Investigating Cross-Device Interaction between a Handheld Device and a Large Display

Paay, Jeni; Raptis, Dimitrios; Kjeldskov, Jesper; Skov, Mikael; Ruder, Eric V.; Lauridsen, Bjarke M.

Published in:

CHI '17, Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems

DOI (link to publication from Publisher):

[10.1145/3025453.3025724](https://doi.org/10.1145/3025453.3025724)

Publication date:

2017

Document Version

Early version, also known as pre-print

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Paay, J., Raptis, D., Kjeldskov, J., Skov, M., Ruder, E. V., & Lauridsen, B. M. (2017). Investigating Cross-Device Interaction between a Handheld Device and a Large Display. In *CHI '17, Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems* (pp. 6608-6619). Association for Computing Machinery. <https://doi.org/10.1145/3025453.3025724>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Investigating Cross-Device Interaction between a Handheld Device and a Large Display

Jeni Paay, Dimitrios Raptis, Jesper Kjeldskov, Mikael B. Skov,
Eric V. Ruder and Bjarke M. Lauridsen

Department of Computer Science / Research Centre for Socio+Interactive Design
Aalborg University, Denmark
{jeni, raptis, jesper, dubois}@cs.aau.dk

ABSTRACT

There is a growing interest in HCI research to explore cross-device interaction, giving rise to an interest in different approaches facilitating interaction between handheld devices and large displays. Contributing to this, we have investigated the use of four existing approaches combining touch and mid-air gestures, *pinching*, *swiping*, *swinging* and *flicking*. We look specifically at their relative efficiency, effectiveness and accuracy in bi-directional interaction between a smartphone and large display in a point-click context. We report findings from two user studies, which show that swiping is both most effective, fastest and most accurate, closely followed by swinging. What these two approaches have in common is the ability to keep the pointer steady on the large display, unaffected by concurrent gestures or body movements used to complete the interaction, suggesting that this is an important factor for designing effective cross-device interaction with large displays.

Author Keywords

Handheld Devices; Large Displays; Touch; Mid-air Gestures; Kinect; Cross-Device Interaction.

ACM Classification Keywords

H.5.2. User Interfaces: Interaction styles

INTRODUCTION

As the number of interactive computing devices around us continue to grow, and commercial products and services begin integrating interactions across these more closely, there is an increasing interest in “cross-device” and “digital ecosystem” interaction, and a growing body of research. This includes a lot of work on new interaction techniques, for example [2, 5, 6, 9, 14, 18, 20, 33, 34, 38, 40], their use in extreme multi-surface environments [1, 8], and tools for their implementation and management [11, 12, 29]. It also includes work in a broader scope, such as how multi-device

ecologies naturally emerge in practice [17], the role of space and spatial interactions for multi-device scenarios [28], and how cross-device interaction can be characterized conceptually to inform design [22, 27, 36].

Yet, as argued by Marquardt et al. [22], although cross-device interaction should now be commonplace, it is still surprisingly difficult. This means that further research into making cross-device interaction useable and useful is still warranted and very relevant. Nacenta et al. [26] summarize the challenge of cross-device interaction very well by saying that one of the crucial requirements particular to cross-device or multi-display environments is the ability for the user to move objects from one device or display to another. This interaction is a frequent and fundamental part of any cross-device application, and therefore it is of great importance to find suitable techniques. As many candidate approaches already exist in the literature, we find it is important to begin investigating their relative strengths and weaknesses in order to refine these techniques in an informed way, and base new techniques on experience with existing ones. This is a big task because of the sheer number of approaches and specific techniques, and the many different types of tasks and usage scenarios reported. Nonetheless, it is an important one.

In this paper, we contribute to the building of knowledge on cross-device interaction approaches with an empirical investigation of four existing ones, combining mobile devices and mid-air gestures. Specifically, we have investigated the use of *pinching*, *swiping*, *swinging* and *flicking* for moving objects between a handheld device and a large display, looking at their effectiveness, efficiency and accuracy. The contribution of this research are our findings about the relative performance of these four existing approaches in the bi-directional interaction between these devices.

In the following, we provide an overview of related work in cross-device interaction research. We then present the four approaches that we have investigated, and describe the two laboratory studies we conducted in order to learn about and compare their strengths and weaknesses. This is followed by a presentation of our findings. The findings show that the swiping approach was most effective, fastest, and most accurate for interactions both *to* and *from* the display, closely followed by the swinging approach. We then discuss what these findings mean, and what their implications are for the design of cross-device interactions.

RELATED WORK

Much of the related work on cross-device interaction has its origin in the work of Rekimoto [32] who envisioned so-called “multiple-computer user interfaces”. In the following we present some of this related research.

Combining Mobile Devices with Secondary Displays

Much of the work on cross-device interaction involves the use of a mobile device, such as a smart phone, in combination with a secondary surface, either as a pointing device or as a complementary device that data can be exchanged with.

As an example of using mobile devices for pointing at a secondary display surface, Boring et al. [4] used a mobile phone to control a pointer on a large display, comparing three different techniques: buttons on the phone screen, tilting the phone, and moving the phone. Seifert et al. [34] likewise experimented with PointerPhone for interacting with a large display at a distance. Based on observations of use, and qualitative feedback from users, a series of design recommendations for such interactions were derived. Similarly, Rashid et al. [31] explored two different ways of interacting with elements on a large display using a handheld device, both representing an area of the large display on the handheld device where the user could then interact with it.

Looking specifically at coarse vs. fine grained interaction in pointing, Myers et al. [25] experimented with laser pointers as devices for target acquisition on large displays, either on their own or attached to a PDA or a toy gun. Findings from this study suggested that laser pointers are useful for coarse-grained interaction but should be combined with other techniques for fine-grained selection. Also investigating precision, Nancel et al. [27] compared two approaches to target acquisition on a wall sized display, using either a handheld device or head-tracking for coarse pointing, with precision pointing done through the handheld device’s touch screen. They also compared their approaches with state-of-the-art techniques, such as LaserGyro [37] and SmoothPoint [7], and found them to perform comparatively well.

As an example of using mobile devices for data exchange with a secondary display surface, Rekimoto’s “pick-and-drop” [32] allowed users to move items between wall-sized displays and PDAs, mimicking the physical action of moving pieces of paper between tables and pin-up boards. For a similar task, Dachsel et al. [6] presented a “throwing” approach where the user can put content onto a large display from a distance. Looking at both pointing and data exchange, Bragdon et al. [5] investigated cross-device interaction for collocated team collaboration where participants could share information between their individual mobile devices and laptops and a large display.

Interactions with Multiple Mobile Devices

In other related work, cross-device interaction has involved several mobile devices used concurrently, with research focusing on appropriate interaction techniques and viable ways of implementation. As an example, Rädle et al. [29]

present the HuddleLamp, which uses a camera for tracking several mobile devices as well as the user’s hands and gestures. Woźniak et al. [39] present Thaddeus, which is a combined mobile phone and tablet system allowing ad-hoc interaction with various types of datasets using multiple mobile devices. Also focusing on several mobile devices, Jokela et al. [16] compare three methods for moving virtual objects from mobile phones and tablets, looking at their efficiency, novelty, and learnability. Similarly, Skov et al. [35] compare six different interaction techniques for two-way exchange of virtual objects between personal mobile phones and a shared tablet in a card game scenario.

Extending the scope of mobile devices beyond PDAs and smartphones, Houben and Marquardt [11] present a toolkit for utilizing the in- and output capabilities of smart watches in cross-device interactions with secondary surfaces. Also, facilitating real-world cross-device interaction in ecologies of multiple mobile devices and secondary displays, Houben et al. [12] introduce a configuration tool, ActivitySpace, enabling users to integrate and work across mobile devices and secondary display surfaces.

Interactions with Large Displays

Much of the work on cross-device interaction with large displays builds naturally on related work on interactions with large displays. This work goes back to Bolt’s [3] seminal “Put-that-there” system from the early 1980s, where he demonstrated the use of pointing gestures in combination with voice commands for interaction with a large display. While most of this work has investigated new possible interaction techniques, others have investigated ways of measuring the physical side effects of mid-air gestural interaction, such as arm fatigue [10].

More recently Vogel et al. [37] experimented with different ways of pointing with one’s hand to acquire targets on a very large display. Comparing the performance of absolute and relative approaches to ray casting, it was found that absolute pointing was faster while relative pointing resulted in fewer errors. Also looking at pointing tasks, Mayer et al. [23] investigated precision without visual feedback and studied the use of three ray casting methods using the index finger, the forearm, and the eye-finger relation for input. In a study of user preferences, Jacobsen et al. [15] compared touch and mid-air gestures for large display interaction, revealing that most participants preferred touch when interacting in close proximity to the large display, while most preferred mid-air gestures when interacting at a distance.

Extending the scope of large display interaction beyond the pointing task, Hespanhol et al. [9] proposed a set of five mid-air gestures, *push*, *grab*, *dwell*, *lasso*, and *enclose*, for selecting and rearranging items on a large display, and studied their performance. Going beyond the physical boundaries of displays, Markussen et al. [19] investigated an interaction concept that extends the input space for mid-air pointing to include the immediate area around the display. Specifically, they compare off-screen pointing to on-screen

pointing and touch interaction, showing a number of performance benefits of their off-screen approach.

Investigating cross-device interaction on a much larger scale of implementation, Beaudouin-Lafon et al. [1] present a concept of “interaction instruments” for the so-called WILD Room, specifically designed to facilitate distributed multi-surface interaction with large datasets. Following on from this, Gjerlufsen et al. [8] present middleware for developing flexible interactive multi-surface applications, again using the WILD Room as their development platform

Models of Cross-Device Interaction

Complementing the technical and experimental work on cross-device interaction techniques, their implementation, and respective performance and usability, a smaller body of research is aimed at producing more conceptual guidance on cross-device interaction. As an example, Marquardt et al. [22] propose a design pattern, called gradual engagement, to inform and inspire future interaction design. As a design pattern, it synthesizes generalizable interaction strategies and provides a vocabulary for discussing design solutions. With a similar aim of providing higher level interaction models, Nacenta et al. [26] present a taxonomy that classifies techniques for cross-device object movement according to a number of dimensions. Finally, Sørensen et al. [36] divide interactions with multiple devices, or digital ecosystems, into four categories, and show that each of these has very different suitable types of interaction techniques. Going into detail with one of these four categories, Raptis et al. [30] investigate how Apple’s new “Continuity” facility, for cross-device interaction between MacOS and iOS, is experienced by people in practice.

FOUR CROSS-DEVICE TECHNIQUES

Based on the related work described above we have identified four different overall approaches to cross-device interaction between a handheld device and a large display. While they have different names in different studies, we describe them as *pinching*, *swiping*, *swinging* and *flicking*. Four specific techniques were chosen for this study in order to represent the variety of approaches reported in related work. In the description of each technique below we describe where in the related work they were found. This is also summarized in Table 1.

Technique	Example origins in related work
Pinching	Benko and Wilson [2], Hespanhol et al. [9], Markussen et al. [20], Ikemasu et al. [14]
Swiping	Bragdon et. al [5], Seifert et al. [34]
Swinging	Walter et al. [38], Yatani et al. [40], Scheible et al. [33], Dachsel et al. [6]
Flicking	Lucero et al. [18]

Table 1. The four techniques and origins in related work.

In order to investigate the relative use of these approaches, we have implemented each of them as a specific interaction technique capable of facilitating bi-directional transfer of data. As such, the implemented techniques are neither new nor a contribution in themselves. They merely represent a class of existing techniques, for the purpose of the study.

Pinching

The pinching approach is based on midair gestures of “picking up” an icon on either the handheld device or large display, and then “releasing” it on to the other. This simulates the real world action of picking up an object in one location and placing it in another. Specific techniques exploring this approach can be found in, for example, Benko and Wilson [2], Hespanhol et al. [9], Markussen et al. [20] and Ikemasu et al. [14]. In our instantiation of the pinching approach (Figure 1) data can be moved between the handheld device and the large display through a series of steps, combining touch screen input and midair gestures, and requiring the use of both hands. This is done as described below.

Pinching *to* the large display is done by the user first picking up the object of interest from the smartphone screen with his fingers and then closing his hand as if virtually holding the object (Figure 1a). The user then raises his closed hand and uses it as a pointer on the large display to indicate where he wants it to be put (Figure 1b). Finally, he opens his hand and releases the object onto the large display where he is pointing (Figure 1c). Pinching *from* the large display is similar but done in reverse. Here, the user first points to an object on the large display with his hand (Figure 1d). He then picks it up by closing his hand (Figure 1e). Finally, he places the object on his phone by touching the screen (Figure 1f).

Figure 1. Pinching to (a-c) and from (d-f) the large display.

Swiping

The swiping approach is based on a combination of a mid-air gesture for pointing at an icon or target location on the large screen, and a swiping gesture on the smartphone touch screen. This simulates the action of swiping, or sliding, an object from one location to another. Specific techniques exploring this approach can be found in, for example,

Bragdon et. al [5] and Seifert et al. [34]. In our instantiation of the swiping approach (Figure 2) data can be moved between the handheld device and the large display using one hand, as described below.

Swiping *to* the large display is done by the user first pointing with the smartphone at the desired location on the large display (Figure 2a). She then swipes the object to be moved from the smartphone away from herself toward the display (Figure 2b), after which it appears there (Figure 2c). Swiping *from* the large display is done in reverse. Here, the user first points at an object on the large display with her smartphone (Figure 2d). She then does a swipe gesture on the phone, away from the large display and toward herself (Figure 2d), after which the object appears on the phone (Figure 2c).

Figure 2. Swiping to (a-c) and from (d-f) the large display.

Swinging

The swinging approach is based on a combination of mid-air gestures of “pointing” with one hand and “swinging” with the other. This simulates the action of propelling an object in a certain direction. Specific techniques exploring this approach can be found in, for example, Walter et al. [38], Yatani et al. [40], and Scheible et al. [33]. In our instantiation of the swinging approach (Figure 3) data can be moved between the handheld device and the large display through a series of steps, requiring the use of both hands. This is done as described below.

Figure 3. Swinging to (a-c) and from (d-f) the large display.

Swinging *to* the large display is done by the user first pointing at the target position (Figure 3a), and then selecting the object to be moved on the smartphone (Figure 3b). He then performs a swinging gesture with the hand holding the phone away from himself toward the large display (Figure 3c), after which the object appears there. Swinging *from* the large display to the handheld device is similar but done in reverse. First the user points at the object on the large display (Figure 3d). He then selects the matching object on his phone (Figure 3e) and performs a swinging gesture with that hand away from the display toward himself (Figure 3d).

Flicking

The flicking approach is based on a combination of a mid-air gestures for “pointing” at an icon or target location on the large screen, and then physically tilting the smartphone. This simulates the action of flicking an object away from or towards oneself. Specific techniques exploring this approach can be found in, for example, Lucero et al. [18]. In our instantiation of the flicking approach (Figure. 4) data can be moved between the handheld device and the large display using one hand, as described below.

Flicking *to* the large display is done by the user first pointing with the smartphone at the desired location on the large display (Figure 4a). She then rapidly tilts the phone in the direction away from herself toward the large display (Figure 4b), after which it appears there (Figure 4c). Flicking *from* the large display is done in reverse. Here, the user first points at an object on the large display with her smartphone (fig. 4d). She then rapidly tilts the phone, away from the large display and toward herself (Figure 4d), after which the object appears on the phone (Figure 4c).

Figure 4. Flicking to (a-c) and from (d-f) the large display.

The four techniques all require a combination of mid-air gestures, touch, and physical movement of the phone. Two of them require two-handed interaction (pinching and swinging) whereas the other two can be done with one hand (swiping and flicking). Two require physical movement of the phone (swinging and flicking) whereas the other two require touch screen input (swiping and pinching).

Implementation

All techniques were implemented using the Microsoft Kinect V2 and the accelerometer and touch sensors on the phone. From the Kinect's depth camera, we get information about the user's location in physical space, allowing us to track the position of their hands. We are also able to determine if the user's hand is open or closed. Using the touch sensor on the phone we can recognize touch and swipe gestures, and using the accelerometer we are able to detect significant or rapid physical movement of the smartphone.

For the large screen application, we implemented a grid system where we could control the size of the grid, with each cell acting as a possible target. On top of the grid, we implemented a blue dot, looking similar to a laser pointer, which could move based on the position of the hand held closest to the large display, as tracked by the Kinect. For the mobile phone we implemented a simple screen where the user could either see, or chose, the correct shape to interact with, depending on the direction of the interaction, as described in the study section. Gesture data from the touch screen and accelerometer were collected through their APIs, and used to trigger the interaction technique being tested. Pinching and swiping were triggered by touch screen input, while swinging and flicking were triggered by input from the accelerometer. A video of the four interaction techniques can be found here: https://youtu.be/MMK3w_0Lmyk.

STUDY

We investigated the use of the four approaches through two separate studies. Study One investigated effectiveness (hit or miss) and efficiency (time taken). Study Two investigated accuracy (distance from target). Both studies used the same test application, setup, and tasks. The reason for doing two studies was, simply, that after the first study we decided to collect additional data on accuracy as well. We therefore added a second experiment focusing only on accuracy.

Test application

The four approaches were implemented in a test application designed to mimic moving data from a smartphone to a large display, and the other way around. In order to study the effect of target size on the large display, the test application was implemented so that we could shift between two different grid sizes. One had 5×10 large cells measuring 122 pixels (14.6 cm) square. The other had 10×20 small cells measuring 61 pixels (7.3 cm) on each side. On the large display, a small blue dot was used to show where the user was pointing, mimicking the beam of a laser pointer.

Moving data to the large display was enacted by presenting the user with a shape (square or circle) on the large display and two on the smartphone (Figure 5). The user then had to select the correct shape on the phone, and move it to the large display to the grid location of the matching shape using the interaction approach being tested.

Figure 5. Shapes on the handheld device and large display when moving an object *to* the large display.

Moving data from the large display was enacted by presenting the user with two shapes on the large display, and one on their smartphone (Figure 6). The user then had to select the correct shape on the large display, and move it to the smartphone using the interaction approach being tested.

Figure 6. Shapes on the handheld device and large display when moving an object *from* the large display.

When the test application was activated, target shapes (circle or square) would appear in a grid cell, such that 80% of the cell was filled. These shapes appeared in a predetermined series of locations to ensure an equal distribution of distances between targets for the complete test. As far as the user was concerned the sequence would appear random.

Setup

The experiment was conducted in a laboratory setting using a Samsung Galaxy S2 smartphone and a Panasonic 65" display for cross-device interaction (Figure 7). A Microsoft Kinect v2 was mounted below the large display (81 cm above the floor), and a mark on the floor 2 meters in front of the display (the optimal operating distance for the Kinect), indicated where the participant should stand.

Figure 7. Experimental Setup.

On the right hand side, a 42" display was used for showing instructional videos on the four techniques used.

Experimental Design

A within-subjects design was used for the two studies, with the four techniques and two different target sizes (large, small) as the independent variables.

Participants

Participants were recruited using posters placed around the University campus and through social networks. In total we recruited 84 participants.

For Study One there were 51 participants. These ranged in age from 21 to 52 (M: 27.98) and were between 1.56m and 1.98m tall (M: 1.79m). 15.7% were women, and 84.3% were male. All of the participants owned smartphones and had owned one for 2-12 years (M:5.9).

For Study Two there were 33 participants. These were between 20 and 55 years old (M: 23.18) and were between 1.56m and 2m tall (M: 1.77m). 30.3% of the participants were female, while 69.7% were male. They had all owned smartphones for between 1 and 9 years (M: 5.5).

Task & Procedure

The task and procedure were the same for Study One and Study Two, with the additional instructions in Study Two to be as precise as possible and aim for the white cross placed in the centre of the target shape. Before a participant started the test, the general purpose of the study was explained to them. A demonstration video of a technique was then shown once on the large screen, and then repeated in a loop on the smaller display on the right hand side (Figure 7). The trial then commenced, with the participant attempting to “hit” all targets using the assigned technique. After completing one technique the test application would shift to the next one, and follow the same procedure described, until the participant had used all four techniques.

For each technique there were 18 consecutive targets (9 small and 9 large), with 3 additional practice targets at the beginning, allowing participants to familiarize with the technique before data collection. The techniques and target sizes were presented to participants in a mixed order to minimize the learning effect.

For both Study One and Study Two, 51% of the participants started the test by moving data *to* the large display, and 49% started by moving it *from* the display. For Study One, 27% of the participants started the test with pinching, 21% with swiping, 25% with swinging, and 27% with flicking. Similarly, for Study Two, 22% of the participants started the test with pinching, 27% with swiping, 24% with swinging, and 27% with flicking. For Study One, a total of 7344 data points was registered (2 target sizes $\times$ 4 techniques $\times$ 2 directions $\times$ 9 repetitions $\times$ 51 participants). For Study Two, a total of 4752 data points was registered (2 target sizes $\times$ 4 techniques $\times$ 2 directions $\times$ 9 repetitions $\times$ 33 participants).

The average time for completing the study was 25 minutes. After each study the participant completed a demographic questionnaire including age, height, gender, current phone, years of smartphone use, and prior experience with systems like the Nintendo Wii or Microsoft Kinect.

Prior to analysis some data points were removed due to equipment registering false activations, or when an

interaction was characterized as an outlier using the Outlier Labeling rule [8]. For Study One, 176 attempts were removed due to system errors, and 406 due to outliers, resulting in a total of 6762 data points. For Study Two, 111 attempts were removed due to system errors, and 138 due to outliers, resulting in a total of 4503 data points. Table 2 shows the final number of attempts that were included in the analysis for the two different studies.

	Pinching		Swiping		Swinging		Flicking	
	To	From	To	From	To	From	To	From
Study 1	830	784	835	893	814	867	862	877
Study 2	551	533	576	576	561	569	582	555

Table 2. Final number of data points per technique.

FINDINGS

In following section, we report our findings. We begin with findings from Study One on hits or misses, and time taken. We then continue with findings from Study Two on distance from target.

Study One. Hits or Misses

For each technique we recorded whether the target had been successfully hit or missed. We then calculated the effectiveness of each user by summarizing the number of successful attempts per technique. The average successful attempts of a user per technique, along with standard deviations are presented in Table 3.

	Pinching	Swiping	Swinging	Flicking
To	15.60 (2.44)	15.72 (2.59)	14.88 (3.02)	14.08 (2.73)
From	14.49 (3.18)	17.08 (1.48)	16.43 (1.58)	12.35 (3.08)

Table 3. Average user's successful attempts per technique (max value=18, standard deviations in parentheses).

We examined effectiveness by performing two-way repeated measures ANOVAs both for the *to* as well as for the *from* scenario. The independent variables included in the analysis were *target size* and *techniques*. For the *to* scenario, *technique* had a significant effect on effectiveness ($F(3,150)=6.793$, $p<.001$) and the same was the case for *target size* ($F(1,50)=90.159$, $p<.001$) as well their interaction ($F(3,150)=4.715$, $p=.001$). From the pairwise comparisons we identified that the effect of *target size* was significant for all techniques (pinching, $p=.031$, swiping, $p=.002$, swinging, $p=.006$, and flicking, $p<.001$). For all of them the larger the target the more effective the users were. Users were most effective while swiping and pinching, while the worst performing technique was flicking. Swiping was significantly different from flicking ($p=.013$), and so was pinching ($p=.025$).

For the *from* scenario, Mauchly's Test of Sphericity was significant so we present the Greenhouse-Geisser corrected data. *Technique* had a significant effect on effectiveness

($F(2.466, 123.307)=28.275$, $p<.001$) and the same was the case for *target size* ($F(1, 50)=4.320$, $p=.043$) and their interaction ($F(2.441, 122.059)=83.787$, $p<.001$). The effect of target size was again significant for all techniques (for all $p<.001$). As opposed to sending data *to* the large display, when getting it *from* the large display, swiping and swinging were the most effective and flicking the least. We identified statistically significant differences between the pairs, pinching and swinging ($p<.001$), and swiping and the rest of techniques (for all $p<.001$).

Study One: Time Taken

The *time taken* to complete each attempt was recorded from the point at which users finished one data transfer until the completion of the next. The average time taken for each technique as well as standard deviations can be found in Table 4. We performed a linear mixed effects (LME) model analysis on the data to see how the different aspects of our experiment affected the time taken to hit a target. We used the linear mixed effect models because our experiment had both fixed and random effects and LME allowed us to take into consideration the variability of each user. In the model we included the users and *hit success* (hit or miss), *direction* (to or from), *target size* (small or large) and *technique*.

	Pinching	Swiping	Swinging	Flicking
To	5.05 (1.39)	3.9 (0.98)	4.68 (1.11)	5.06 (1.79)
From	5.74 (1.57)	3.67 (0.86)	4.12 (1.01)	4.55 (1.61)

Table 4. Time taken: average completion time in seconds and (standard deviations) per technique.

We found that neither *hit success* nor *direction* had an effect on the time each user took per attempt. However, *target size* ($F(1, 6695.228)=91.634$, $p<.001$) as well as *technique* ($F(1, 6695.228)=91.634$, $p<.001$) did have significant effects on time taken. We also performed a post-hoc LSD pairwise comparison and found that all techniques were significantly different ($p<.001$) from each other. Other significant interactions between the variables were also identified. *direction* $\times$ *technique* ($F(3, 6694.657)=52.272$, $p<.001$), *hit success* $\times$ *technique* ($F(3, 6696.169)=5.227$, $p<.001$) and finally *hit success* $\times$ *direction* $\times$ *technique* ($F(3, 6696.038)=10.235$, $p<.001$) all showed significant interactions. A post-hoc LSD pairwise comparison on *direction* and *technique* showed that for all techniques the difference in time between *to* and *from* were significant. We then did a post-hoc LSD pairwise comparison on *hit success* for each *technique* and *direction* to see where that significance was. The only significantly different pair was between a successful and unsuccessful attempt of the pinching *from* technique ($p<.001$) meaning that an unsuccessful pinching *from* interaction takes a significantly longer time to perform than a successful pinching *to* interaction.

Figure 8 shows the mean time taken for each technique, for both directions, in graphical form to give an overview, including standard deviation indicators.

Figure 8. Time taken – Comparison of mean and standard deviation for all approaches.

Study Two: Distance from Target

We calculated the *distance from target* in pixels, with regard to how far from the centre of the target the pointer was when the participant activated the data transfer. The mean distance in pixels and standard deviations can be found in Table 5. Again, we performed a LME model analysis on the dataset to see if each technique had a significant effect on the *distance from target* of each attempt.

	Pinching	Swiping	Swinging	Flicking
To	16.8 (10.5)	14.14 (9.32)	16.24 (10.29)	28.03 (19.7)
From	17.65 (10.92)	12.4 (8.59)	15.07 (10.25)	32.71 (21.49)

Table 5. Distance from Target: mean distance in pixels and (standard deviations).

We found that *technique* ($F(3, 4458.26)=193.869$, $p<.001$) and *target size* ($F(1, 4462.203)=100.016$, $p<.001$) had an effect on *distance from target*, but *direction* did not have an effect. We then performed a post-hoc LSD pairwise comparison to see where the differences lay, and found that all techniques were significantly different from each other ($p<.003$). We also found that the *direction* $\times$ *technique* interaction had a significant effect on the *distance from target* ($F(3, 4457.354)=8.882$, $p<.001$). We then performed another LSD pairwise comparison between *technique* for each *direction*. We found that the only pair that was not significantly different from each other was pinching *to* and swinging *to* ($p=.508$). All others were significantly different from one another ($p<.004$). Lastly, we did a LSD pairwise comparison between *direction* for each *technique* and the only two that were not significantly different from each other were pinching *to* and pinching *from* ($p=.355$). The difference between *to* and *from* interactions for all other techniques were significant ($p<.044$).

Figure 9 shows the mean *distance from target* for each technique for both *to* and *from* scenarios, in graphical form, with standard deviation indicators.

Figure 9. Distance from target – Comparison of mean and standard deviation for all approaches.

DISCUSSION

Our results, based on empirical data collected through the experiment reported, make an important contribution to the knowledge on cross-device interaction by providing new insights with respect to knowing more about how these four different approaches compare in terms of effectiveness, efficiency and accuracy. We use the term effectiveness to refer to hit success, that is, how often the user successfully hit the target. We use the term efficiency to refer to time taken to interact. Both effectiveness and efficiency were investigated in Study One. Finally, when we talk about accuracy, we refer to how precise an interaction was done measured by the distance between the position of the pointer and the centre of the target, as investigated in Study Two.

Effectiveness

Our results for effectiveness (hit or miss) show that there is in fact some association between each approach and whether or not an attempt was successful. Looking at the hit success rates for the different techniques we can see that swiping outperforms all other techniques, for both directions of interaction. The effectiveness ranking for techniques sending data *to* the large display is: (1) swiping, (2) pinching, (3) swinging, and (4) flicking. For getting data *from* the large screen display to the smartphone, the ranking is: (1) swiping, (2) swinging, (3) pinching, and (4) flicking. This means that for both directions of interaction, swiping is the most effective technique for hitting a target correctly, whereas flicking is the least effective. In looking at the characteristics of the swiping approach, this is a 1-handed technique where the user points with the handheld device, and initiates the moving of an object *to* or *from* the large display using an on-screen gesture. Looking at the characteristics of the flicking approach, this is notably *also* a 1-handed technique where both pointing and selection is done with the hand that holds the mobile device. The difference, however, is that physically flicking the handheld device to move an object notably interferes with where the user is pointing, thereby moving the pointer on the large display at the exact moment of selection. This leads us to conclude that in order to achieve effectiveness with 1-handed techniques, the gesture chosen for selection must be one that requires minimal physical movement of the pointing device. If using a 2-handed technique, such as pinching or swinging, the physical movement of the handheld device is not so important.

Efficiency

Looking at the results with regard to the efficiency (time taken) for each technique, some rather interesting things come to light. We see that swiping is the fastest approach by far. With a mean time per attempt of 3.67 seconds for swiping *from* and 3.9 seconds for swiping *to*. The next fastest approach is swinging, with a mean time of 4.12 seconds for swinging *from* and 4.68 seconds per attempt for swinging *to*. One might have expected both one-handed approaches to be the fastest, since they take less steps and physical bodily movement in use, but swinging is notably more efficient than flicking. This could be due to the observation that people felt very uncomfortable with the flicking approach. When performing it, users appeared very cautious because the pointer tended to move large distances unexpectedly during some attempts (as discussed under effectiveness). This led people to sometimes perform a very cautious and slow tilting movement with the phone, and consequently sometimes not triggering the interaction, and having to repeat the tilting movement. Not surprisingly, pinching was the slowest approach, since this required the greatest number of steps. This means that it is not the number of hands required to use a particular interaction technique that determines its efficiency. Rather it is the combination of how to point and how to select that is crucial. As with effectiveness, 1-handed approaches can be made very efficient if pointing is combined with a subtle gesture for interaction, but is very sensitive to larger physical gestures. For 2-handed techniques, these may not be as efficient, but are then less sensitive to the physical gesture used for selection.

It is also worth noting that success of the interaction attempt did not have a significant effect on the time for that attempt. This can be explained by considering that having to take aim with a pointer (hand or mobile device) is common for all techniques, and as such should take the same amount of time. Once the user has acquired the target object, s/he then initiates an interaction with this object, which is where the differences in time taken is most notable between the four different techniques. In relation to this, and as expected, target size did have significant effect on the time taken for all techniques since aiming takes a little longer for smaller targets, as also shown in other studies. However, this affected all techniques equally. This means that, in essence, the possible size of the objects on the large display depends on the precision of the pointing device, at the distance of interaction, and not so much on the technique used for selection – as long as selection does not interfere with pointing, as discussed above.

The direction (*to* or *from*) of each attempt did not have a significant effect on the efficiency of the techniques. This is because not all techniques were affected similarly by the direction. While swiping, swinging and flicking all took less time to perform when moving objects *from* the large display, pinching performed in the opposite way. This can be explained by examining the interaction between direction, technique and success in more detail. Our pairwise

comparisons show that the significant interaction between direction, technique and success comes from the difference between the successful and unsuccessful attempts for the pinching *from* technique. For this technique we saw that an unsuccessful attempt with pinching *from* is significantly slower than a successful one. This means that the most likely factor at play here is the specific implementation of the technique. Once the user closes their hand in an attempt to pinch the object from the screen, it is not possible to retry, even if the target was missed. This is counter-intuitive to real world pinching where if a person misses an object they are pinching, they will simply just pinch it again. This limitation of the interaction technique clearly confused our participants. Whenever they missed an object when using the pinching technique, they would try to have another attempt, which was not possible. This limitation to the technique was, of course, an intentional design decision for the sake of the experiment, since none of the other techniques had the opportunity of retrying an interaction attempt if the target was missed. However, we speculate that the pinching technique may have been particularly disadvantaged by the experimental setup. This observation also means that, in hindsight, we would probably not impose such artificial limitation on the different interaction techniques, but simply allow the users to retry an interaction if they failed. This way we could also collect data on the number of retries required for each technique

Accuracy

We see from our results that out of the 4 approaches, swiping was the most accurate technique, having a mean of 12.4 pixels' distance from the centre of each target for swiping *from* and 14.14 pixels for swiping *to*. This is closely followed by the swinging techniques, with 15.07 pixels for swinging *from* and 16.24 pixels for swinging *to*, and pinching *to* with 16.8 pixels and pinching *from* with 17.65 pixels. The flicking technique trail far behind, with a mean distance of 28.03 pixels for flicking *to* and 32.71 pixels from the centre of each target for flicking *from*. This means, again, that a 1-handed technique can be made very accurate when combined with a subtle selection gesture, such as swiping on a display, but will drop significantly in performance when combined with a strong physical gesture, such as flicking the phone. Again, 2-handed techniques, however, can be made almost as accurate, and appear less sensitive to which selection gesture is used. Having observed the tests, it was not surprising to us that flicking was so far behind, since it requires the user to do quite a big physical movement with the same hand that they are pointing with. As discussed earlier, this usually led to the pointer moving away from its original position during the transfer gesture, resulting in a missed target. In fact, several users even placed the pointer slightly off of the target to compensate for this movement.

Even though we asked our participants to be as precise as possible when aiming at the white cross in the centre of the target shape in Study Two, we observed that target size influenced how precise they tried to be. This is likely due to the fact that, to some degree, they regarded the shape itself

as the target. We observed that when users had the pointer as close as possible to the centre of the shape, they would then switch their attention to performing the selection gesture. While doing this, the pointer often started to drift, as it was difficult to hold it completely still for very long. Since it did not take much movement before the pointer went outside the small target's shape area, users would actively work to realign the pointer with the centre of the shape before performing the technique. This was not the case with the larger targets. Here the pointer could drift further from the centre cross before it actually went outside the shape, and as such did not prompt the users to realign while doing the selection gesture. This suggests that even though users took great care to initially aim precisely at the cross, once their focus had shifted to activating the transfer, they would go ahead with that as long as the pointer was still inside the shape. This means that it might be interesting to experiment with a mechanism that detects if the user is trying to take aim at an object, e.g. holding the pointer still, and then reduce the sensitivity of the pointer, so that a more accurate selection can be done, as also explored by Myers et al. [25] and Nancel et al. [27] when combining coarse and precision pointing.

Common Traits

We selected our set of four approaches to have a spread of different traits in terms of number of hands and number of steps required to interact. We also varied whether pointing was done using one's hand or a smartphone, and the types of physical movements required to select an object. From this we can see traits that are common between techniques, and from here compare traits of those techniques that proved to be the most effective, efficient and accurate in our studies.

Looking at the common trait of the two approaches, swiping and swinging, which consistently outperformed the others, we can see that the means of pointing (hand or phone) can be held relatively motionless by the user throughout performance of the motion required to execute the technique. Conversely, the other two approaches, flicking and pinching, require the means of pointing (hand or phone), to be moved during the selection part of the technique, causing unwanted movement of the pointer on the large display, resulting in users finding it difficult to control. As discussed earlier, these traits had significant influence on both the effectiveness, efficiency and accuracy of the interactions afforded by the individual techniques. This means that, independent of the number of hands required by a specific technique, the ability to keep the means of pointing (hand or phone) stable during the interaction plays a major role in whether or not the technique will perform well in terms of effectiveness, efficiency and accuracy.

CONCLUSIONS

One of the crucial requirements particular to cross-device or multi-display environments is the ability for the user to move objects from one device or display to another [26]. Yet, this is still surprisingly difficult [22]. In this paper we have contributed to a better understanding of what makes

particular approaches to cross-device interaction for moving objects from one device or display to another better than others in terms of effectiveness, efficiency and accuracy. We have presented results from a comparative empirical study investigating four approaches to cross-device interaction between a smartphone and a large display. Our findings are based on two studies with users, comparing the use of four different techniques, representing the approaches of *pinching*, *swiping*, *swinging* and *flicking*. Our results show that for interactions both *to* and *from* the large display, swiping and swinging were the most successful approaches out of the four investigated, meaning that the number of hands required for the technique is not crucial. We contribute this to the common factor between the two approaches being that they do not inflict movement on the pointing device (hand or phone) during the selection part of the interaction. This indicates that an important contributor to the performance of a technique, in terms of effectiveness, efficiency and accuracy, is the user's ability to stabilize the pointer on the large display while interacting. This insight is valuable for future design of interaction techniques for cross-device interaction because it narrows down the combinations of gestures for pointing and selecting quite considerably, while still encouraging both 1- and 2-handed approaches.

FUTURE WORK

The research presented in this paper leaves a number of opportunities for future work. Firstly, we have chosen to do a laboratory study of single user interaction with a non ad-hoc task as a starting point for learning about the relative strengths and weaknesses of the four different approaches identified from the literature. For possible future work, in order to study the external validity of our findings, we suggest expanding this scope to include multiple users, and realistic tasks, for example ad-hoc ones that emerge unplanned out of real use of a cross-device system. Secondly, the age of our participants ranged quite broadly from 21 to 52 years (M: 27.98) in study one, and from 20 to 55 years (M: 23.18) in study 2. As these means show a small tendency toward the younger end of the age range, it would be relevant to test our findings with other age groups, including older users. Finally, it might be fruitful to conduct a gesture elicitation study [24] in order to let users inform the gestures.

ACKNOWLEDGMENTS

We thank everyone who participated in the two studies, and Kasper Hornbæk for valuable feedback on an earlier draft. We also thank Elias Ringhauge and Ivan Penchev for their contribution to the implementation of the four techniques. The work is part of Center for Data-Intensive Cyber-Physical Systems (DiCyPS) funded by Innovation Fund Denmark.

REFERENCES

1. Michel Beaudouin-Lafon, Stephane Huot, Mathieu Nancel, Wendy Mackay, Emmanuel Pietriga, Romain Primet, Julie Wagner, Olivier Chapuis, Clement Pillias, James Eagan, Tony Gjerlufsen, Clemens Klokmoose. 2012. Multisurface interaction in the wild room. *Computer* 45, no. 4 (2012): 48-56.
2. Hrvoje Benko and Andrew D. Wilson. 2010. Multi-point Interactions with Immersive Omnidirectional Visualizations in a Dome. In *ACM International Conference on Interactive Tabletops and Surfaces (ITS'10)*. ACM, New York, NY, USA, 19-28. <http://dx.doi.org/10.1145/1936652.1936657>
3. Richard A. Bolt. 1980. "Put-that-there": Voice and gesture at the graphics interface. In *Proceedings of the 7th annual conference on Computer graphics and interactive techniques (SIGGRAPH '80)*. ACM, New York, NY, USA, 262-270. <http://dx.doi.org/10.1145/800250.807503>
4. Sebastian Boring, Marko Jurmu, and Andreas Butz. 2009. Scroll, Tilt or Move It: Using Mobile Phones to Continuously Control Pointers on Large Public Displays. In *Proceedings of the 21st Annual Conference of the Australian Computer-Human Interaction Special Interest Group: Design: Open 24/7 (OZCHI '09)*. ACM, New York, NY, USA, 161-168. <http://dx.doi.org/10.1145/1738826.1738853>
5. Andrew Bragdon, Rob DeLine, Ken Hinckley, and Meredith Ringel Morris. 2011. Code Space: Touch + Air Gesture Hybrid Interactions for Supporting Developer Meetings. In *Proceedings of the International Conference on Interactive Tabletops and Surfaces (ITS '11)*. ACM, New York, NY, USA, 212-221. <http://dx.doi.org/10.1145/2076354.2076393>
6. Raimund Dachsel and Robert Buchholz. 2009. Natural throw and tilt interaction between mobile phones and distant displays. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems (CHI EA '09)*. ACM, New York, NY, USA, 3253-3258. <http://dx.doi.org/10.1145/1520340.1520467>
7. L. Gallo and A. Minutolo. 2012. Design and Comparative Evaluation of Smoothed Pointing: A Velocity-oriented Remote Pointing Enhancement Technique. *Int. J. Hum.-Comput. Stud.* 70, 4 (April 2012), 287-300. <http://dx.doi.org/10.1016/j.ijhcs.2011.12.001>
8. Tony Gjerlufsen, Clemens Nylandsted Klokmoose, James Eagan, Clément Pillias, and Michel Beaudouin-Lafon. 2011. Shared substance: developing flexible multi-surface applications. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 3383-3392. <http://dx.doi.org/10.1145/1978942.1979446>
9. Luke Hespanhol, Martin Tomitsch, Kazjon Grace, Anthony Collins, and Judy Kay. 2012. Investigating Intuitiveness and Effectiveness of Gestures for Free Spatial Interaction with Large Displays. In *Proceedings of the 2012 International Symposium on Pervasive Displays (PerDis '12)*. ACM, New York, NY, USA, <http://dx.doi.org/10.1145/2307798.2307804>

10. Juan David Hincapié-Ramos, Xiang Guo, Paymahn Moghadasian, and Pourang Irani. 2014. Consumed endurance: a metric to quantify arm fatigue of mid-air interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '14). ACM, New York, NY, USA, 1063–1072. <http://dx.doi.org/10.1145/2556288.2557130>
11. Steven Houben and Nicolai Marquardt. 2015. WatchConnect: A Toolkit for Prototyping Smartwatch-Centric Cross-Device Applications. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, ACM Press, 1247–1256. <http://doi.org/10.1145/2702123.2702215>
12. Steven Houben, Paolo Tell, and Jakob E. Bardram. 2014. ActivitySpace: Managing Device Ecologies in an Activity-Centric Configuration Space. In *Proceedings of International Conference on Interactive Tabletops and Surfaces* (ITS '14). ACM, New York, NY, USA, 119–128. <http://dx.doi.org/10.1145/2669485.2669493>
13. David C. Hoaglin and Boris Iglewicz. 1987. Fine-Tuning Some Resistant Rules for Outlier Labeling. *J. Amer. Statist. Assoc.* 82, 400 (1987), 1147–1149. <http://www.jstor.org/stable/2289392>
14. Kaori Ikematsu and Itiro Siio. 2015. Memory Stones: An Intuitive Information Transfer Technique Between Multi-touch Computers. In *Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications* (HotMobile '15). ACM, New York, NY, USA, 3–8. <http://dx.doi.org/10.1145/2699343.2699352>
15. Mikkel R. Jakobsen, Yvonne Jansen, Sebastian Boring, and Kasper Hornbæk. 2015. Should I Stay or Should I Go? Selecting Between Touch and mid-air gestures for large-display interaction. In *Proceedings of the 15th IFIP TC 13 International Conference on Human-Computer Interaction* (INTERACT 2015), 455–473. http://dx.doi.org/10.1007/978-3-319-22698-9_31
16. Tero Jokela, Jarno Ojala, Guido Grassel, Petri Piippo, and Thomas Olsson. 2015. A Comparison of Methods to Move Visual Objects Between Personal Mobile Devices in Different Contexts of Use. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services* (MobileHCI '15). ACM, New York, NY, USA, 172–181. <http://dx.doi.org/10.1145/2785830.2785841>
17. Tero Jokela, Jarno Ojala, and Thomas Olsson. 2015. A Diary Study on Combining Multiple Information Devices in Everyday Activities and Tasks. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). ACM, New York, NY, USA, 3903–3912. <http://dx.doi.org/10.1145/2702123.2702211>
18. Andrés Lucero, Jussi Holopainen, and Tero Jokela. 2012. MobiComics: Collaborative Use of Mobile Phones and Large Displays for Public Expression. In *Proceedings of the 14th International Conference on Human-computer Interaction with Mobile Devices and Services* (MobileHCI '12). ACM, New York, NY, USA, 383–392. <http://dx.doi.org/10.1145/2371574.2371634>
19. Anders Markussen, Sebastian Boring, Mikkel R. Jakobsen, and Kasper Hornbæk. 2016. Off-Limits: Interacting Beyond the Boundaries of Large Displays. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (CHI '16). ACM, New York, NY, USA, 5862–5873. <http://dx.doi.org/10.1145/2858036.2858083>
20. Anders Markussen, Mikkel Rønne Jakobsen, and Kasper Hornbæk. 2014. Vulture: A Mid-air Word-gesture Keyboard. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems* (CHI '14). ACM, New York, NY, USA, 1073–1082. <http://dx.doi.org/10.1145/2556288.2556964>
21. Nicolai Marquardt, Ken Hinckley, and Saul Greenberg. 2012. Cross-device interaction via micro-mobility and f-formations. In *Proceedings of the 25th annual ACM symposium on User interface software and technology* (UIST '12). ACM, New York, NY, USA, 13–22. <http://dx.doi.org/10.1145/2380116.2380121>
22. Nicolai Marquardt, Till Ballendat, Sebastian Boring, Saul Greenberg, and Ken Hinckley. 2012. Gradual engagement: facilitating information exchange between digital devices as a function of proximity. In *Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces* (ITS '12). ACM, New York, NY, USA, 31–40. <http://dx.doi.org/10.1145/2396636.2396642>
23. Sven Mayer, Katrin Wolf, Stefan Schneegass, and Niels Henze. 2015. Modeling Distant Pointing for Compensating Systematic Displacements. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). ACM, New York, NY, USA, 4165–4168. <http://dx.doi.org/10.1145/2702123.2702332>
24. Meredith R. Morris, Andreea Danielescu, Steven Drucker, Danyel Fisher, Bongshin Lee, m. c. schraefel, and Jacob O. Wobbrock. 2014. Reducing legacy bias in gesture elicitation studies. *interactions* 21, 3 (May 2014), 40–45. <http://dx.doi.org/10.1145/2591689>
25. Brad A. Myers, Rishi Bhatnagar, Jeffrey Nichols, Choon Hong Peck, Dave Kong, Robert Miller, and A. Chris Long. 2002. Interacting at a Distance: Measuring the Performance of Laser Pointers and Other Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '02). ACM, New

York, NY, USA, 33–40.

<http://dx.doi.org/10.1145/503376.503383>

26. Miguel A. Nacenta, Carl Gutwin, Dzmitry Aliakseyeu, Sriram Subramanian. 2009. There and Back again: Cross-Display Object Movement in Multi-Display Environments. *Journal of Human-Computer Interaction*, vol. 24 no. 1, 170-229. <http://dx.doi.org/10.1080/07370020902819882>
27. Mathieu Nancel, Olivier Chapuis, Emmanuel Pietriga, Xing-Dong Yang, Pourang P. Irani, and Michel Beaudouin-Lafon. 2013. High-precision Pointing on Large Wall Displays Using Small Handheld Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '13). ACM, New York, NY, USA, 831–840. <http://dx.doi.org/10.1145/2470654.2470773>
28. Roman Rädle, Hans-Christian Jetter, Mario Schreiner, Zhihao Lu, Harald Reiterer, and Yvonne Rogers. 2015. Spatially-aware or Spatially-agnostic?: Elicitation and Evaluation of User-Defined Cross-Device Interactions. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (CHI '15). ACM, New York, NY, USA, 3913-3922. <http://dx.doi.org/10.1145/2702123.2702287>
29. Roman Rädle, Hans-Christian Jetter, Nicolai Marquardt, Harald Reiterer, and Yvonne Rogers. 2014. HuddleLamp: Spatially-Aware Mobile Displays for Ad-hoc Around-the-Table Collaboration. In *Proceedings of the Ninth ACM International Conference on Interactive Tabletops and Surfaces* (ITS '14). ACM, New York, NY, USA, 45-54. <http://dx.doi.org/10.1145/2669485.2669500>
30. Dimitrios Raptis, Jesper Kjeldskov, and Mikael B. Skov. 2016. Continuity in Multi-Device Interaction: An Online Study. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction: Game-Changing Design* (NordiCHI '16). ACM, New York, NY, USA.
31. Umar Rashid, Jarmo Kauko, Jonna Häkkinä, and Aaron Quigley. 2011. Proximal and Distal Selection of Widgets: Designing Distributed UI for Mobile Interaction with Large Display. In *Proceedings of the 13th International Conference on Human Computer Interaction with Mobile Devices and Services* (MobileHCI '11). ACM, New York, NY, USA, 495–498. <http://dx.doi.org/10.1145/2037373.2037446>
32. Jun Rekimoto. 1997. Pick-and-drop: A Direct Manipulation Technique for Multiple Computer Environments. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology* (UIST '97). ACM, New York, NY, USA, 31–39. <http://dx.doi.org/10.1145/263407.263505>
33. Jürgen Scheible, Timo Ojala, and Paul Coulton. 2008. MobiToss: A Novel Gesture Based Interface for Creating and Sharing Mobile Multimedia Art on Large Public Displays. In *Proceedings of the 16th ACM International Conference on Multimedia* (MM '08). ACM, New York, NY, USA, 957–960. <http://dx.doi.org/10.1145/1459359.1459532>
34. Julian Seifert, Andreas Bayer, and Enrico Rukzio. 2013. PointerPhone: Using Mobile Phones for Direct Pointing Interactions with Remote Displays. In *Proceedings of the 14th IFIP TC 13 International Conference on Human-Computer Interaction* (INTERACT 2013). Springer Berlin Heidelberg, 18–35. http://dx.doi.org/10.1007/978-3-642-40477-1_2
35. Mikael B. Skov, Jesper Kjeldskov, Jeni Paay, Heidi P. Jensen, and Marius P. Olsen. 2015. Investigating Cross-Device Interaction Techniques. *Proceedings of the Australian Conference on Human-Computer Interaction, OzCHI '15*, ACM Press, 446–454. <http://doi.org/10.1145/2838739.2838763>
36. Henrik Sørensen, Dimitrios Raptis, Jesper Kjeldskov, and Mikael B. Skov. 2014. The 4C framework: principles of interaction in digital ecosystems. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (UbiComp '14). ACM, New York, NY, USA, 87-97. <http://dx.doi.org/10.1145/2632048.2636089>
37. Daniel Vogel and Ravin Balakrishnan. 2005. Distant Freehand Pointing and Clicking on Very Large, High Resolution Displays. In *Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology* (UIST '05). ACM, New York, NY, USA, 33–42. <http://dx.doi.org/10.1145/1095034.1095041>
38. Robert Walter, Gilles Bailly, Nina Valkanova, and Jörg Müller. 2014. Cuenesics: using mid-air gestures to select items on interactive public displays. In *Proceedings of the 16th international conference on Human-computer interaction with mobile devices & services* (MobileHCI '14). ACM, New York, NY, USA, 299-308. DOI: <http://dx.doi.org/10.1145/2628363.2628368>
39. Paweł Woźniak, Lars Lischke, Benjamin Schmidt, Shengdong Zhao, and Morten Fjeld. 2014. Thaddeus: A Dual Device Interaction Space for Exploring Information Visualisation. *Proceedings of the 8th Nordic Conference on Human-Computer Interaction - NordiCHI '14*, ACM Press, 41–50. <http://doi.org/10.1145/2639189.2639237>
40. Koji Yatani, Koiti Tamura, Keiichi Hiroki, Masanori Sugimoto, and Hiromichi Hashizume. 2005. Toss-it: Intuitive Information Transfer Techniques for Mobile Devices. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '05). ACM, New York, NY, USA, 1881–1884. <http://dx.doi.org/10.1145/1056808.105704>