

Multi-Display Composition: Supporting Display Sharing for Collocated Mobile Devices

Anonymous

Anonymous

Abstract. Multi-display composition is a technique that enables several mobile devices to join together over a wireless network to create a larger logical display. This logical display can be created in an *ad hoc* manner for use when and where it is needed out of a group of users' existing mobile computers. In this work we present our multi-display composition system and discuss our implementation. Furthermore, we present findings from a study of collocated groups of individuals using multi-display composition on two different types of mobile computers. In particular, we found mixed results with respect to the effect of the resulting area of the composed display. The use of two devices by a pair of participants tended to be rated more favorably than a tiled array of four devices used by a corresponding group of four participants. While the system provides additional screen real estate for smaller UMPCs, tablets were rated more favorably when using our system. Finally, we discuss usage themes that emerged from participants' use of multi-display composition.

1 Introduction

In recent years, personal computing has evolved from a primarily desktop activity to a highly mobile one. Laptop computers are an extremely popular computing platform, and the tremendous success of mobile phones indicates that the adoption rates that we observe for smart phones and Mobile Internet Devices (MIDs) will likely continue. While these mobile devices have ever increasing processing, storage, and network capabilities, they also tend to have limited input and output. One key challenge for enabling the full utilization of the capabilities of these devices will be overcoming the limitations of their interfaces. Dynamic Composable Computing is one approach for overcoming some of the limitations of mobile devices by enabling the impromptu assembly of a logical computer from the best available set of nearby wireless components [12] [19].

The display characteristics of a mobile computer are one of the most defining aspects of the device. The size of the display has implications for both the mobility and the form factor of the computer. For example, there are several laptop computers on the market that offer very similar computing capabilities, but are packaged differently and offer different screen sizes and resolutions. Tablet PCs often offer similar performance to laptops but are designed to be operated with a stylus. Ultra Mobile PCs (UMPCs) provide users with significant computing resources and run standard desktop operating systems and applications. However, these systems are designed to be more portable and as such have much smaller displays. Smart phones are smaller still and require highly



Fig. 1. Four individual tablet computers, linked only by a wireless network, forming a multi-display composition resulting in a single logical display.

tailored applications to accommodate the limited screen real estate. While in a modern version of each device the power of the processor subsystem scales moderately with physical size, the display limitation provides the largest differentiating factor in the user experience offered by each device.

In this work, we are exploring how the displays of several mobile computers can be joined together to gain more display area. In particular, multi-display composition is a system that uses the screen real estate of several mobile computers to form a larger logical display. Originally, this technique was developed as a mechanism for overcoming the display size limitations of small devices like UMPCs and smart phones. The intent was to provide a mechanism for obtaining enough screen space on which to run traditional legacy desktop applications. However, multi-display composition can also be applied to larger mobile computers. For example, Figure 1 shows the formation of a single logical display from four separate tablet PCs effectively resulting in the *ad hoc* creation of a tabletop display.

The contributions of this paper are twofold: First, we present the *ad hoc* wireless multi-display composition technique which supports running legacy applications on a logical display formed by combining the display resources of several mobile computers. Second, we discuss the results from a study where groups of collocated individuals completed several tasks using a multi-display composition on two different types of mobile computers.

2 Related Work

Research into alternative display technologies reveals several techniques which could decouple the size of the display from the size of the device. Technologies such as electronic ink and organic LEDs may eventually allow for the mass market production of displays that have a large area and high resolution while also being easily rolled or

folded to fit in a pocket or bag for portability. Micro projectors [2] [1] and head mounted displays [15] [18] use miniature displays and optics to create large images out of very small packages. In addition to these novel display technologies, another approach to gaining more display area on a mobile device is to utilize multiple traditional LCDs. The Nintendo DS portable gaming system utilizes a clam-shell design with two smaller displays. The hinged design offers a small form-factor to support mobility while not in use. However, when open, the two displays can be used to increase the amount of available screen real estate. Similarly, Chen *et al.* demonstrated an electronic book reader with two displays that attach and fold against each other [3], and Siftables explores the use of several very small displays [11].

More directly related to multi-client composition is the work on ConnecTables [16] which demonstrated the ability to dynamically link two mobile computers together to gain increased screen area. Hinckley's work on Synchronous Gestures [6] and Stitching [7] brings a similar concept to tablet PCs and explores different mechanisms for initiating the link between the computers. This set of related work all relied upon custom applications to utilize the display resources of the mobile computers. Furthermore, these only demonstrate connecting two devices together. In contrast, as we will show with our multi-display composition system, we provide support to run legacy applications unmodified across multiple mobile displays and demonstrate the use of four displays in our study. Finally, our user study, in contrast to Hinckley *et al.*'s [7], is explicitly focused on the multi-display aspects of the system and uses groups of participants with pre-existing social relationships.

3 Multi-Display Composition

Multi-display composition performs its display sharing at the windowing system level, and as such, is related to a large body of work exploring the interplay between displays and high-speed networks. The X Window System and Virtual Network Computing (VNC) [14] are two well established examples of systems that support sending graphical data over the network to a remote computer system. In this work, we continue this trend of using network enabled displays but focus on how it can be applied to a group of *ad hoc* mobile devices. The advances in the capabilities of mobile computers have resulted in the ability to perform similar types of display sharing across wireless networks such as WiFi (e.g. IEEE 802.11n) and on devices such as laptops and tablets. Even smaller handhelds, such as UMPCs and MIDs, are gaining enough computational power for this type of display sharing.

By using a multi-display composition on mobile devices, several novel usage scenarios emerge. With a composition in which the displays of several devices mirror a single source display, a group of collocated users can each use their own device to access the same information. For example, instead of passing a camera phone between members of a group of people to view a captured photograph, a multi-display composition could let every person view the photo from their own phone similar to the system created by Clawson *et al.* [5]. Alternatively, several tablet computers can be placed on a table and used together to form one large aggregate display surface. With an extended multi-display composition, the displays of several mobile devices

are bound together in a similar way creating one display which spans the computers. When applied to a group of four tablets, an *ad hoc* tabletop display is formed (Figure 1). These examples utilize the devices from a group of individuals to create a larger display system. It is possible that a single individual may also have access to multiple mobile devices. For example, if Weiser's vision of ubiquitous tabs and pads comes to pass [20], a single user could form multi-display compositions using devices found serendipitously in the environment. We are already starting to see signs of this type of usage where people use both a laptop and mobile phone or PDA in combination [7] [9].

Our implementation of multi-display composition is built on top of the VNC remote display sharing protocol [14]. This protocol allows a user to see and interact with the framebuffer of one computer using another remote computer connected by an IP network. We extended the VNC protocol and created a custom X server that is headless and not attached to a physical framebuffer. By decoupling the display from the host device, we are able to extend the server so that the framebuffer can be arbitrarily resized at run-time and shared over the network.

This design choice is a key contribution in that it allows for any legacy X application to be used in this multi-display environment that can dynamically grow and shrink as needed; custom applications do not need to be written nor existing software modified. The applications are rendered like normal into X's framebuffer, and the multi-display composition system manages all of the issues associated with distributing the framebuffer, and therefore the applications' display, across the devices. While we chose X and VNC as the basis of our system, several alternative implementations could also be explored where the display sharing is implemented at other layers in the windowing system. For example, two alternatives would be to share and distribute OpenGL [8] or the X protocol to multiple devices. Similarly instead of the custom X server, a modern compositing windowing system such as Mac OS X or Windows Vista already has the required decoupling between the rendering of windows and graphical information and the framebuffer.

Our VNC server extensions build on the TurboVNC¹ implementation of VNC. The performance of this software has been improved by using a high-speed vector optimized JPEG library. Our implementation incorporated these optimizations for the VNC server into LibVNCServer², an open source library that supports the creation of custom VNC servers. In turn, this library was linked against our custom X server with multi-client support. Two types of VNC clients were created. The native Linux optimized TurboVNC client was modified to support the additions needed for multi-client mode. Furthermore, the Java version of the Tight VNC client was modified to provide a cross platform client solution.

The system is started by one mobile device running the server. The software creates a new X session, and VNC client is attached to this server from the same device to show the first portion of the desktop so any applications run within the X server can be shared transparently with other mobile devices. When a client connects in legacy VNC mode, the entire framebuffer is shared. In a multi-client display composition, this configuration allows for a single display to be *mirrored* on several other devices.

¹ <http://www.virtualgl.org/About/TurboVNC>

² <http://libvncserver.sourceforge.net/>

As additional clients connect the new *extended* mode, the server grows its logical framebuffer by the size of the connecting device. The client's viewport is set to the newly created region to display its portion of the overall desktop. The user is then free to utilize the new screen real estate shown on the new device as if it were part of the original display. The system can handle an arbitrary number of displays, limited only by the processing capability of the device acting server and available bandwidth. We have tested the system with up to six displays in a grid over IEEE 802.11a. The policy for choosing a direction to grow the display (horizontally or vertically) and where to place the incoming display connection is currently programmed into the server. We have a prototype implementation of a mechanism that allows a user to manually reposition the viewports of each connected device, and in the future we want to explore the use of sensors to automatically determine the relative location of devices [10]. When devices of the same resolution are connected, the system fully tiles the space in a grid. Devices with heterogeneous display sizes are also placed into the grid using the resolution of the connecting device to grow the display as needed. However in this case, it is likely that there will be inaccessible portions of the display (just as when using two monitors of different resolutions in a multi-monitor configuration). Currently these portions of the display are rendered, but not visible. Solutions for existing multi-monitor systems would likely be useful to implement in this system to address this issue.

Finally, we are using the Composition Framework [12] [19] to manage the multi-display composition. The Composition Framework provides a mechanism for the system to discover the devices and sharable services connected to the same (wireless) network. It also provides a user interface for managing the sets of connections needed to form the multi-display composition. Previous work pilot testing the interface has shown that participants could effectively use it to create and manage compositions of different platform services.

4 Evaluation

As discussed above, we have a generic implementation of multi-display composition that runs on mobile devices connected to a wireless network. The Composition Framework supports the formation of *ad hoc* display compositions, and the geometry management of the system allows the display to dynamically grow and shrink across devices as needed. Given the above implementation, we wanted to explore how the system might be used and if it effectively allowed for the sharing of multiple mobile display resources. We conducted an evaluation on the composition of a large logical display and decided to focus on pen-based mobile computers and how a group of users might come together to use a set of devices as an *ad hoc* tabletop display. In our study, participants used the multi-display system simulating a situation where each person would carry and contribute their own mobile device to work jointly on a task.

4.1 Participants

We recruited participants from our organization (primarily interns) by word of mouth. None of the participants had any previous experience with our system, and we recruited

participants so that each person knew all of his or her group members for at least one month prior to the study. Each participant was compensated with a \$50 gift card for a single 90 minute session.

We recruited five groups of four people, and five groups of two, for a total of thirty participants. Eight of the participants were female. The median age of the participants was 23 (ranging from 19 to 31) and they knew each other a median of 2.75 months (ranging from 1 to 12 months). All of the participants reported extensive computer use with a median of 10.5 years of use (ranging from 6 to 21 years). Twenty-six of the participants used a laptop as their primary computer while the remaining four indicated they used a desktop. Only five participants indicated they had used a tablet more than once, and only one participant indicated daily use. All 30 participants had used some form of remote display sharing application (VNC, Windows Remote Desktop, NetMeeting, LiveMeeting, Citrix, X). Nine of the participants indicated they used it for sharing a display remotely with another individual, fourteen reported using remote display sharing to gain access to a remote system for individual use. Four participants used display sharing for both of these usages, and three did not specify their usage.

4.2 System Configuration

We used two different sets of mobile devices for the study: Lenovo X61 Tablet PCs and Sony Vaio UX71 UMPCs. As indicated above, we recruited two different group sizes (two and four people) and maintained a one-to-one mapping between the number of devices and the size of the group. The devices for the groups of four people were pre-configured in an array of two columns by two rows (2x2) (Figure 1), while the devices for the groups of two people were placed in a single row (2x1). For all of the configurations, the devices were positioned side-by-side in a landscape orientation with minimal space between them.

The X61 Tablets have resolution of 1400x1050. The active portion of the display measures 24.6x18.5cm resulting in a pixel density of approximately 140 dpi. Discounting screen real estate used by the system, the 2x1 configuration has a total resolution of 2780x970, while the 2x2 has a resolution of 2780x1940. The UX71 UMPCs have a resolution of 1024x600. The active portion of the display measures 9.9x5.8 cm resulting in pixel density of approximately 260 dpi. Again, discounting screen real estate used by the system, the 2x1 configuration has a total resolution of 2028x520, while the 2x2 has a resolution of 2028x1040. The devices communicate through a dedicated 802.11a wireless access point.

In addition to the basic multi-display composition system, we implemented a special mechanism to allow users to drag and drop between devices using pen input. Even though the displays logically present a single screen, the bezels of the devices present a barrier that one would not be able to cross with a standard pen-based input mechanism. In particular, it would be impossible to start a drag on one device and finish on another. Inspired by the Rekimoto's Pick-and-Drop technique [13], we utilized a hardware button on the device's bezel to allow a user to enter a mode where they could lift the pen mid-drag and place it down on another device to finish the drag operation.




Fig. 2. Sample study data from the flower image sorting task (left) nutrition spreadsheet task (center) and movie graph task (right).

4.3 Tasks

We designed three tasks for the groups of participants to perform in our study: a sorting task of flower photographs, the use of a spreadsheet containing nutritional data, and an analysis of a graph of movie actors/actresses and movie titles (Figure 2). The primary purpose of the tasks was to have a structured way in which we could engage the participants with the system. Furthermore, the tasks represented examples that had the potential to benefit from the system’s larger aggregate display size. Finally, these are tasks that use traditional desktop applications which might be started by an individual, then grown across a set of collaborators’ devices. While we collected data on the group’s performance of each individual task, we were primarily interested in the overall experience of the group using the system. As such, the tasks were chosen so they spanned a range of visual representations of data (images, text, and abstract graphics) as well as in the nature and amount of required interaction. Furthermore, the tasks were constructed so that they could be performed in a limited number of ways to reduce variation and dependence upon creative user input. For each type of task, three different versions (one practice, two test) were created to minimize learning effects.

For the flower image sorting task, participants were asked to sort pictures of 32 flowers into six categories using drag and drop. Seven windows were displayed with one window being the source of all images and the remaining six were destinations. Each window showed image thumbnails, and the image could be opened for closer inspection in a separate window by double clicking on the thumbnail. A total of 132 images were obtained with each image ranging in size between 1600x1200 and 4368x2912 pixels. The images for each category were randomly shuffled and divided evenly into three sets. Finally, two representative images were manually selected from each category to serve as examples and were provided to the participants to facilitate sorting. Participants were also aided by a paper printout which showed the examples printed in color and the final number of flowers for each category.

For the nutrition spreadsheet tasks, participants were asked to perform operations on the U.S. Department of Agriculture’s abbreviated spreadsheet of nutritional information [17]. This spreadsheet contained 7520 rows of data, each representing a different type of food. The 51 columns of the spreadsheet contained various nutritional data about each food item. The groups of participants were asked to perform several operations on this data including sorting, copying rows of data, and searching for items with specific

characteristics. A paper form was provided to the group with the instructions for the task and provided space to record answers.

For the final task, the group of participants interacted with a large undirected graph which contained information about actors, actresses, and the movies in which they performed. The source data is from the IEEE Infovis 2007 contest³. The data were filtered and rendered into static images of undirected graphs with two types of nodes. The first type contained the title of the movie, while the second type presented the name of an actor or actress. Edges linked people to the movie(s) in which they performed. This resulted in graphs with between 20 and 22 actor nodes, 76 and 95 movie nodes, and 114 and 121 edges. The rendered images were between 3007x2658 and 3194x3572 pixels. The groups of participants filled out a paper form asking them to list all of the movies which contained actors or actresses in exactly six movies, and to list the movies which had either two or three actors shown in the graph.

4.4 Procedure

The study began with the group members filling out a survey for demographic data and information about their computer usage. The system was described, and the participants were provided with some usages scenarios of how the system might be used and constructed from individuals' mobile devices. Next, a practice session began. Here a researcher showed the participants how to use the stylus, how the system presented a single logical display spread across the devices, and how to use the cross device drag and drop capability. Each participant was encouraged to interact with the system to gain some experience. Next, the experimenter explained each of the three tasks in detail and allowed the group to practice each task. This practice provided further time to become familiar with the system and ensured that the participants understood the tasks they were performing. After answering any questions, the practice session came to a conclusion.

At that point, the three tasks were performed on the first set of devices (either the UMPCs or the Tablets). The order of tasks was selected randomly, while the order of devices was counterbalanced across groups. The participants were given five minutes to complete as much of each task as quickly and accurately as possible. While five minutes is a short duration, we chose this amount to reflect our scenario of the group of individuals joining their displays for a short period of time to perform a specific task.

After each task, the participants individually completed a questionnaire. The questionnaire asked them to list three positive and three negative aspects of using the system for the given task. Furthermore, it contained Likert questions (on a nine point scale) from the Questionnaire for User Interaction Satisfaction (QUIS) [4]. In particular it contained the questions about the Overall User Reactions and a subset of questions related to the display. Additionally, several questions were created in a similar style asking for ratings on the ability to participate in the task, feedback about the spacing between group members, etc.

After all three tasks and questionnaires were completed for the first type of display, the researcher swapped devices for the final part of the experiment. The researcher showed the participants the differences between the computers (how the devices had

³ <http://eagereyes.org/InfoVisContest2007Data.html>

different types of pen input, etc.). After answering any questions, the last three tasks were performed by the group, again filling out the post-task questionnaires. At the conclusion of the study, the participants completed one more questionnaire asking them to rate each device and condition pairing performed (six in total), to subjectively rate their satisfaction with each of the devices, and to comment on their overall experience.

5 Findings

We analyzed the data collected during the study examining trends about the participants' perception of the system in the quantitative measures collected from the post-task Likert scale questionnaire and the exit questionnaire. We also report themes in usage and feedback about the multi-display composition system. In particular, we examine the positive and negative comments recorded on the post-task questionnaire. This information is further substantiated by the overall comments participants provided at the end of the study and by the observations recorded in the researcher's notes during the trials.

5.1 Overall Utility

Overall, the participants rated the extended multi-display composition system positively. Aggregated across all of the data, the QUIS Overall User Reaction score was moderately positive with a mean score of 6.05 (SD=1.72) from a range of 1 to 9 with higher values being more positive. Similarly, the absolute rankings of the devices provided at the end of the experiment about how helpful the system was for performing the tasks show similar results with a mean "helpfulness" rating of 7.08 (SD=1.29) (again on the 9 point scale).

Examining these results in more detail reveals differences based on the type of device and the size of the group (Figures 3 and 4). For the overall QUIS user reaction score, there was a significant difference found based on condition (Kruskal-Wallis, $p < 0.001$). Further analysis reveals participants preferred the tablets (M=6.71, SD=1.47) relative to the UMPCs (M=5.38, SD=1.70) (Wilcoxon Signed Ranks, $p < 0.001$) and groups of two rated the multi-display composition higher (M=6.72 SD=1.45) than groups of four (M=5.71 SD=1.76) (Mann-Whitney $p < 0.001$). The responses about device preference from the exit interview follow the same pattern with a significant difference between the conditions (Kruskal-Wallis Test, $p < 0.001$). Tablet conditions again rated more favorably (M=7.09 SD=1.29) than the UMPC conditions (M=4.94 SD=2.04) (Wilcoxon Signed Ranks Test, $p < 0.001$), and the groups of two rated the task more favorably (M=6.82 SD=1.39) than groups of four (M=5.62 SD=2.16) (Mann-Whitney Test, $p < 0.001$). These trends also continue for our other subjective measures with the tablets generally rated higher than the UMPCs and similarly, the groups of two rating the system more favorably than the groups of four (Figures 5 – 10).

Examining differences between tasks also reveals significant differences. The overall QUIS user reaction score shows a significant difference based on task (Friedman, $p = 0.024$). Further analysis shows there are pairwise differences between the nutrition spreadsheet task and the flower photograph sorting task as well as between the nutrition




Fig. 3. QUIS Overall User Reaction scores (mean and 95% C.I.).




Fig. 4. Overall helpfulness of the display for a task.




Fig. 5. Ability to participate in task.




Fig. 6. The spacing between yourself and other participants.

task and the movie graph task (Wilcoxon Signed Ranks: $p = 0.007$ and $p = 0.010$ respectively). The nutrition spreadsheet task was rated least favorably (photo: $M=6.20$ $SD=1.62$, nutrition: $M=5.64$ $SD=1.63$, movie: $M=6.20$, $SD=1.85$). Similar results are true for the helpfulness ratings collected with the exit questionnaire. There is an overall effect for task type (Friedman, $p = 0.005$) and significant differences between the nutrition task and both the photo and movie tasks (Wilcoxon Signed Ranks: $p = 0.001$ and $p = 0.003$ respectively). Again the spreadsheet nutrition task was rated least favorably (photo: $M=6.58$ $SD=2.49$, nutrition: $M=5.20$ $SD=2.43$, movie: $M=6.27$, $SD=2.60$).

Together, these quantitative data show interesting results. While we originally inspired to create this system to increase the capabilities of UMPCs, it appears overall that it provides more benefit for tablets. Also, these data indicate that more pixels are not always better. While the use of four devices offered more screen area for the applications, it also meant there were four people trying to use the devices simultaneously, negating the advantages of having a larger display area.

5.2 Usage Themes

Given the overall trends in the quantitative data, we next turn to the qualitative information. This data was collected by asking the participants to list three positive and three negative aspects of the system after each task (for a total of six times per participant). The data from general comments provided by the participants at the end of the experiment and observation recorded during the study are further used to examine the usage patterns.

Display Size: One of the first positive comments participants often made was about having more screen real estate with the multi-display composition on the mobile devices. This aspect of the system was by design, but the participant comments and other data reveal several different ways the screen space was used. For the tasks that primarily utilized a single window (the spreadsheet and the movie tasks), nearly all of the groups immediately maximized the window to fill all of the available displays. Participants commented that this let them see more of the data and led to less scrolling.




Fig. 7. Ability to provide input.




Fig. 8. Visibility of computer screen.




Fig. 9. Characters on the screen.




Fig. 10. Amount of info. displayed.

This held true even though most participants thought the space between displays taken up by the bezel was a negative aspect. For the photo sorting task, several groups used the screen space to spread out the windows to minimize overlapping. A few groups also commented on the advantage of the larger screen for seeing the larger view of a photograph. These comments held true for both the tablet and UMPC conditions: “the screens were small, but did fit a lot of things for a small screen”.

Awareness: The participants also indicated that they appreciated the awareness of the group’s activity that the system provided. The participants commented that they liked the ability to see what the others were doing, being able to quickly check with other group members about an activity, and to point out information on the composed display. Given that only one person could provide input at a given time, there was often the ability for the other group member(s) to help guide the navigation through the interface either looking ahead or for confirming the proper input. The shared display was also used for a common frame of reference where participants would point to the screen with either the stylus or a finger so that everyone could focus on the same information. This awareness was used for confirming a selection in the spreadsheet task, for pointing to a name in the movie task, and for building consensus about the type of a given flower in the photograph sorting task. This awareness results from using the multiple mobile computers as a single logical display, and would likely be absent if other display sharing models were used.

Collaboration: Many of the participants liked the ability to parallelize the task. Given the time limit of the tasks imposed by the study, there was an incentive for the groups to attempt to optimize for efficiency. In doing so, one common strategy adopted was to divide the tasks in different ways. The participants also leveraged the fact that the single logical display was spread across several physical devices, and used the device as a unit for dividing the task or interaction. For example, some of the groups of two very quickly adopted language such as “mine” and “yours” referring to either the device or data on the device directly in front of them or in front of their partner respectively.

The division of labor also let participants manage the physical area of the display. Several participants commented negatively about the difficulty of reaching to the other side of the display to provide some needed input. While this was seen as a problem

by some, others commented that the group nature of the task also provided a solution where “everyone clicked on their closest screen.” And indeed, observations revealed that sometimes a participant would attempt to reach across all of the devices and fail to provide the needed input, so a closer person would finish the interaction. Some groups developed the strategy for some of the tasks of having a person “assigned” either implicitly or explicitly to a region of the interface, for example operating the scroll bar or menus. While the participants commented on using this approach, observations indicated that these divisions were very flexible and fluid.

Input: Another form of shared interaction which occurred spontaneously in several groups centered around the functionality for performing a drag across device boundaries. During the flower sorting task which required many drag and drop operations, some of the groups developed a strategy where one person would start the drag, another would press the button on the tablet bezel to initiate the needed mode, and potentially a third person (for the groups of four) finished the drag on another device. Beyond this split operation, a few of the groups performed this task without speaking. The shared objective and visibility of the operations on the single logical display provided sufficient information for the group to successfully perform this operation in a very fluid fashion.

Another issue revealed in the user comments and observations related to the single input nature of the system. The devices, the underlying windowing system, and applications were not multi-cursor aware. On the resistive touch screen of the UMPC, problems only occurred when multiple people touched the display. However, the tablets also tracked pen hover which could cause erratic mouse behavior when more than one person put their pen in the proximity of a device. While the tablets presented this problem, the overall trend of the tablets being rated more favorably continued when participants were asked to rate their ability to provide input to the system (Figure 7). As devices such as the iPhone and associated application that support multi-touch become more common, these issues may become less important.

Physical Aspects: Many participants commented about their position around the table and their ability to view the composition of displays. For example, some participants made comments about the “crowding of people”. Others remarked about the viewing angle and effectively looking at the display from the side. It was also observed that some of the participants, especially in the groups of four, would stand up or kneel on their chair to be in a position where they could lean over the table. While participants commented that their spacing within the group was less than ideal, it appears not to have been too negative of a factor. The quantitative data reveals that participants rated the spacing favorably overall ($M=6.50$ $SD=2.20$). The overall trend of tablets being rated better than UMPCs, and groups of two having higher scores continues (Figure 6).

Some participants also commented that some of the text was small or that the visibility could be improved. Here the quantitative data again shows that the UMPC performed worse than the tablet PC (Figures 8 and 9). In addition to the UMPC having fewer pixels and a smaller screen, it also has a higher pixel density. Together, those factors result in the bitmapped information appearing smaller on these devices. And while the multi-display system provided more pixels to see more content, it did not overcome this problem for the participants. This data reveals that an alternative approach for using multiple devices might be worth exploring. In particular, instead of

just expanding the size of the virtual screen to fill all of the available pixels, the system could use the extra screen real estate to also provide some magnification of the screen. Depending on the number of displays and magnification applied, the system could still increase the number of available pixels for displaying information, but also use more area for each pixel.

6 Conclusions

Overall, our study provided insight that groups of participants were able to effectively use a multi-display composition involving several mobile devices as a unified display. The study revealed an interesting trade-off between the size of a composite display and correspondingly the number of participants in the group. Even though groups of four had more screen real estate to perform the tasks, they generally rated the use of the system less positively than the groups of two. There could be a number of factors leading to this result ranging from the need to coordinate more people, reduced visibility and the need to provide input to a physically larger device.

The trend in our results for the differences between devices is less surprising. Each tablet provides a larger individual display which is likely more usable than a smaller UMPC display. And while the UMPCs were not the preferred devices, our data indicates that the system was still usable. When using smaller devices such as UMPCs or MIDs, an especially rich area for future work will be to explore how different configurations of multi-display composition (mirroring and extending) compare to more traditional collocated collaboration techniques on the same devices.

The spacing between the individuals in the group also presents an interesting finding. In previous work by Hinckley *et al.*, which also examined the joint use of two tablets by two individuals, the participants were reluctant to keep the devices in contact with each other [7]. While we did receive some negative feedback about the spacing between group members, the overall ratings for this issue were positive. Furthermore, none of the participants asked to separate the devices during our study. One possibility for the alternate findings is the different visual presentation and functionality of the two systems. It is possible our tiling of the entire desktop across the computers increased the need or desire to keep the devices in direct physical proximity. Secondly, Hinckley *et al.* used pairs of participants who did not know each other before the study, whereas in our study there was a preexisting relationship of at least one month. Future work will be needed to better understand the dynamics of group interaction with the joint use of mobile devices.

In conclusion, we have presented multi-display composition, a technique for supporting the collocated display sharing of mobile devices. We described our implementation of the system and discussed how it can be used to run legacy applications on a logical display formed from several mobile computers. Our study examined a specific usage scenario where a group of collocated users collaborate using an *ad hoc* tabletop display composed from pen-based computers. Our findings indicate that the system was generally rated favorably and groups of two people using tablet computers provided the most positive results. We also found several interesting themes of usage relating to the collaborative practices adopted by the groups using our system and some of the

technical challenges that should be addressed. Overall, this work shows multi-display composition provides a useful technique for opportunistically overcoming the display limitations of mobile devices.

References

1. G. Blasko, S. Feiner, and F. Coriand. Exploring interaction with a simulated wrist-worn projection display. In *Proceedings of the ISWC '05*, pages 2–9. IEEE, 2005.
2. X. Cao and R. Balakrishnan. Interacting with dynamically defined information spaces using a handheld projector and a pen. In *Proceedings of UIST '06*, pages 225–234. ACM, 2006.
3. N. Chen, F. Guimbretiere, M. Dixon, C. Lewis, and M. Agrawala. Navigation techniques for dual-display e-book readers. In *Proceeding of CHI '08*, pages 1779–1788. ACM, 2008.
4. J. P. Chin, V. A. Diehl, and K. L. Norman. Development of an instrument measuring user satisfaction of the human-computer interface. In *Proceedings of CHI '88*, pages 213–218. ACM, 1988.
5. J. Clawson, A. Volda, N. Patel, and K. Lyons. Mobiphos: a collocated-synchronous mobile photo sharing application. In *Proceedings of MobileHCI '08*, pages 187–195. ACM, 2008.
6. K. Hinckley. Synchronous gestures for multiple persons and computers. In *Proceedings of UIST '03*, pages 149–158. ACM, 2003.
7. K. Hinckley, G. Ramos, F. Guimbretiere, P. Baudisch, and M. Smith. Stitching: pen gestures that span multiple displays. In *Proceedings of AVI '04*, pages 23–31. ACM, 2004.
8. G. Humphreys, M. Houston, R. Ng, R. Frank, S. Ahern, P. D. Kirchner, and J. T. Klosowski. Chromium: a stream-processing framework for interactive rendering on clusters. In *Proceedings of SIGGRAPH '02*, pages 693–702. ACM, 2002.
9. H. M. Hutchings and J. S. Pierce. Understanding the whethers, hows, and whys of divisible interfaces. In *Proceedings of AVI '06*, pages 274–277. ACM, 2006.
10. G. Kortuem, C. Kray, and H. Gellersen. Sensing and visualizing spatial relations of mobile devices. In *Proceedings of UIST '05*, pages 93–102. ACM, 2005.
11. D. Merrill, J. Kalanithi, and P. Maes. Siftables: towards sensor network user interfaces. In *Proceedings of Tangible and embedded interaction*, pages 75–78. ACM, 2007.
12. T. Pering, R. Want, B. Rosario, S. Sud, and K. Lyons. Enabling pervasive collaboration with platform composition. In *Proceedings of Pervasive '09*, 2009.
13. J. Rekimoto. Pick-and-drop: a direct manipulation technique for multiple computer environments. In *Proceedings of UIST '97*, pages 31–39. ACM, 1997.
14. T. Richardson, Q. Stafford-Fraser, K. R. Wood, and A. Hopper. Virtual network computing. *IEEE Internet Computing*, 2(1):33–38, 1998.
15. M. Spitzer, N. Rensing, R. McClelland, and P. Aquilino. Eyeglass-based systems for wearable computing. In *Proceedings of the ISWC '97*. IEEE, 1997.
16. P. Tandler, T. Prante, C. Müller-Tomfelde, N. Streitz, and R. Steinmetz. Connectables: dynamic coupling of displays for the flexible creation of shared workspaces. In *Proceedings of UIST '01*, pages 11–20. ACM, 2001.
17. U.S. Department of Agriculture Agricultural Research Service. *USDA National Nutrient Database for Standard Reference, Release 20.*, 2007. <http://www.nal.usda.gov/fnic/foodcomp/Data/SR20/dnload/sr20abxl.zip>.
18. K. Vadas, N. Patel, K. Lyons, T. Starner, and J. Jacko. Reading on-the-go: a comparison of audio and hand-held displays. In *Proc. of MobileHCI '06*, pages 219–226. ACM, 2006.
19. R. Want, T. Pering, S. Sud, and B. Rosario. Dynamic composable computing. In *Proceedings of HotMobile 2008*, February 2008.
20. M. Weiser. The computer for the twenty-first century. *Scientific American*, pages 94–104, 1991.