or by forming a bounding box around the tangibles with two fingers. Besides the interaction commands mentioned above, users can move or rotate a group of tangibles by applying direct rotation or movement to a group member. The projected movement path and the destination of each Tangible Bots are displayed as a line and a circle, respectively (see Figure 9).

STUDY #1: USABILITY
Study #1 aims at investigating how active tangibles affects interaction speed and precision. The goal of the study was to investigate how many tangibles are needed for the interaction techniques to be advantageous with active tangibles and to get general feedback on the interaction techniques and Tangible Bots.

Participants
Participants were recruited among students at our department using invitations by e-mail and posters. A total of 16 persons participated. All were male and aged between 26 and 43 (M = 31.5). Thirteen were right-handed and three were left-handed. The participants received no payment or compensation for their participation.

Apparatus
The study was done using six Tangible Bots and the hardware and software previously described.

Procedure
First, participants were welcomed and explained the purpose of the evaluation. Then participants were asked to complete a set of four tasks using both active tangibles and passive tangibles. The interaction techniques were explained to the participants before each new task, and participants were given the time needed to get confident with them. The tasks, which asked participants to move or rotate tangibles to specified positions or angles, where displayed textually in a box in the upper middle part of the screen. Completion time was measured as the time span from the moment a participant tapped a button labelled "Begin" until the moment when all tangibles had been correctly positioned and a button labelled "Finish" had been tapped. If users tapped "Finish" without having correctly positioned all tangibles, an error was logged.

After completing the task set with either type of tangible, participants rated the interaction with a questionnaire. Finally, the participants were interviewed and asked to compare aspects of the interaction with passive and active tangibles. The experiment was recorded using two cameras: One captured the facial expression of the participants, the other captured the surface of the tabletop.

Experimental design
The experiment used a within-subject design with three independent variables, described below.

User interface (UI). The experiment evaluated two user interfaces. With the active tangibles UI (AT), participants could group and control multiple tangibles, whereas the passive tangible UI (PT) required participants to move tangibles manually. Robots were used as tangibles in both AT and PT to maximize similarity between conditions.

Task (T). The experiment required participants to carry out a task set using both AT and PT. The task set consisted of four tasks that each evaluated separate interaction tech-
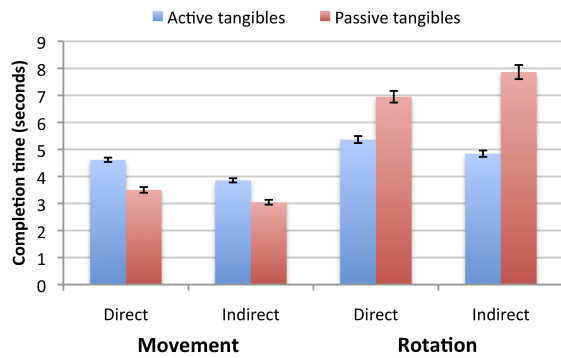
Figure 11. Completion time (+/- SEM) per task

Figure 12. Completion time (+/- SEM) per number of tangibles and UI for direct and indirect movement tasks (DM, IM).

niques: Direct movement (DM), indirect movement (IM), direct rotation (DR), and indirect rotation (IR). Figure 10 illustrates the four tasks. In the movement tasks, participants had to move all tangibles from one position to another, preserving their initial spatial configuration (DM) or aligning the tangibles on a straight line (IM). In the rotation tasks, participants were asked to add/subtract a value to/from the values of the tangibles (DR) or rotate all tangibles to the same value (IR).

Number of tangibles (NT). The tasks that evaluated the indirect interaction techniques (IM, IR) were completed four times with a changing number of tangibles (NT = 1, 2, 3, 6), whereas the direct interaction tasks (DM, DR) were only completed three times (NT = 2, 3, 6). The reason is that the direct interaction techniques are only applicable to groups of tangibles, that is, more than one tangible. The constants were chosen to compare active and passive interactions when participants were able to manipulate all tangibles at once (NT = 1, 2), had one more tangible than they could control with two hands (NT = 3), or had far too many tangibles to manipulate at once (NT = 6).

The starting UI alternated between AT and PT. The task order was shifted systematically for each participant, whereas the order of NT was randomly determined. For each combination of user interface, task, and number of tangibles, the participants were asked to carry out five trials. In summary, the experiment consisted of: 2 UI × (2 T × 3 NT + 2 T × 4 NT) × 5 trials = 140 trials per participant.

Dependent measures
Completion time and error rate were logged for each trial for later statistical analysis. In the two questionnaires following each UI, the participants rated the interaction using six questions from QUIS [4]: Wonderful/horrible, easy/difficult, satisfying/frustrating, flexible/rigid, fun/boring, clear/confusing. A 9-step scale was used. Moreover, the users rated their perceived mental and physical task load using questions from NASA TLX [6].

STUDY #1: RESULTS
We found a main effect of UI on task completion time ($F_{1,15} = 9.48, p < .01$), and an interaction between UI
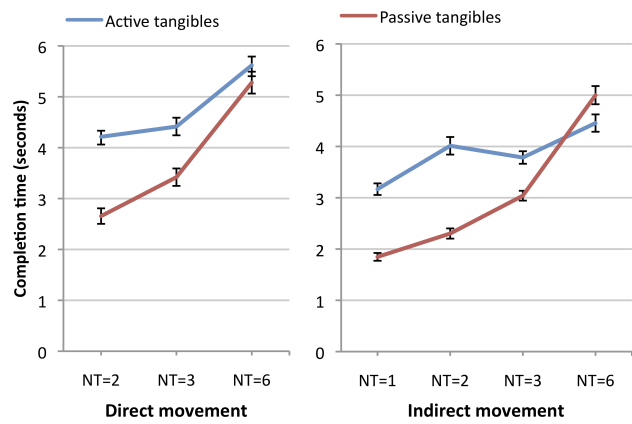
and task ($F_{3,13} = 64.61, p < .0001$). Post-hoc tests showed clear differences: Active tangibles increased task completion time for movement tasks (DM, IM) and reduced task completion time for rotation tasks (DR, IR), see Figure 11. On average the completion time for DM and IM increased by 1113 ms (31.8%) and 808 ms (26.5%), respectively, when using active tangibles. For DR and IR the completion time was reduced with 1583 ms (22.8%) and 3019 ms (38.4%), respectively.

The error rate was very low (1.6% of all trials) and did not differ significantly between active and passive tangibles.

As excepted, there was an UI x NT interaction ($F_{3,13} = 96.28, p < .0001$). Post-hoc tests showed that more tangibles increased task completion time.

Movement
Figure 12 shows the average completion time per NT and UI for movement tasks (DM, IM). Post-hoc pairwise comparison (Bonferroni adjusted) showed that for for NT ≤ 3, participants completed the movement tasks significantly faster when using passive tangibles ($p < .05$). For NT = 6, the data showed no significant time difference for neither direct movement nor indirect movement.

The difference in completion time between active and passive tangibles decreased as the number of tangibles increased, suggesting that active tangibles would be fastest if NT was increased sufficiently. The participants were much faster at moving tangibles than rotating them. Analysis of videos from the experiment show that most participants used both hands simultaneously when moving tangibles, resulting in a maximum of three steps for all movement tasks.

Rotation
Figure 13 shows the average completion time per NT and UI for rotation tasks (DR and IR). For direct rotation, active tangibles result in significantly faster completion times for NT ≥ 3 ($p < .01$), whereas active tangibles are faster for indirect rotation for NT ≥ 2 ($p < .001$). For NT = 6, active
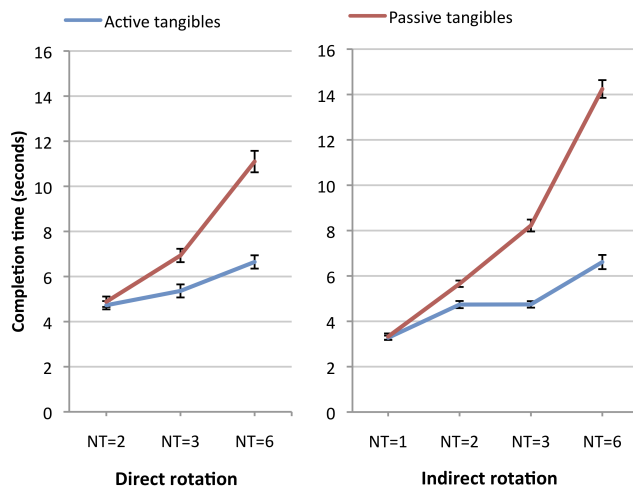
Figure 13. Completion time (+/- SEM) per number of tangibles and UI for direct and indirect rotation tasks (DR, IR).

Figure 14. Results from QUIS and TLX questions. For TLX, less is better.

tangibles on average resulted in 1.66x and 2.15x faster task completion times for DR and IR tasks, respectively.

Analysis of video from the experiment show that participants only rotated one tangible at a time when using passive tangibles. Initially, some participants would try rotating two tangibles simultaneously, but would abandon this strategy after one or two trials.

Grouping
When doing tasks with two or more active tangibles, participants had to form groups before rotating or moving. On average, this step took 1661 ms (SD = 840 ms). Because we only asked participants to perform a single action after having formed a group, the grouping step accounted for a large amount of the total completion time (34.9%). Users most frequently formed groups by using the bounding box technique mentioned earlier. This technique was used in 97.2% of the trials, and showed to be the fastest way to form groups.

Satisfaction and comments
Figure 14 show the results from the questionnaires. We used multivariate analysis of variance to analyze the QUIS and TLX questions and found a main effect for UI ($F_{1,15} = 39.53, p < .001$). Pairwise comparison on each measure showed that participants found it significantly easier to interact with active tangibles ($p < .05$). Even though participants controlled multiple tangibles at once when using active tangibles, they did not find this more confusing. Participants were significantly more satisfied with active tangibles ($p < .001$) and found them to be both more wonderful ($p < .001$) and fun ($p < .005$) to use. Active and passive tangibles were perceived as being equally flexible. Additionally, participants generally found active tangibles to be significantly less physical demanding to use ($p < .001$) and found them to lower the mental task load ($p < .001$).

Fifteen of the participants (88%) preferred active tangibles and believed that they had completed the task set fastest with
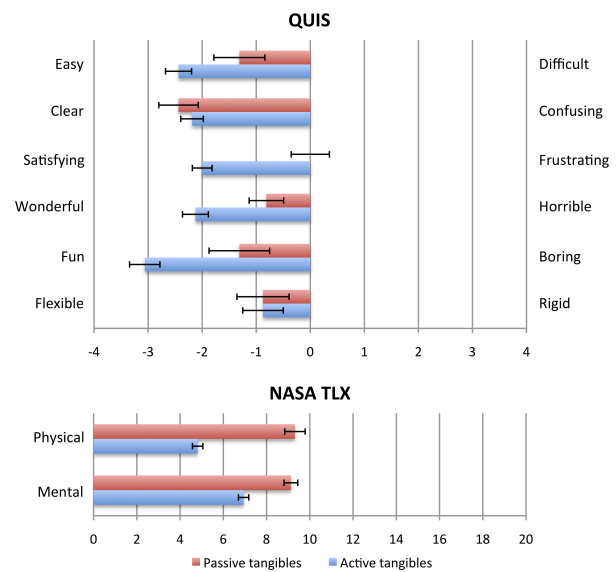
this UI. Participants on average found it to be an advantage to use grouping for NT≥ 3 (SD = 0.8). Though participants noted that the gain of using active tangibles was much higher with the rotation techniques than with the movement techniques, they believed that they had completed the movement tasks fastest with active tangibles.

STUDY #2: USEFULNESS
Study #2 aims at investigating how the active tangible interaction techniques support expert users. We invited seven electronic musicians to evaluate a bidirectional version of mixiTUI, a tangible sequencer for electronic live performances [17] (see Figure 15). mixiTUI is originally an unidirectional interface, but had been modified to integrate a subset of the interaction techniques described earlier. We wanted to observe how the musicians would react to an interface involving active tangibles: Do active tangibles increase the complexity of interactions? How do the indirect interactions affect the physicality and tangibility of the system? What kind of new musical expressions emerge?

Domain and participants
We choose musical performance as an evaluation domain as tangible interactions have shown to suit musical performance well [8, 16, 17]. Also, musicians are very critical users: When interacting with music systems they are not only concerned with whether systems allow them to get from A to B, but also how getting from A to B feels and sounds. The seven musicians participating had 2-10 years of experience with live performances of electronic music. Three of the musicians had participated in a previous evaluation of mixiTUI.

Interaction design
In mixiTUI, musicians create music using *loop tokens* and *effect tokens*. Loop tokens are the sound generating part of

Figure 15. An electronic musicians plays a bidirectional version of mixiTUI. Loop tokens (white, at left and top) allow musicians to arrange and mix precomposed samples. Effect tokens (dark, at right and bottom) alter the sound of the loop tokens.

mixiTUI and allow musicians to play and mix precomposed samples, say, a drum beat. By placing an effect token above a loop token, the musician can alter the sound of that loop. The routing of sound signals is shown with waveforms connecting tokens. Volume levels of loop tokens and the parameters of effect tokens are shown as arcs around the tokens. Moreover, the duration and progress of each loop are shown in terms of beats and bars in order to assist the musician in keeping track of time.

In the bidirectional version of mixiTUI, *haptic feedback* was used to prevent musicians from turning both loop tokens and effect tokens below minimum or above maximum. *Rotation imitation* allowed musicians to record and playback rotation patterns, thereby creating automation of effect parameters on the fly. Imitation was activated by double tapping with two fingers below an effect token. With *movement imitation* musicians could specify a movement path for the effect tokens and, for instance, have it move from loop token to loop token. Finally, musicians could group loop tokens together and use *direct rotation* and *indirect rotation* to adjust the volume of several loop tokens simultaneously.

Procedure
To allow the musicians to play their own music, mixiTUI had been set up with loops provided by the musicians beforehand. In the first part of the user study, musicians were presented with the unidirectional version of mixiTUI. This allowed them to get a feel of its functionality and to get confident with the basics of performing electronic music on mixiTUI. In the second part of the user study, the musicians were presented with the bidirectional version of mixiTUI. After having learned the interaction techniques, the musicians played for thirty minutes, during which they were asked to comment on the system and propose ideas for new functionality and new interaction techniques. The comments were transcribed by the evaluator during the evaluation. Finally, to emulate the pressure of a real live concert and to encourage the musicians to challenge the boundaries of the system, they were asked to perform a three minute piece of their choice, using loops from their own music. The perfor-

mance was recorded using a single camera that captured the surface of the tabletop.

STUDY #2: RESULTS
Results for Study #2 are based on two sources of data. First, we analyzed the videos showing each of the musician's performance. We focused on their use of interaction techniques, analyzing how Tangible Bots were used in creating music, and on which interaction techniques musicians used for what purpose. Analysis of video was done collaboratively between the authors to identify significant themes (as in [1]). After that, themes were checked by recoding the videos, and frequencies of observations were obtained. Second, we analyzed the comments of musicians as they where learning to use the system and during the debriefing. These were abstracted using meaning condensation [12], and themes were identified in collaboration between the authors.

All musicians found haptic feedback useful, and felt that it helped them avoid turning tangibles below their minimum or above their maximum setting. One musician explained that haptic feedback added "the feel of a physical knob" to the tangible. In the videos from the evaluation, it was difficult to identify and count when haptic feedback was activated and useful to the musician.

All musicians used rotation imitation in their performance, four more extensively than the rest. These musicians would have several tangibles playing back imitation repeatedly, while interacting with loop tokens. One musician reported a feeling of "having an extra pair of hands". The motions recorded were either small fine-grained adjustments or sweeps from one endpoint of the scale to the other. In the debriefing, one musician noted that rotation imitation allowed him to create greater musical variation in his performance, as it applied constant changes to the otherwise static loops. Two musicians made reports to similar effect. The musicians did not find movement imitation to be useful, and none of them used it during their performance. We suspect that the reason for this is that mixiTUI is primarily operated through rotation.

We observed two distinct uses of grouping of Tangible Bots. First, musicians formed temporary groups to allow multiple loops to be faded in or out singlehandedly. After having completed the fade, musicians would immediately break up the group. Second, musicians grouped together musically related loops (i.e., a drum loop and a percussion loop) in order to mute a group of instruments quickly or to crossfade between groups of instruments. Musicians would typically preserve musically related groups throughout their performance. Thus, group interactions would be both related to short, efficient manipulations and to longer-term, semantically meaningful groupings.

Six musicians used direct and indirect rotation alternately when adjusting the volume of loops. By analyzing the videos from the evaluation, we found a general pattern as to when they chose one over the other. Indirect rotation was used for discrete volume changes where the musical outcome could

be predicted (i.e., turn volume from min to max). One musician explained that he regarded indirect manipulation as a "shortcut" that allowed him to carry out volume changes more quickly. When performing delicate adjustments that required musicians to listen to the change in music and interact alternately, they typically chose direct rotation. On several occasions we noticed that the musicians mixed direct and indirect interaction to accomplish multitasking. As an example, one musician recorded imitation while turning down the volume of two loops using indirect rotation.

In the debriefing, we asked musicians to compare the complexity of passive and active tangibles. Two musicians compared the learning curve of the active tangibles to that of a traditional music instrument. They felt that the active tangibles, though more complex to master, allowed them to create a more complex and variated musical performance.

DISCUSSION
Our paper has introduced Tangible Bots, an implementation of active tangibles on tabletop tangible interfaces using simple, off-the-shelf robots. We have contributed and evaluated a set of interaction techniques for Tangible Bots, including support for (a) feedback during interaction, (b) command-based control of tangibles, (c) grouping of tangibles, and (d) model-based interactions. The usability study (Study #1) showed that rotation of tangibles (relatively fine-grained movements) benefitted from the grouping and command-based control of Tangible Bots. Movements of six tangibles were performed equally quickly with Tangible Bots and passive tangibles, but for fewer tangibles, coarse movements of passive tangibles were more efficient. Participants preferred active tangibles, and found that active tangibles required less workload and were more fun and satisfying to use. The study of usefulness (Study #2) was exploratory and focused on electronic musicians. It showed that active tangibles could be used in a simulated live performance in a satisfactory manner. Active tangibles supported the performance by allowing musicians to manipulate several tangibles singlehandedly (through grouping) and to record and replay movements through imitation. Feedback during interaction was shown useful in Study #2, whereas model-based interactions were not evaluated.

Our results provide initial evidence that Tangible Bots and the associated interaction techniques are usable for general-purpose manipulation of tangibles and useful to electronic musicians using tabletop tangible interfaces. In the following, we discuss the two most successful interaction techniques.

The study of musicians shows that group interactions were supporting efficient interaction, for instance by helping musicians create ad-hoc groups for efficient manipulation. Grouping, however, was not only used with efficiency in mind. The musicians also formed semantically meaningful groups based on the similarity of instruments. Typically, these groups comprised two active tangibles, which could have been manipulated by one hand each (i.e., without grouping). This observation suggests that group interaction

can also be used to organize and to overview, as the number of tangibles grow. Overall, the ability to group tangibles seems to change interaction with tabletops in important ways.

Both studies show that direct manipulation of tangibles and indirect manipulation were often intertwined. We believe that an important contribution of our work is to show how tabletop interfaces may be extended to use both hands-on manipulation of tangibles (as in other tangible user interfaces) as well as indirect manipulation using multi-touch. This observation raises a fundamental question about tangible interfaces: When is the direct, physical manipulation of a tangible preferable, and when is the indirect manipulation of for instance a group of tangibles preferable? We have introduced in tabletop tangible interfaces a form of command-based, indirect manipulation that may be seen as the antithesis to the original idea of Tangible Bits [7]. Our empirical data help discuss this seeming tension. The study of musicians showed that some types of manipulation were done indirectly, for instance adjusting a tangible to a specific value. Such manipulations were used when the setting were well known (e.g., turn off). Conversely, musicians often chose to do direct, physical handling of a tangible when they were listening for how their changes affected the music. Study #2 showed musicians to be fairly consistent in when they chose which interaction approach (direct manipulation of a tangible, indirect manipulation through multi-touch): One interpretation of this consistency is that these approaches supplement each other, rather than being mutually exclusive.

The present work is limited in several ways, for which we aim future work to compensate. First, the hardware and software comprising Tangible Bots could be improved. In particular, omni-directional and smaller robots would remove some limitations of the Pololu robots. Second, the study of usefulness is closely tied to the practice of performing electronic music. While electronic music is a widely used testbed for tangible interfaces [8, 16, 17], future work should look at the use of Tangible Bots in other domains and with other groups of professional or recreational users. Finally, we consider the potential of model-based interactions with Tangible Bots particularly promising; we are eager to see empirical data on the usefulness of this type of interaction.

CONCLUSION
Active tangibles in tabletop interfaces promise to reduce conflicts between the digital model and the spatial layout of tangibles, and to afford new ways of interacting. We have presented four groups of interaction techniques for active tangibles and implemented them with Tangible Bots.

We evaluated the usability of Tangible Bots in an experiment requiring participants to do a variety of rotation and movement tasks with one to six robots. The study showed that rotation benefited from Tangible Bots, whereas coarse movements could be done more quickly with non-active tangibles. A second study investigated the usefulness of Tangible Bots to electronic musicians. It showed how active tangibles changed the interaction and the music created.

Active tangibles allow for new interaction techniques that we think are promising as an extension of research on tangible user interfaces. New questions arise, however, in particular about the relative benefits of passive tangibles, physical manipulation of active tangibles, and indirect manipulation of groups of tangibles.

REFERENCES

1. Aronson, J. 1994. A pragmatic view of thematic analysis. Qualitative Report 2:1

2. Brave, S., Ishii, H., and Dahley, A. 1998. Tangible interfaces for remote collaboration and communication. In *Proc. CSCW 1998*. ACM, New York, 169-178

3. Bischof, M., Conradi, B., Lachenmaier, P., Linde, K., Meier, M., Pötzl, P., and André, E. 2008. Xenakis: combining tangible interaction with probability-based musical composition. In *Proc. TEI 2008*. ACM, New York, NY, 121-124

4. Chin, J. P., Diehl, V. A., and Norman, K. L. 1988. Development of an instrument measuring user satisfaction of the human-computer interface. In *Proc. CHI 1988*. ACM, New York, 213-218.

5. Frei, P., Su, V., Mikhak, B., and Ishii, H. 2000. curlybot: designing a new class of computational toys. In *Proc. CHI 2000*. ACM, New York, 129-136

6. Hart and Staveland. Development of NASA-TLX: results of empirial and theoretical research. In Hancock, P.A., Meshkati, P. (Eds.), *Human Mental Workload*, Elsevier, Amsterdam. 139-183

7. Ishii, H. and Ullmer, B. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proc. CHI 2007*. ACM, New York, NY, 234-241.

8. Jordà, S., Geiger, G., Alonso, M., and Kaltenbrunner, M. 2007. The reacTable: exploring the synergy between live music performance and tabletop tangible interfaces. In *Proc. TEI 2007*. ACM, New York, 139-146

9. Kaltenbrunner, M. and Bencina, R. 2007. reacTIVision: a computer-vision framework for table-based tangible interaction. In *Proc. TEI 2007*. ACM, New York, 69-74.

10. Kato, J., Sakamoto, D., Inami, M., and Igarashi, T. 2009. Multi-touch interface for controlling multiple mobile robots. In *Proc. CHI 2009*. ACM, New York, NY, 3443-3448.

11. Kojima, M., Sugimoto, M., Nakamura, A., Tomita, M., Inami, M., and Nii, H. 2006. Augmented Coliseum: An Augmented Game Environment with Small Vehicles. In *Proc. TABLETOP 2007*. IEEE Computer Society, Washington, 3-8

12. Kvale, S. 1996. Interviews: An Introduction to Qualitative Research Interviewing, Sage Publications, Thousand Oaks California

13. Micire, M., Desai, M., Courtemanche, A., Tsui, K. M., and Yanco, H. A. 2009. Analysis of natural gestures for controlling robot teams on multi-touch tabletop surfaces. In *Proc. ITS 2009*. ACM, New York, 41-48.

14. Pangaro, G., Maynes-Aminzade, D., and Ishii, H. 2002. The actuated workbench: computer-controlled actuation in tabletop tangible interfaces. In *Proc. UIST 2002* ACM, New York, 181-190

15. Patten, J. and Ishii, H. 2007. Mechanical constraints as computational constraints in tabletop tangible interfaces. In *Proc. CHI 2007*. ACM, New York, 809-818

16. Patten, J., Recht, B., and Ishii, H. 2002. Audiopad: a tag-based interface for musical performance. In *Proc. NIME 2002*. National University of Singapore, Singapore, 1-6.

17. Pedersen, E. W. and Hornbæk, K. 2009. mixiTUI: a tangible sequencer for electronic live performances. In *Proc. TEI 2009*. ACM, New York, 223-230

18. Piper, B., Ratti, C., and Ishii, H. 2002. Illuminating clay: a 3-D tangible interface for landscape analysis. In *Proc. CHI 2002. ACM*. New York, 355-362

19. Ressler, S., Antonishek, B., Wang, Q., and Godil, A. 2001. Integrating active tangible devices with a synthetic environment for collaborative engineering. In *Proc. Web3D 2001*. ACM, New York, 93-100

20. Richter, J., Thomas, B. H., Sugimoto, M., and Inami, M. 2007. Remote active tangible interactions. In *Proc. TEI 200*. ACM, New York, 39-42

21. Rosenfeld, D., Zawadzki, M., Sudol, J., and Perlin, K. 2004. Physical Objects as Bidirectional User Interface Elements. *IEEE Comput. Graph*. Appl. 24, 1 (Jan. 2004), 44-49

22. Schöning, J., Brandl, P., Daiber, F., Echtler, F., Hilliges, O., Hook, J., Löchtefeld, M., Motamedi, N., Muller, L., Olivier, P., Roth, T. and Zadow, U. von Schöning. 2008. Multi-touch surfaces: A technical guide. In *Technical Report TUM-I0833*, Technical Reports of the Technical University of Munich, (October 2008)

23. Ullmer, B. and Ishii, H. 2000. Emerging frameworks for tangible user interfaces. In *IBM Syst. J.* 39, 3-4 (Jul. 2000), 915-931

24. Underkoffler, J. and Ishii, H. 1999. Urp: a luminous-tangible workbench for urban planning and design. In *Proc. CHI 1999*. ACM, New York, 386-393

25. Weiss, M., Schwarz, F., Jakubowski, S., Borchers, J. 2010. Madgets: Actuating Widgets on Interactive Tabletops. In *Proc. UIST 2010*. ACM, New York

26. Yanco, H. A. and Drury, J. L. Classifying Human-Robot Interaction: An Updated Taxonomy. In *Proceedings of the IEEE Conference on Systems, Man and Cybernetics*, October 2004.