

A Comparison of Techniques for Multi-Display Reaching

Miguel A. Nacenta¹, Dzmitry Aliakseyeu², Sriram Subramanian¹, and Carl Gutwin¹

¹Department of Computer Science

University of Saskatchewan

110 Science Place, Saskatoon, SK, Canada

{nacenta, sriram, gutwin}@cs.usask.ca

²Faculty of Industrial Design

Eindhoven University of Technology

Den Dolech 2, Eindhoven, The Netherlands

d.aliakseyeu@tue.nl

ABSTRACT

Recent advances in multi-user collaboration have seen a proliferation of interaction techniques for moving digital objects from one device to another. However, little is known about how these techniques work in realistic situations, or how they compare to one another. We conducted a study to compare the efficiency of six techniques for moving objects from a tablet to a tabletop display. We compared the techniques in four different distance ranges and with three movement directions. We found that techniques like the Radar View and Pick-and-Drop, that have a control-to-display ratio of 1, are significantly faster for object movement than techniques that have smaller control-to-display ratios. We also found that using spatial manipulation of objects was faster than pressure-based manipulation.

ACM Classification: H.5.2 [Information Interfaces and Presentation]: User Interfaces - *Interaction Styles; Input devices and strategies; Theory and methods.*

Keywords: Pen-based interaction techniques; tabletop interaction; co-located collaboration; ubiquitous computing.

INTRODUCTION

In recent years, personal digital assistants (PDAs) and tablet PCs have become commonplace at meetings. Ubiquitous access to personal devices has led to situations where multiple people interact through multiple computers in the same room. In addition, rooms are being augmented with group displays such as large interactive walls or tabletop displays that can act as shared public spaces to facilitate group work [14].

An important component of these meetings is the ability to share information between devices. Some examples are giving files to another person, moving artifacts to the group display, or sending documents to a printer or meeting

recorder. In these situations, it is often necessary to move objects (files, icons, or task artifacts) between the different devices in the room. One of the goals of computer support for these situations is that it should be as easy to move virtual objects around as it is to slide papers or tools around a large table.

Recognizing the need for interaction techniques that allow transfer of objects between various devices, researchers have proposed several techniques for multi-display reaching, such as Pick-and-Drop [11], Hyperdrag [12], Throwing [5] and SyncTap [13]. However, little is known about how reaching techniques work in realistic usage situations and less still is known about how they compare to one another. To investigate these issues, we carried out two experiments that compared the performance of six techniques: Pick-and-Drop, Corresponding-Gestures, Radar Views, Pantograph, Slingshot and Press-and-Flick. We looked at the techniques over different distance ranges and different target angles.

We found that there are substantial differences between the techniques, and that there are several factors designers should consider before choosing and implementing a technique for multi-display reaching. There is clear evidence that techniques like the Radar View and Pick-and-Drop work better than techniques like Pantograph or Slingshot. We also found that pressure based techniques, though promising, are difficult to implement in a real-world scenario.

In this paper, we first describe multi-device collaboration in more detail, and then present a design framework that classifies existing reaching techniques, identifying their major characteristics and categories. We then report on the experiments we carried out, our findings, and the implications of our results for design.

BACKGROUND

Despite a growing number of digital devices, information exchange between devices is still in its infancy. There exist only a few integrated multi-device environments that go beyond basic inter-networking. One of the first attempts in mediated collaboration was PARC's CoLab Project [16] which created 'public windows' on personal workstations, and provided a large shared public space using the Liveboard display. Roomware (i-Land) [17] is a more recent and more extensive system that provides a set of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2005, April 2-7, 2005, Portland, Oregon, USA.

Copyright 2005 ACM 1-58113-998-5/05/0004...\$5.00.

artifacts (interactive electronic wall, interactive table, chairs with built-in computers) to support individual and group work. The iRoom environment [4] uses a platform-independent approach which emphasizes the ability to easily create and add new displays and input devices.

Although multi-device infrastructures are not common, several research groups have looked at supporting information exchange in such environments. Rekimoto [11] proposed the Pick-and-Drop technique, where the user can ‘pick up’ an object by touching it with a digital pen (or any other suitable device), and then ‘drop’ the object anywhere in the workspace by repeating the touch action in the desired location. A different approach is used in the SyncTap technique [13] which does not require that the devices have display surfaces. When the user wants to connect two devices she presses and releases a connection button on both devices at the same time.

The most common situation where such interaction techniques will be useful is a co-located collaborative environment like a boardroom meeting. Group work in this kind of environment is only possible if it is easy to exchange and share different types of artifacts between the participants and the different spaces. For example, consider a situation where one member of a financial team wants to share a spreadsheet with all the participants. To do so, she usually has to connect her computer to the projector, or find a shared folder to put the document into. A more natural and faster way to share the document is to simply pick the document up from the personal computer and drop it onto the shared tabletop display.

DESIGN FRAMEWORK OF REACHING TECHNIQUES

We built a design framework to classify multi-display reaching techniques based on their characteristics and requirements. The framework uses nine attributes to characterize the techniques: topology of the underlying interaction space, reaching range, nature of the destination display, feedback provided to the user, input devices used, display and input area requirements, implicit privacy rules, sidedness, and symmetry.

Topology of the environment

Topology considers how the physical space and its virtual conception relate to one another in the underlying networked environment. Different topologies enable and constrain different types of interaction techniques. While many interaction techniques rely on the existence of a particular topology, some of them work regardless of the underlying organization. We identify four main topology paradigms:

- *Virtual space*. In this topology, the interaction and transitions between the different displays are determined by a virtual space defined by the system. Each device provides a viewport into this virtual space. The Virtual meeting room [2] and virtual take techniques like those described in [3] are examples of such techniques.

- *Coupled virtual and physical space*. This combination is possible when the system keeps track of the position and/or orientation of all relevant elements in the environment (e.g. users, displays, devices). The Slingshot [7], Push-and-Throw (Pantograph), Hyperdrag [12], Drag-and-Pop [1], Radar View, Beam [5], and Press-and-Flick techniques are all based on this topology.
- *Physical space*. In order to avoid having to monitor the location of physical elements in real time, systems can simply use the physical world to interact with the virtual objects. Examples of techniques that use such a topology are Pick-and-Drop [11], Taking [5], and Tangible mediaBlocks [18].
- *Discrete* (or no topology at all). This group includes all the interaction techniques that do not rely on any spatial concept; in other words, they define their own topology. Examples of such techniques include basic shared folders and SyncTap [13].

Nature of destination space

In some multi-display reaching tasks, the user needs to send an object to a destination display but does not care about the exact location in that display. In these cases, we can consider the destination as a discrete location.

However, in the case of large or shared displays, the final location of the object could be important. For example, if people are sharing an interactive wallboard, the location of a particular object can convey meaning (e.g., important objects are centered; less important objects occupy the periphery). Here it is not sufficient to treat the destination as a discrete location. The user needs continuous access to all locations within the destination display.

While some interaction techniques are well-suited for moving objects to a particular position in the same or another display (continuous destinations), others like *wormholes* [19] are better suited to discrete destinations.

Range

All reaching techniques have different useful reach range. For example, *Pick-and-Drop* works well when the target is within arm range, but with more distant targets, this technique requires that the user stand up and walk to reach the display. This is true for all techniques based on physical space, since in each of these techniques, some physical object must actually travel from the starting point to the target. Similarly, *coupled-space* based techniques are limited by the availability of remote feedback, which helps the user to know if the object is actually going to the intended destination.

Based on the Gibsonian perception-action coupling paradigm, Stappers et al., [15] partitioned interaction space in virtual environments into three ranges: short, medium, and long. At short range, users can perform detailed manipulations using fingers and hands, and can receive feedback through detailed vision and haptics. At medium range they use their hands, arms and posture to grasp

objects and make gestures. Objects at long range are out of reach and are used mainly to provide an overview and a feeling of presence in the environment.

We adapt this partition to a meeting room scenario and identify three ranges of interaction: *within hand's reach*, *beyond hand's reach*, and *discrete points in the periphery*. This classification comes from a meeting room setting where group members work around a large table. Each person's personal devices may act as discrete points in the periphery of the interaction space, and the large table is a shared public display with detailed control available within hand's reach.

Accuracy and resolution

Techniques that allow the user to continuously position an object in a display can be categorized based on targeting accuracy. The accuracy of the continuous techniques is influenced by the resolution of the destination space. Similarly, techniques that use discrete destinations can be categorized by how accurately a user can acquire particular destinations.

The accuracy of a particular technique also depends on the other attributes described in this section (such as the type of feedback given, and the type of input device used).

Feedback

The task of repositioning an object can be defined as a control task, and feedback is fundamental in control tasks as it allows the user to correct the action being performed [8]. The quality of feedback and the user's dependence on feedback are likely to affect both efficiency and accuracy.

The way that feedback is given is one of the defining factors of a technique, and can differentiate between different implementations of the same technique. For example, while the *Radar View* presents the feedback in the sender's display, *Pantograph* uses the destination display (plus any other display en route) to provide feedback.

Input device

Another critical component of control tasks is the input device used. In the boardroom scenario and for our experimental task, the available graphical input devices will be pen input devices, touch screens, mice, and other pointing devices. If the particular technique relies on a *coupled virtual and physical space* topology, the position of the different devices tracked by the environment will be also a source of input.

Techniques relying on *physical space topologies* might not work if the input device is not recognized by all possible destination devices (e.g., *Pick-and-Drop* with pen input will not work if the pen technology is not the same in all devices, or if pens from different users cannot be distinguished.)

Display and input area requirements

Since the environment for multi-device reaching can be composed of many heterogeneous displays, it is impractical to assume a fixed or minimum size for the display or for the input. Some techniques (e.g., *Radar views* or *wormholes*) can be seriously impaired if the available area for system feedback or user input is not large enough. In addition, we cannot assume that the display will only be used to perform the reaching task as many other activities will be going on; therefore, display and input space must be shared with other applications.

Implicit privacy rules

Social aspects that surround the use of a particular interaction technique are important for user acceptability. Privacy is particularly relevant when several people are working together in a room using personal devices. Shen and colleagues [14] identify three kinds of space: *private*, where data are not visible or accessible to others; *personal*, where data may be visible but not electronically accessible by others; and *public*, where others have visibility and access to data. Greenberg and colleagues [6] further distinguish two kinds of artifacts. Personal artifacts are things created, manipulated, and owned by one and only one person. Public artifacts differ, as they are created by cooperating group members, are considered to be owned by the group rather than any individual member, and can be viewed and manipulated by all.

Some reaching techniques, by the very nature of their interaction, either intrude upon a user's privacy (e.g., *Pick-and-Drop* and most techniques with physical or virtual-physical topology) or protect it (e.g., *wormholes*). Other techniques based only on a *virtual space topology* have to implement an external mechanism to enforce privacy (e.g., restrict movement of avatars in certain parts of the space).

One-sided vs. two-sided techniques

For some techniques like the *Corresponding-Gestures*, actions are required from both the sender and the recipient. Although two-sided techniques like these could have some advantages, such as the implicit enforcement of privacy, they cannot be used to send objects to shared public displays where there is no recipient.

There are also other destinations (e.g., printers, projectors or meeting recorders) that cannot support two-sided techniques. In these cases it is better to select one-sided techniques for the design of the interactive environment.

Symmetric vs. asymmetric techniques

Symmetry considers whether reaching techniques can be used in two directions. Some techniques are suitable only for moving objects from one kind of display to another. For example, with the *Beam* techniques it is easy to point a PDA towards a shared wall, but not the other way around. In many cases, however, it is possible to construct inverse versions of asymmetric techniques.

As a first step in evaluating the interaction between the various categories of this framework, we carried out an experiment that compared the performance of six different techniques for multi-display reaching, and studied the influence of topology, range, and feedback on efficiency.

STUDY: A COMPARISON OF REACHING TECHNIQUES

In terms of the framework described above, we considered interaction techniques that used physical-space or coupled-space topology with different kinds of feedback for different ranges. Only one-sided symmetric techniques were considered. The input device used was a pressure-sensitive pen (although pressure was not used for all techniques). Privacy rules were not considered, and the display and input areas were held constant.

Due to practical limitations in the number of techniques that can be compared, we picked those that are most relevant for the boardroom scenario. The six techniques that we compared are: *Pantograph*, *Slingshot*, *Pick-and-Drop*, *Corresponding-Gestures*, *Radar view*, and *Press-and-Flick*.

Based on a pilot study, all techniques with partial or no feedback were excluded from the design, as they seemed to be much less effective than techniques with more feedback. Although techniques that are imprecise may be faster and adequate for some kinds of tasks, we decided to focus on techniques that have higher accuracy.

We conducted two experiments, one to compare the techniques for continuous ranges that are *within hand's reach* and another for continuous ranges that are *beyond hand's reach*.

All six techniques were tested in the within hand's range. However, in the beyond hand's reach range, *Pick-and-Drop* and *Corresponding-Gestures* were not included. As described in our framework, these techniques are not well suited to targets that lie out of hand's reach.

Experimental Setup for both Experiments

Apparatus

A multi-display system was set up with two computers: a pen-activated tablet PC with a 15.5cm by 21cm screen (768x1024 pixels) and a desktop PC controlling two projectors capable of projecting an image of 121 by 81cm (1536x1024 pixels) combined. The projectors were mounted vertically to project onto a large white table (151x178cm) that incorporated a pen-driven digitizer also connected to the desktop PC. Both the digitizer and the tablet PC were pressure sensitive. Figure 1 depicts the experimental setup. In our experiments, the tablet PC was fixed into position so that the software could always provide a coherent visual output. In the *within hand's reach* experiment, the tablet PC was attached to the table on the side closest to the projected surface (A in Figure 1) and in the *beyond hand's reach* experiment to the opposite side (B in Figure 1). In the second experiment the user had visual feedback only in the projected surface area at the far end of

the table. The two PCs ran separate programs that communicated with each other through TCP sockets over an ad-hoc wireless network.

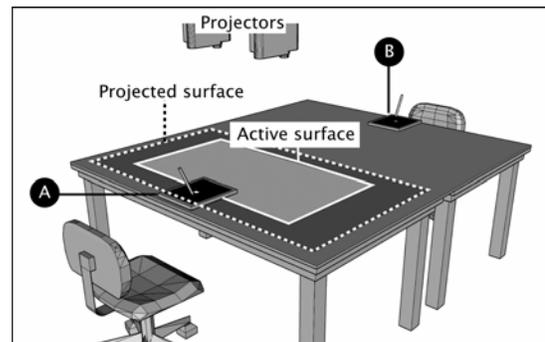


Figure 1: The experimental setup. In the within hand's reach experiment the tablet PC was placed at A and in the next experiment it was placed at B.

Task

Each subject was instructed to move several objects from the tablet PC to various targets on the table, using the different interaction techniques. The object to be moved was a 15x15mm graphical icon, and the targets were 80x80mm squares placed in the projected space of the table.

The trial was considered a 'hit' if the object ended up completely inside the area of the target, and a 'miss' otherwise. The system provided distinctive sound feedback for hits and misses. The subjects were asked to perform repetitions of this task using the different interaction techniques, as fast as possible but without missing the target.

Reaching Techniques Used

Pick-and-Drop: The Pick-and-Drop technique was implemented as defined by Rekimoto [11]. The user can select an object by touching it with a digital pen and then lifting the pen. The selected object then disappears and is visible only when the pen is moved close to the surface. As the user touches the surface again the object is dropped at that location.

Corresponding-Gestures: Corresponding-Gestures is similar to Pick-and-Drop, but instead of touching the surface to select and deselect an object, the user makes a special predefined gesture. To select the object the user draws the gesture starting inside the object (see Figure 2) and then lifts the pen. When the object is selected, it is visible only when the pen is moved close to the surface. To deselect (drop) the object the user draws the same gesture, and the object is then dropped so that its centre coincides with the starting point of the gesture (see Figure 2).

Slingshot: This technique is based on the metaphor of a physical slingshot and is similar to the drag-and-throw technique [7]. It was implemented as follows: the pen touches the object, then without losing contact with the surface the pen is moved backward and released. The object

is ‘thrown’ forward proportionally to the size of pen’s backward stroke. The dependence is linear. In order to keep the effective distance moved by the pen the same for the Pantograph, Slingshot and Radar, the amplification factors were equal across these techniques and different in both the experiments. In the within hand’s reach experiment, the amplification coefficient was equal to 15, in the beyond hand’s reach it was 40.

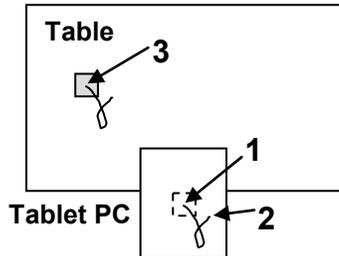


Figure 2: Corresponding-Gestures. 1-Starting point of the selecting gesture, 2 – end point of the selecting gesture, 3 – dropping gesture.

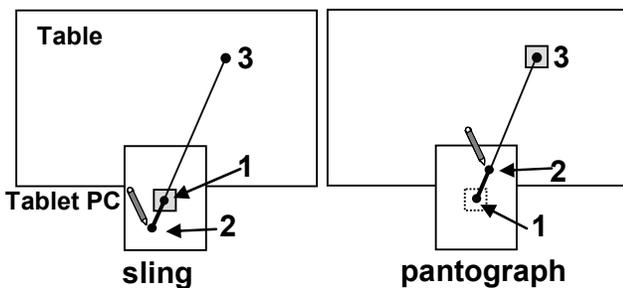


Figure 3: Slingshot and Pantograph. 1-initial position of the pen, 2-current position of the pen, 3-destination of the object.

While moving the pen backwards the user also can move it to the right and left to define the direction of the shot. The continuous visual feedback shows a line indicating where the object will end up if the pen is released (see Figure 3). The figure corresponds to the within hand’s reach experiment; in the beyond hand’s reach experiment, there was a feedback-blind area between the tablet and the projected surface (see Figure 1). This holds also for the Pantograph and Press-and-Flick techniques.

Pantograph: The implementation of the Pantograph technique is similar to the push-and-throw technique [7]. Here, the short movement of the pen is mapped into the long movement of the object. Therefore this technique is very similar to the Slingshot, but instead of a backward movement it uses forward movement. In the within hand’s reach experiment the amplification coefficient was equal to 15, in the beyond hand’s reach, 40.

The user selects an object by touching it, then without lifting the pen drags in the desired direction. As with the Slingshot, a feedback line shows the destination. The further the object is moved from the initial position, the further it will be thrown in the same direction (Figure 3).

Press-and-Flick: In this technique the throwing distance for an object is defined by pen pressure, and the direction is

defined by a single straight stroke toward the desired target. These actions are separate. The user first defines the distance (by pressing the pen against the surface) and then moves the pen slightly from the starting point to fix the throwing distance. This is signaled to the user by changing the color of the circle. At this point, only the direction will now be measured (see Figure 5). The dependence between the pressure and the throwing distance was a monotonic function adjusted using pilot testing and chosen based on the guidelines suggested in [8].

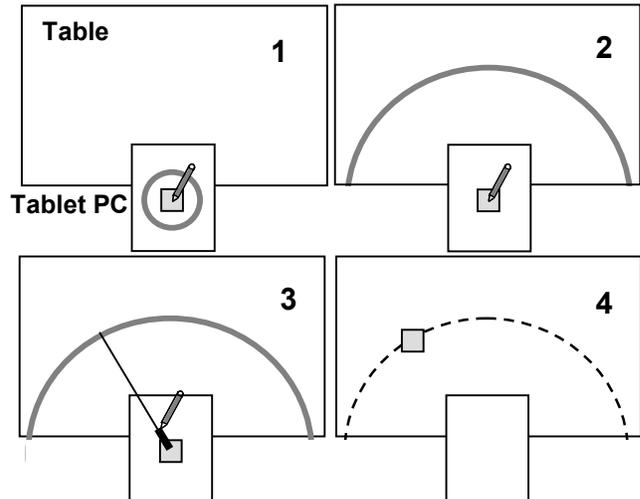


Figure 4: Press-and-Flick. 1-Pen touches the surface (low pressure); 2-user increases the pressure; 3-user starts to move the pen - the distance is automatically fixed, the color of the circle changes and the line which indicates where the object will end up is displayed; 4-pen is released.

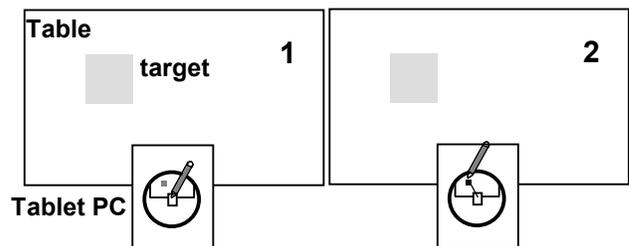


Figure 5: Radar. 1-Pen touches the object, reduced representation (map) of the surrounding environment appears, 2- user moves the pen to the representation of the target within the map and lifts the pen.

Radar View: The Radar technique uses a miniature representation (a map) of the surrounding environment (i.e., the tabletop display and the tablet PC). When the pen touches the object the map appears. The map is placed so that the position of the pen is the same in both representations (see Figure 6). The map provided is similar to that in Drag-and-Pick [1] but allows continuous positioning without distorting the shared space. When the user starts to move the pen (without lifting it from the tablet) a small line connects the starting point of the stroke with the actual position of the pen’s tip. The place on the

map where the pen is released determines the final position of the object (see Figure 5).

Experiment 1: Within hand's reach

In this experiment the different interaction techniques were compared for targets that were within reach of the user's hand.

The experiment was conducted with 10 right-handed participants (3 females and 7 males) between the ages of 18 and 44. All subjects had previous experience with graphical interfaces. No subject participated in more than one experiment and all were tested individually. The experiment used a 6x2x3 within-participants factorial design with a variety of planned comparisons. The factors were:

- Technique (Radar, Pantograph, Pick-and-Drop, Corresponding-Gestures, Press-and-Flick, Slingshot)
- Target distance (25cm, 50cm)
- Target angle (-45°, 0°, +45°)

For each technique and each location, participants completed two training trials and eight test trials, for a total of 72 training trials and 288 test trials.

At the end of the experiment participants were also asked in an exit questionnaire to rank the different techniques in order of preference. Trial completion times were used as the main measure to compare the different techniques, although number of hits and misses was also recorded.

Results

Some trials were deleted when participants accidentally released the pen from the object, causing a series of errors. In techniques like Pick-and-Drop, and Corresponding-Gestures, the users sometimes repeated their gesture to re-pick the object causing it to be dropped close to the initial point. In the case of Press-and-Flick in an attempt to change the pressure the user sometimes removed the pen from the surface causing the object to be dropped close to the initial point. These errors were due to the implementation rather than the design of the techniques, and so were removed from the data. Approximately 15 trials per subject were deleted. The distribution of deleted trials over various interaction techniques is shown in Table 1.

One-way repeated-measures ANOVA tests showed that the interaction technique had a significant effect on the trial completion time for both the near targets ($F_{5,13}=294$, $p<0.001$) and far targets ($F_{5,13}=712$, $p<0.001$). Figures 6 and 7 show the mean trial completion times with standard error for all interaction techniques in this experiment.

Post-hoc pair-wise comparisons (Tamhane, unequal variances) of interaction techniques yielded significant differences (all $p<0.02$) in trial completion times for all pairs except for the pair *Slingshot* / *Corresponding-Gestures* in the near targets and *Pantograph* / *Corresponding-Gestures* in the far targets. The techniques can be ranked in decreasing order of performance as shown in Figure 8.

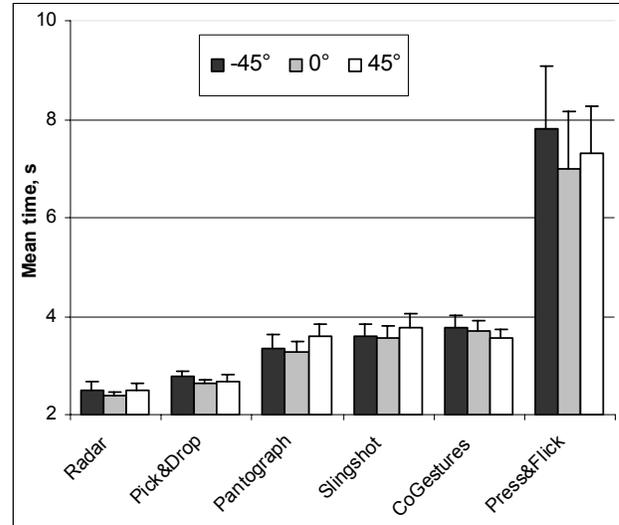


Figure 6: Mean trial completion times for different interaction techniques of experiment 1 (25 cm).

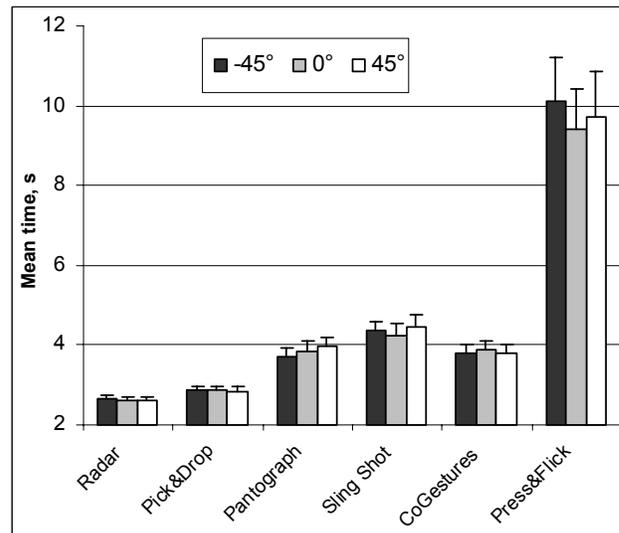


Figure 7: Mean trial completion times for different interaction techniques of experiment 1 (50 cm).

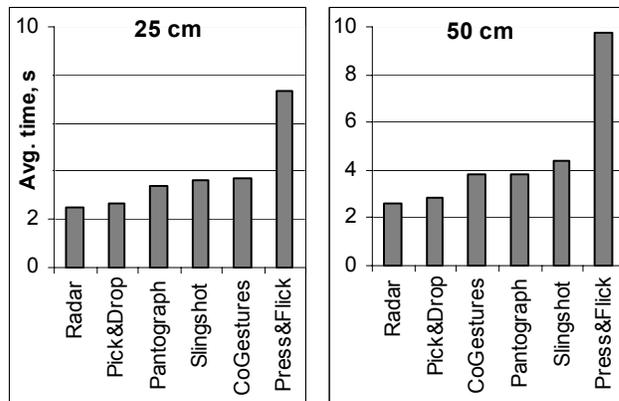


Figure 8: Efficiency ranking of the interaction techniques.

We further asked people to rank the techniques in terms of which technique the participant felt was fastest, and which they preferred overall. Results for these two questions were very similar: most people ranked the radar highest (7 of 10 people), with a few preferring either Pick-and-Drop (2 people) or Corresponding-Gestures (1 person); in addition, all participants ranked Press-and-Flick last. Ordered by average rank, the techniques are Radar, Pick-and-Drop, Corresponding-Gestures, Pantograph, Slingshot, and Press-and-Flick. These rankings matched the efficiency ranking of Figure 8.

The number of misses was significantly different between the different interaction techniques in both the 25 cm range ($\chi^2_{5, 13}=16.784, p<0.001$) and the 50cm range ($\chi^2_{5, 13}=35.895, p<0.001$). In both the 25 cm and 50 cm ranges, Radar and Pick-and-Drop had significantly lower error than either Press-and-Flick or Corresponding-Gestures.

Experiment 2: Beyond hand's reach

In this experiment, the interaction techniques were compared for targets beyond the reach of the user's hand. The experimental setup, the user task, and the interaction techniques were identical to the first experiment. However, *Pick-and-Drop* and *Corresponding-Gestures* were not included, since they would have required users to stand up and walk to the target location.

This study was conducted with 8 right-handed participants (3 females and 5 males) between the ages of 18 and 44. All participants had previous experience with graphical interfaces. The experiment used a 4x2x3 within-participants factorial design with planned comparisons. Factors were:

- Technique (Radar, Pantograph, Press-and-Flick, Slingshot)
- Target distance (80cm and 140cm)
- Target locations (-50cm, 0 cm, 50cm)

The target locations were chosen such that they are close to the target angles of the previous experiment and still within the projected space. For each location and technique, participants completed two training trials and ten test trials (for a total of 240 test trials and 48 training trials).

Results

Again, about 15 trials per participant were discarded due to inadvertent errors in releasing the pen.

One-way repeated-measures ANOVA tests showed a significant main effect of interaction technique on trial completion time for targets in both the 80 cm range ($F_{3,9}=248, p<0.001$) and the 140 cm range ($F_{3,9}=568, p<0.001$). Figure 9 shows the mean trial completion times with standard error for all interaction techniques in this experiment. Post-hoc pair-wise tests were all significant except for the difference between Pantograph and Slingshot; the pairwise tests yield the ranking reflected in Figure 10.

We also asked for people's preferences in the second study, and again perception of speed and preference were very similar. The Radar was again ranked first most frequently (7 of 8 people), and Press-and-Flick was ranked last by all participants. In order of mean rank, the techniques are: Radar, Pantograph, Slingshot, and Press-and-Flick.

The overall success rate was lower for this experiment: 1690 hits and 164 misses overall (91.2%). The number of misses were significantly different between different interaction techniques in both the 80 cm ($\chi^2_{3,9}=23.579, p<0.001$) and the 140cm ranges ($\chi^2_{3,9}=30.882, p<0.001$). Pairwise comparisons showed that Radar had significantly more hits than either Press-and-Flick or Slingshot for the 140cm range, and that Pantograph had significantly more hits than Press-and-Flick.

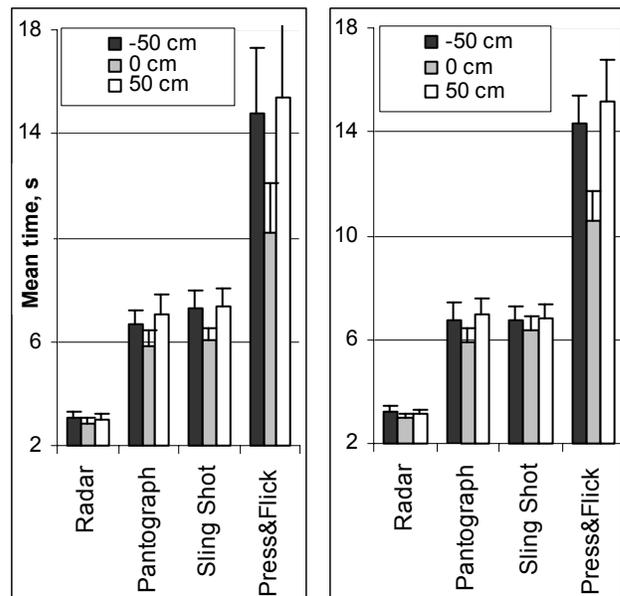


Figure 9: Mean trial completion times for different interaction techniques in Experiment 2, 80 (left) and 140 (right) cm.

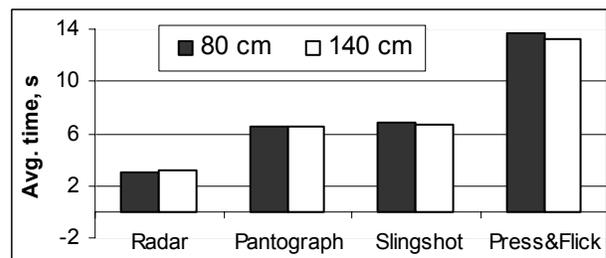


Figure 10: Efficiency ranking of interaction techniques for beyond hand's reach experiment.

DISCUSSION

Figure 11 summarizes the effect of target distance on trial completion times for the different interaction techniques of both experiments. Figure 12 and Table 1 show the target positions used in both the experiments and the corresponding miss and hit rates for each interaction technique. Techniques that were faster were also more

accurate. Radar was the most accurate interaction technique while Press-and-Flick was the least accurate. Below we discuss the performance of several techniques in particular

Radar View

The Radar View was significantly faster than all other techniques in the ranges tested, even though the actual physical movement of the pen was the same for Pantograph and Slingshot. Reaching with the Radar took about 3/4 and 1/2 the completion time of Pantograph and Slingshot in the two experiments. The Radar was also the most preferred technique: 14 out of the 18 subjects preferred it to the others.

This suggests that having visual feedback close to the subject and coupled to the input (as in the Radar View) is better than having it in the environment, even though users have to work with a duplicate representation of the target instead of the real thing. Furthermore, Radar views have a control to display ratio of 1 while other techniques like Pantograph and Slingshot had smaller control to display ratios.

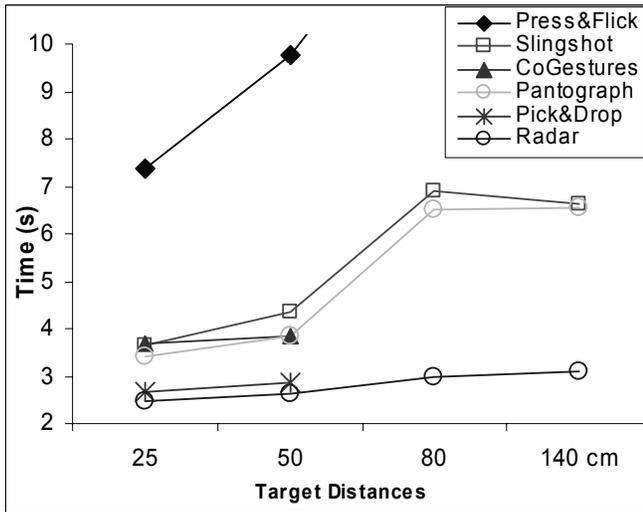


Figure 11: Effect of target distance on completion times for different interaction techniques.

However, in real-world tasks the Radar’s performance might not be as exceptional as seen here. The experimental implementation was built under the assumption that the user would be able to easily find the right destination in the radar’s map. In a real scenario, finding the correct destination could be difficult to do if the map lacks proper reference points or if it is too small for the user to easily perceive the analogy between the map and the real space.

Pick-and-Drop

Pick-and-Drop was only tested in the within hand’s reach range. It is slower than Radar but significantly faster than Slingshot or Pantograph.

Pick-and-Drop is a technique to consider when the interactions are within hand’s reach, especially if other constraints apply that would reduce Radar’s efficiency (e.g.,

very small input/output space), or if the system does not allow tracking of position and orientation. Furthermore, Pick-and-Drop’s performance is not likely to change when used in a real world setting. However, as pointed out in the framework, Pick-and-Drop can be intrusive on the recipient’s personal space.

When we compare the performance of Radar, Slingshot and Pantograph (which are based on a coupled topology) with Pick-and-Drop (based on a physical topology) we can conclude that factors like feedback and input device are more influential on performance than the different topologies. However, we did not test all topologies.

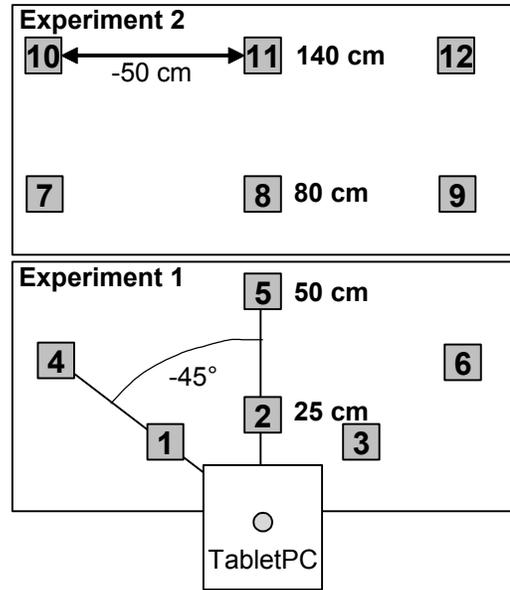


Figure 12: Target locations in both experiments.

	Radar	Pantograph	Slingshot	Press-and-Flick	Pick-and-Drop	Corresponding-Gestures
1	0.0%	1.3%	1.3% (1)	4.1% (3)	0.0% (4)	5.5% (3)
2	0.0%	2.6%	2.6%	5.4% (2)	1.3% (3)	6.9% (3)
3	0.0%	2.6%	6.7%	2.6% (1)	1.3% (4)	1.3% (2)
4	1.3% (1)	1.3%	3.9%	16.7% (3)	0.0% (5)	4.1% (4)
5	0.0%	5.3%	3.9%	11.1%	1.3% (1)	2.7% (5)
6	1.3%	3.9% (1)	3.9%	4.0% (2)	0.0% (2)	2.7% (3)
7	6.7%	1.3% (1)	1.4% (7)	1.4% (8)	-	-
8	1.3%	9.7% (1)	9.6%	15.6% (6)	-	-
9	0.0% (2)	5.6% (4)	17.9% (1)	1.4% (7)	-	-
10	6.7%	12.9% (1)	15.9%	37% (10)	-	-
11	8.1%	9.6%	17.6%	20.6% (4)	-	-
12	2.6%	9.6%	9.6%	27.6% (6)	-	-

Table 1: Error rate for all interaction techniques in both experiments corresponding to target locations in Figure 12. The numbers in parenthesis indicate the trials that were discarded due to inadvertent errors in releasing the pen.

Pantograph vs. Slingshot

Even though the Pantograph was significantly faster than Slingshot in the first experiment, the difference in means was only about 0.2 seconds. This marginal difference is partly due to the different feedback and the different device mappings used in the two techniques.

In the Pantograph, the user saw the object move in the tablet PC as they moved the pen, whereas in the Slingshot the user only saw a line. Secondly, in the Slingshot the object's motion was opposite to the pen's movement. This negative mapping, in spite of a good metaphorical basis, may not be as intuitive as the Pantograph.

The general strategy adopted by the users was to make an initial estimate of how far the pen should be dragged and then use the projected line on the table to adjust the target destination. The users' initial estimates were always in front of the intended target (i.e., they undershot the target).

However, in experiment 2, when the target appeared 80cm from the object, participants did not have sufficient feedback to use this strategy. The experimental setup was such that user only had visual feedback from the farther half of the table (see Figure 1). This meant that the users had to over-estimate the initial position of the object and then correct. Participants found this strategy harder to use than the previous one.

Corresponding-Gestures

As expected, the Corresponding-Gestures technique was slower than Pick-and-Drop. This difference can be attributed to the time for drawing the gesture at both the object and the target locations. The marginal difference in performance between Corresponding-Gestures and Slingshot or Pantograph might be partly due to different control-to-display ratios.

Unlike Pick-and-Drop, however, Corresponding-Gestures can be implemented as a two-sided technique and used even if the two devices use different input technologies. For example, to use Pick-and-Drop one needs the same pen to work on multiple displays, whereas Corresponding-Gestures can be used with a pen on one display and a mouse on the other. When Corresponding-Gestures is used as a two-sided technique, the time between the 'acquire' gesture and 'release' gesture can be shortened, further reducing the difference with Pick-and-Drop.

Press-and-Flick

Press-and-Flick was consistently and by far the least efficient technique. The two-phase interaction, the unintuitive relationship between pressure and distance and the poor control of pressure with the pen used in the experiment are possible reasons for the inferior performance. Pressure based input devices need to be studied carefully to make them easier to control. These observations complement the findings of Ramos and colleagues [10].

Applicability of the experimental setup

In our experiments the users were required to perform the task as fast as possible and with strict accuracy requirements (i.e., pre-determined target sizes and specific well-identified locations). In a real world meeting room, performance requirements are often not so strict. It is not always the aim of the user to send an object as fast or as accurately as possible. The experimental task might have also influenced users' preferences. Users will probably prefer techniques that strike a balance between convenience, speed, and accuracy.

Lessons for Practitioners

There are several lessons that designers can take from these experiments:

- In a spatially-aware environment, designers should consider Radar Views for multi-display reaching;
- If the environment cannot be made spatially aware, then Pick-and-Drop is probably the best option, as long as targets are within hand's reach;
- Providing feedback locally (as in the Radar views or Pick-and-Drop) appears to be more effective than providing distant feedback (as in Pantograph, Slingshot and Press-and-Flick);
- A control to display ratio of 1 (as in Radar Views and Pick-and-Drop) appears to be better than providing acceleration by decreasing the ratio (as in Pantograph and Slingshot);
- Spatial manipulations should be used over pressure manipulation;
- Discontinuities in feedback, even if these are at non-critical locations, should be avoided.

FUTURE WORK

In this study we only compared techniques for reaching to continuous spaces. In the next study we will include reaching to discrete locations. We also intend to look at user actions and preferences when speed and accuracy requirements are relaxed. Even though this study is focused on multi-user collaboration, the experimental task was carried out with a single user. We plan to extend our study to include multi-sided object reaching which will allow us to explore privacy and cooperation issues.

CONCLUSION

Multi-device multi-person collaborative computing situations are becoming more common, and a frequent task in these ad-hoc environments is that of moving objects from one display to another. Although several techniques have been proposed for multi-display reaching, little is known about how these techniques compare. We carried out a study of the performance of several different techniques in several task situations. We found that Radar Views and Pick-and-Drop work better than techniques like Pantograph or Slingshot. We also found that users had difficulty controlling pressure based input devices. The results of the experiment suggest that providing local feedback can be

more effective than providing feedback at the distant target. Finally, depending on the kind of feedback provided, it might be better to maintain the control to display ratio at 1 than to reduce it to provide acceleration. These results and the design framework that we have introduced provide the first empirical and analytical evidence to assist designers in determining which interaction techniques to use in multi-system co-located environments.

ACKNOWLEDGMENTS

This research was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC), and Department of Industrial Design, Eindhoven University of Technology, The Netherlands. We thank the reviewers for insightful comments and suggestions.

REFERENCES

1. Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P. Bederson, B., and Zierlinger, A. Drag-and-Pop and Drag-and-Pick: Techniques for Accessing Remote Screen Content on Touch- and Pen-operated Systems. In *Proc. of Interact 2003*, August 2003, 57-64.
2. Benford, S.D, and Fahlen, L.E. A Spatial Model of Interaction in Large Virtual Environments. In *Proc. of the Third European Conference on CSCW (ECSCW'93)*, Kluwer (1993).
3. Bowman, D.A., Johnson, D.B., Hodges, L.F. Testbed Evaluation of Virtual Environment Interaction Techniques. Presence, Vol. 10, No 1 February 2001, 75-95.
4. Fox, A., Johanson, B., Hanrahan, P., Winograd, T. Integrating Information Appliances into an Interactive Workspace. *IEEE Computer Graphics and Applications*, May/June 2000, 54-65.
5. Geißler, J. Shuffle, throw or take it! Working Efficiently with an Interactive Wall. In *Proc. of CHI'98*, ACM Press (1998), 265-266.
6. Greenberg, S., Boyle, M. and LaBerge, J. PDAs and Shared Public Displays: Making Personal Information Public, and Public Information Personal. *Personal Technologies*, Vol.3, No.1, March 1999, Elsevier.
7. Hascoët, M. Throwing models for large displays. In *Proc. of HCI'2003*, British HCI Group, 73-77.
8. Norman, D. *The Design of Everyday things*, Currency Doubleday publishers (1990), 197-199.
9. Prante T., Magerkurth, K., Streitz, N. Developing CSCW tools for idea finding: empirical results and implications for design. In *Proc. of the CSCW'02*, ACM Press (2002), 106-115.
10. Ramos, G., Boulos, M., Balakrishnan, R. Pressure Widgets. In *Proc. of CHI'04*, ACM Press (2004), 487-494.
11. Rekimoto, J. Pick-and-Drop A Direct Manipulation Technique for Multiple Computer Environments. In *Proceedings of UIST 1997*, ACM Press (1997), 31-39.
12. Rekimoto, J. and Saitoh, M. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *Proc. of CHI '99*, ACM Press (1999), 378-385.
13. Rekimoto, J., Ayatsuka, Y., Kohno, M. SyncTap: An Interaction Technique for Mobile Networking. In *Proc. of Mobile HCI 2003*, September 2003.
14. Shen, C., Everitt, K., and Ryall, K. UbiTable: Impromptu Face-to-Face Collaboration on Horizontal Interactive Surfaces. In *Proc. of UbiComp 2003*, 218-288.
15. Stappers, P.J., Adriaanse J., Keller, A.I., User-centered Conceptual Design with and for Virtual Reality Environments”, in “User Centered Design and Implementation of Virtual Environments” special issue of the *International Journal of Human-Computer Studies*, 2000.
16. Stefik, M. J., Foster, G., Bobrow, D.G., Kahn, K., Lanning, S., and Suchman, L. Beyond the chalkboard: Computer Support for Collaboration and Problem Solving in Meetings. *Communications of the ACM*, 30:1, January 1987, ACM Press (1987), pp. 32-47.
17. Streitz, N.A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., Steinmetz, R., Steinmetz, R., and Steinmetz, R. i-LAND: An interactive Landscape for Creativity and Innovation. In *Proc. of CHI '99*, ACM Press (1999), 120-127.
18. Ullmer, B., Ishii, H., Glas, D. mediaBlocks: physical containers, transports, and controls for online media. In *Proc. of the 25th annual conference on Computer graphics and interactive techniques*, ACM Press (1998), 379-386.
19. Wu, M., Balakrishnan, R. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. In *Proc. of UIST 2003 – the ACM Symposium on User Interface Software and Technology*, ACM Press (2003), 193-202.