A Comparative Evaluation of Touch-Based Methods to Bind Mobile Devices for Collaborative Interactions

Tero Jokela, Andrés Lucero
Nokia Research Center
P.O. Box 1000, FI-33721 Tampere, Finland
{tero.jokela, andres.lucero}@nokia.com

ABSTRACT
We present a comparative evaluation of two touch-based group-binding methods, a leader-driven method and a peer-based method, against a more conventional group-binding method based on scanning and passwords. The results indicate that the participants strongly preferred the touch-based methods in both pragmatic and hedonic qualities as well as in the overall attractiveness. While the leader-driven method allowed better control over the group and required only one participant to be able to form a group, the peer-based method helped to create a greater sense of community and scaled better for larger group sizes and distances. As the optimal group-binding method depends on the social situation and physical environment, the binding methods should be flexible, allowing the users to adapt them to different contexts of use. For determining the order of the devices, manual arrangement was preferred over defining the order by touching.

Author Keywords
Collocated interaction; mobile phones; user interfaces; device ecosystem binding; group association; pairing.

ACM Classification Keywords
H.5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Mobile devices were originally conceived as, and have traditionally been, very personal devices targeted at individual use. Recent advances in sensor and short-range communication technologies offer new opportunities for collaborative use of mobile devices. Groups of collocated users can couple their devices together and create ecosystems of interaction [21]. This allows the users to engage in collaborative activities and experiences with their mobile devices, thus shifting from *personal-individual* towards *shared-multi-user* interactions. Examples of applications that would benefit from such collaborative use of mobile devices include sharing of digital content, collaborative creation and editing of content, and different kinds of games. In many of these applications, it would be natural to utilize spatial interactions in the shared space, for example, throwing virtual objects such as files between devices. However, finding the positions of the devices has presented a challenging problem, requiring the use of special tracking equipment or dedicated infrastructure.

Before a group of users can engage in collaborative interactions with their mobile devices, the multi-device ecosystem must first be set up. This involves initiating the necessary system and application software in all devices. The devices must become aware of the other devices existing in the proximity, and the devices intended to participate in the ecosystem must be identified. A communication channel then needs to be established between the devices participating in the ecosystem, in order to allow exchange of data and coordination of the interactions. Wireless short-range communication technologies such as WLAN or Bluetooth are typically used to exchange data between devices. The process of setting up the ecosystem is generally known as device binding or ecosystem binding [21] (also known as device association, pairing, or coupling [3]). As the intention is to enable spontaneous interactions, it should be possible to bind devices having no prior knowledge of each other in a fast and easy way. If the process of binding devices is too complicated or tedious, the users might lose interest in using multi-device interactions in the first place. As the wireless connections provide no physical indications (for example, cables) of which devices are actually connected, the binding process should provide sufficient security and cues so that the users can ensure that the right devices are connected.

In this paper, we are concerned with device-binding methods for establishing an ecosystem of mobile devices to support collaborative interactions within small-to-medium-sized groups of collocated users. While the problem of a single user pairing two devices has been extensively studied in prior research, researchers have started to address more complex scenarios involving multiple users and devices only recently. In particular, we focus on methods based on device proximity and touch interactions, which have been found to be intuitive and easy to explain, but which have been little explored in the literature [2]. We present a comparative evaluation of two touch-based group-binding methods, a leader-driven method called Host and a peer-

CHI 2013, April 27–May 2, 2013, Paris, France.

based method called Ring, against a more conventional method called Seek, which is based on scanning the available devices in the proximity and passwords for security. While most earlier studies on device binding have focused on pragmatic aspects such as security and usability, we approach the problem from a broader user experience perspective, covering also hedonic aspects such as social and emotional factors, which have been shown to be important considerations when users select binding methods in real-life situations [10, 18]. We consider the complete group creation process in a realistic application context, including identification of the devices to participate in the group, initiation of the application software in all devices, and authentication of the connection. We also explore options to determine the device order during the group creation phase, in order to allow spatial interactions without dedicated tracking equipment. The evaluation results indicate that the participants strongly preferred touch-based methods over Seek. Several important differences were identified between leader-driven and peer-based methods. The optimal group-binding method was found to depend on various social and environmental factors, suggesting that the binding methods should be flexible to allow users to adopt different group creation strategies in different contexts of use. For determining the order of the devices, manual arrangement was preferred over defining the order by touching.

The rest of this paper is structured as follows. First, we provide a brief overview of the related work. We then give a detailed description of the three group-binding methods and the evaluation procedure. Finally, we present the results of the evaluation, followed by conclusions.

RELATED WORK

The problem of device binding has been extensively studied in the fields of human-computer interaction and security research. A wide range of methods for device binding has been proposed – in security research alone, over 20 different methods have been identified [17]. These methods vary in terms of device hardware requirements, amount of user involvement, and level of provided security.

The problem of device binding can be divided into two subproblems: device *identification* and *authentication*. Device identification involves selecting which of the devices available in the proximity should be bound with each other. The need for device authentication originates from the invisibility of wireless communications. As the users cannot see the wireless communication channels, they cannot be sure that they are really connecting to the other devices intended to, opening the possibility for so-called Man-in-the-Middle attacks. To counter this threat, a wide variety of methods have been proposed that authenticate the wireless connection over auxiliary communication channels (also known as Out-of-Band Channels), which can be perceived and managed by human users.

The most common device-binding methods today, such as those typically used in Bluetooth and WLAN networks, are based on scanning the environment for available devices and then presenting a list of the found devices to the user for selecting the other device to bind with. The authentication is based on short strings (also known as PIN codes) that the user is expected to copy or compare between devices. The authentication strings can be represented as numbers, words, graphical images, or audio signals in the user interface.

The proposed alternative methods include a variety of techniques based on synchronous user actions, for example, pressing buttons on both devices [19] or touching both devices [23] simultaneously, shaking the devices together [8], or bumping the devices together [6]. Bumping is also used in the popular commercial service Bump[1]. Further, device binding can be based on continuous gestures spanning from one device display to another [7]. Methods based on spatial alignment of the devices include pointing, for example, with laser light [15], touching [20], or placing the devices in close proximity of each other [12]. It is also possible to bind devices with various auxiliary devices, for example, tokens [1] or cameras [16]. Some of the proposed methods cover only device identification or authentication, while others combine both identification and authentication into a single user action.

The development of new binding methods has been largely technology-driven with little user involvement. As an example of a more user-centered approach, Chong and Gellersen [2] present a study on users' spontaneous actions for device binding. In the study, the users' were asked to invent methods for binding together low-fidelity acrylic prototypes of different devices. Device proximity and touch based methods were found to be among the most commonly proposed methods, and the physical contact of devices was also considered as the easiest method to describe and teach to another person. Still, there has been little work exploring such techniques in the literature.

Binding methods are not just means for connecting devices – they have strong social and emotional aspects. In real-life situations, the users do not always use the easiest or fastest method available, nor the one they like best. Many factors influence their choice of binding method, including the place, the social setting, the other people present, and the sensitivity of data [10, 18]. Users are willing to take security risks to comply with social norms [10].

The vast majority of prior research has focused on scenarios of a single user binding two devices with each other (for example, binding a headset with a mobile device). Only recently have researchers started to consider more complex scenarios involving multiple users and devices. Such multi-user scenarios differ in many respects from single-user

[1] http://bu.mp/

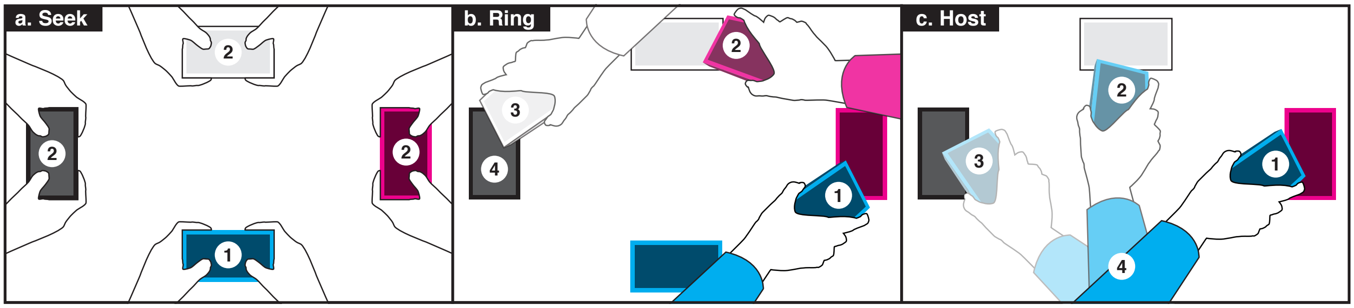

Figure 1. Three group-binding methods. a) Seek: a leader creates a group and shares a password (1), which is then entered in parallel by the other participants (2). b) Ring: one person starts the app and touches the next device to their right to add it (1), then others continue adding the next person to their right (2,3), and the last person completes the group (4). c) Host: a leader starts the app, adds people by touching all devices in counter-clockwise order (1-3), and puts the device on the table to complete the group (4).

scenarios, making the single-user device-binding methods not necessarily applicable to multi-user scenarios. In multi-user scenarios, communication between group members provides an additional source for potential errors. On the other hand, the users are typically willing to help each other and make decisions by mutual agreement, which reduces the amount of errors [11]. Methods that involve physical exchange of devices have been found to be unacceptable unless the users know each other very well, as the users are unwilling to hand in their devices to strangers [22].

Chong and Gellersen [3] present a framework that summarizes and classifies the different factors that influence the usability of spontaneous device binding, identifying technology, user interaction, and application context as the three most important criteria.

EVALUATION OF GROUP-BINDING METHODS

Objectives

In this study, we were primarily interested in three research questions. First, we wanted to compare touch and proximity-based methods for group binding against more conventional methods based on scanning and passwords. Second, we wanted to explore different ways to divide the group-binding task between the participants – in particular, we were interested in differences between leader-driven and peer-based methods. Third, we wanted to investigate possibilities to define the device order as a part of the group creation process, in order to allow implementation of spatial interactions without extensive tracking equipment.

Group-Binding Methods

To study these research questions in practice, we designed three different group-binding methods called Seek, Ring, and Host. The Seek method represented the conventional approach used, for example, in network games and was based on scanning for device identification and passwords for authentication. Both the Ring and Host methods used touch for device identification and authentication. The main difference between Ring and Host was that Ring was peer-based, distributing the group creation task between all participants, while Host was leader-driven, concentrating the group creation task on one participant. Additionally, Host utilized device gestures for some interactions. The Host method was based on the EasyGroups method [14] reported earlier. While all the methods were generic, we decided to study them in the context of a simple photo sharing application in order to provide a more realistic goal for the group creation task during the evaluation. The photo sharing application was a simplified version of Pass-Them-Around [13] and it allowed the users to browse a collection of photos stored in their own devices and supported spatial interactions of throwing photos between devices.

Seek

To set up a new group, one person (the leader) should start the Seek application on their device and create a new group (Fig. 1a). The application prompts the leader to join a WLAN network and enter a name for the new group. The application automatically generates a six-digit password for the group. The application then moves to the Table Overview (Fig. 2) showing all devices that are currently part of the group and their order as well as the group name and password. As new devices join the group, an animation shows how the device order changes on the table. To enable the users to identify the devices, the color of each device is indicated on the screen. The other persons can then join the group in parallel by starting the Seek application, joining the same network as the leader, and selecting the existing group from the list. The application then prompts the user to enter the password. If the password is correct, the device joins the group and moves to the Table Overview. If the order of the devices presented on the screen is different from the order of the devices on the table, the leader can correct it by dragging the devices to the right positions on the screen. The users can move to the Photo Sharing Mode by tapping their own piles of photos on the screen.

If a new person wishes to join an existing group, the person should start the Seek application and join the group in the same way as during the initial group creation phase. The leader can check the order of the devices on the screen and correct it if necessary. To leave the group, the person should press the “Exit” button on the screen.

Figure 2. The Table Overview during Seek.

Ring

To begin group formation, one person should start the Ring application on their device (Fig. 1b). This device automatically enters Discovery Mode and visual feedback is shown in portrait view to suggest holding the device vertically for a more comfortable grip. The person holding the device is instructed to touch the next device to their right. When the person moves their device close to the next device, the device detects the new device and the person holding the device is asked to hold their device still while the new device is added to the group. When the new device has been added to the group, the device exits Discovery Mode and moves to the Table Overview, which shows all devices that are currently part of the group and their order. The new device that was just added to the group now automatically starts the application and enters the Discovery Mode. The owner of that device is instructed to continue in the same way and touch the next device to their right. By asking the user always to connect to the next device to their right, we are able to define the order of the devices on the table based on the touching order. When all the devices around the table have been added to the group, the owner of the last device can complete the group by pressing the "Complete" button on screen. The users can move to the Photo Sharing Mode and start sharing pictures.

If a new person wishes to join an existing group, the person on the left side of the new person can press the "Add Device" button on the screen and touch the incoming person's device. The new person is then added to the right side of the person who just added them. The new group member can then continue adding new devices, or press the "Complete" button if there are no more devices to add. To leave the group, the person should press the "Exit" button.

Host

To set up a new group, one person (the leader) should start the Host application on their device (Fig. 1c). When the device is picked up from the table, it detects the pick-up gesture and enters Discovery Mode, which allows the leader to add new people to the group. The leader should then touch the other devices one by one in counter-clockwise order around the table (Fig. 3). The order of the devices on the table is automatically defined based on the touching order. A similar procedure (and visual feedback) as the one described for Ring is used to detect, connect, start the application, and join the group. When all the other devices have been added to the group, the leader should put their device back on the table. The device detects the gesture and exits the Discovery Mode and completes the group set-up. The persons can then start sharing photos between devices.

If a new person wishes to join an existing group, the person on the left side of the new person should pick up their device to enter Discovery Mode and touch the new person's device. The new person is then added next after the person who just added them. To leave the group, the person should pick their device up from the table and flip it upside down. The device detects the gesture and exits.

Prototype Implementation

We built prototypes of the three group-binding methods on Nokia N9[2] mobile devices running the MeeGo operating system. The prototypes were implemented in C++ on top of the Qt 4.7 software framework. QML and Qt Quick with OpenGL ES hardware acceleration were used for fluent animated user interface graphics. The N9's internal accelerometer was used for gesture detection in Host.

In all methods, the objective was to establish a WLAN connection between the devices. In Seek, each device was manually connected to the WLAN network. The device then scanned the network for available groups and presented a list to the user to choose from. In Ring and Host, touching was detected with Bluetooth-based radio technology, which was able to detect other devices at ranges closer than 20 cm in approximately 5 seconds. While the technology generally worked reliably, there were occasionally longer delays before the other devices were detected or detections of devices further away. The necessary connectivity and initialization information was then sent to the discovered device over Bluetooth. A daemon, which listened to a Bluetooth socket, received the

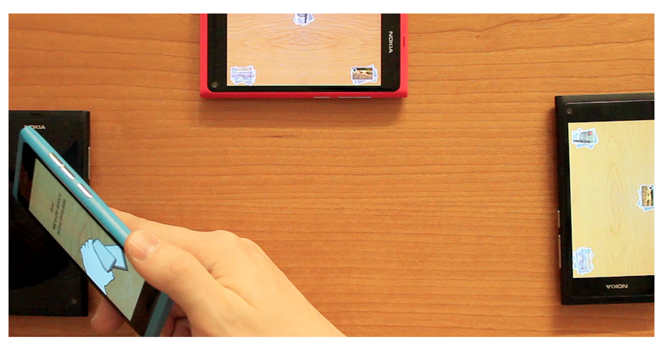

Figure 3. The Host method. The user holding the cyan device has connected the black (right) and magenta devices (top).

[2] http://swipe.nokia.com/

connectivity information on the discovered device and started the actual application, which connected to the correct WLAN network and joined the group. The prototypes were fully functional with real network communication, except for the security protocols, which were only simulated in the user interface.

Participants

We recruited a total of 24 participants for the evaluation by posting an advertisement on a local mailing list. Of the 24 participants, 20 were pairs of users, while the remaining four were individual participants. We preferred to recruit pairs of people who knew each other, so that the participants would feel more comfortable during the evaluation session. We assigned the participants into six groups of four users in the order they registered for the study. Each participant typically knew one other participant in the group, while the two others were strangers. Eight of the participants were female and 16 male. The ages of the participants varied between 23 and 45 years (M=33.6, SD=6.0). Three of the participants were left-handed and 21 right-handed. The participants represented a variety of different backgrounds, with eight participants having a software engineering background, 10 other technical background (for example, mechanical engineering), and six non-technical background (for example, administration or linguistics). The participants were fairly advanced users of technology: on a scale between 1 and 7 (1=novice, 7=expert), the participants rated their familiarity with technology above average (M=5.1, SD=1.2). All participants were active smartphone users and six of the 24 participants had used a Nokia N9 before the study.

Procedure

We organized a series of six evaluation sessions. The evaluation sessions were arranged in a usability laboratory of approximately 40 m^2 (430 sq ft) in size. Fig. 4 shows the evaluation setup. In each session, there were four participants and a moderator present. We used devices of four different colors (black, white, magenta, and cyan) and each participant was assigned a device with a different color. This provided a practical method of identifying the devices of the different participants during the evaluation session. The participants were given seats around a rectangular table of approximately 150x70 cm (60x27 inches) in size, one on each side of the table. The table was carefully selected so that there would be different distances between the participants and that the participants sitting on the short edges would have some difficulty reaching each other. The total durations of the evaluation sessions varied between 100 and 120 minutes.

As the participants arrived in the laboratory, the moderator guided them to their seats around the table and asked them to fill in a background questionnaire form. When all the participants had completed the forms, the moderator introduced the participants to the idea of collaborative use

Figure 4. Evaluation setup with four participants.

of mobile devices and demonstrated it with the photo sharing application. The participants were then given their own devices and they were encouraged to try throwing photos between devices. This small introductory task provided the users with an opportunity to become familiar with their devices. The moderator then explained to the participants that before they could share photos between devices by throwing, they first had to bind their mobile devices together into a group and the objective of the session was to evaluate different methods for that task. Before the actual evaluation started, the moderator informed the participants that some of the methods might require touching other devices and demonstrated how to do it in practice. The participants were then asked to practice touching with their own devices. We saw this training step necessary, because while many of the participants were aware of touching as an interaction technique, few had tried it in practice.

To begin the actual evaluation, the moderator showed a short video clip demonstrating the first group-binding method. The videos were prepared so that they simulated a situation of a participant observing another group of users using the method to create a group. We used video recordings to minimize the variations between the instructions that the different groups received. After the participants had watched the video, the moderator gave them the following task: “By using the method that was just demonstrated to you, create a group so that you can throw photos between your devices.” The moderator then observed as the participants tried out the method and only intervened if the participants clearly could not proceed with the method or there were some technical problems with the devices. The task was considered complete, when the participants could successfully throw photos between all devices. The moderator then asked everybody to leave the group and create another group with a different participant initiating the group creation. The moderator also asked at least one person to leave the group and rejoin it. Overall, each group tried each method two to four times.

After testing the method, the moderator asked the participants to fill in two validated questionnaires. The first questionnaire was NASA-TLX [4], which measures the subjective workload experience when performing a task. To gain a broader view of the methods, we extended the questionnaire with four additional scales: learnability, quickness, security, and overall preference. The second questionnaire was AttrakDiff [5], which measures the attractiveness of interactive products.

The same procedure was then repeated for the second and the third methods. We systematically varied the order in which the six groups were exposed to the three methods to counter-balance any learning effects.

After the participants had tested all the methods, the moderator interviewed the participants about their experiences with the methods. The interview was semi-structured and covered a variety of themes including general feedback about the different methods, perceptions about their learnability and security, as well as specific interaction techniques like touching and hand gestures. The moderator also showed the participants three pictures representing different scenarios and asked them to consider what would be the most appropriate method for creating a group in each scenario. The scenarios were: 1) meeting other family members in the living room at home, 2) meeting representatives of another company in a meeting room at the office, and 3) meeting friends in a busy café. The objective was to encourage the participants to think about different situations and environments and their social and physical characteristics. After the interview was completed, the moderator thanked the participants and gave them a movie ticket each to compensate them for their time.

All sessions were video recorded and interaction with the devices was logged. Two researchers independently analyzed the video recordings and wrote notes about their observations. The same two researchers then analyzed the data and built an Affinity Diagram [9] in a series of interpretation sessions. Each researcher individually studied the notes and grouped them into clusters of related items. The clusters then evolved to broader categories that were naturally revealed and were jointly revisited, discussed, and refined. In the end, the categories were processed into more general findings that form the core of the Results section.

RESULTS

We first give an overview of the quantitative results. We then present the qualitative results and contrast them with the quantitative results when relevant.

Extended NASA-TLX

Fig. 5 illustrates the results of the extended NASA-TLX questionnaire [4]. The main bars indicate the means for each subscale, while the error bars indicate standard errors. The original six subscales of NASA-TLX are presented on the left and the four subscales that we added for the purposes of this study (learnability, quickness, security, and overall preference) are on the right. As the participants were observed in groups, the responses of each participant were influenced by the other participants in the same group.

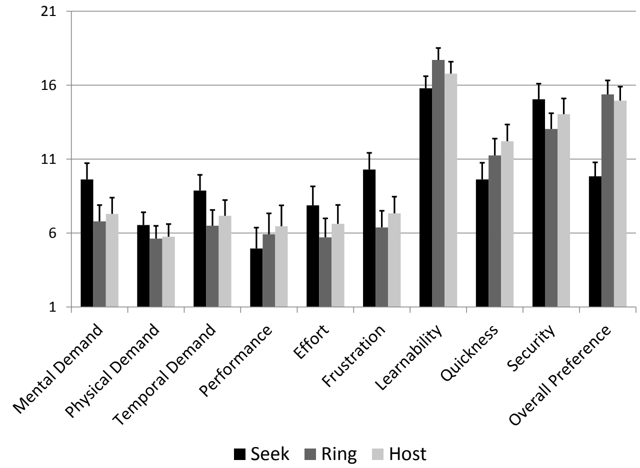

Figure 5. Extended NASA-TLX results.

Therefore, we used mixed model techniques to analyze the data with the binding method as a fixed factor and the groups and the participants nested in the groups as a random component. The results indicate that the binding method had a significant effect on mental demand ($F(2, 44.54) = 8.39$, $p = .001$), frustration ($F(2, 36.92) = 9.54$, $p < .001$), and overall preference ($F(2, 37.18) = 22.16$, $p < .001$). Pair-wise comparisons with Bonferroni correction show that the levels of mental demand ($p = .001$) and frustration ($p < .001$) for Seek were significantly higher compared to Ring, and that the level of overall preference was significantly higher for both Ring ($p < .001$) and Host is ($p < .001$) compared to Seek. There were no significant differences between Ring and Host on any of the subscales.

AttrakDiff

Fig. 6 illustrates the results of the AttrakDiff questionnaire [5] for the three group-binding methods. Pragmatic quality (PQ) refers to the product's ability to support the achievement of behavioral goals (usability). Hedonic quality refers to the users' self: stimulation (HQ-S) is the product's ability to stimulate and enable personal growth, while identification (HQ-I) is the product's ability to address the need of expressing one's self through objects one owns. Perceived attractiveness (ATT) describes a global value of the product based on the quality perception.

We analyzed the AttrakDiff data with the same methodology as the extended NASA-TLX data. The results

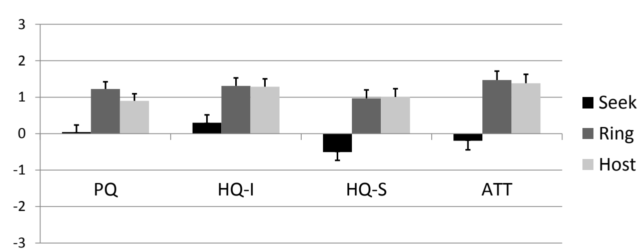

Figure 6. AttrakDiff results.

indicate that the binding method had a significant effect on all dimensions PQ ($F(2, 35.63) = 15.74$, $p < .001$), HQ-I ($F(2, 32.70) = 60.37$, $p < .001$), HQ-S ($F(2, 41.98) = 52.66$, $p < .001$), and ATT ($F(2, 31.52) = 55.79$, $p < .001$). Pairwise comparisons with Bonferroni correction show that the levels of PQ for Ring ($p < .001$) and Host ($p = .002$) were significantly higher compared to Seek, and that the levels of HQ-I, HQ-S, and ATT for both Ring and Host were significantly higher compared to Seek (all $p < .001$). There were no significant differences between Ring and Host on any of the dimensions.

Performance

Seek was the most reliable method with all group creation attempts succeeding without moderator assistance. With Ring, two of the six groups failed their initial attempts because several participants started the application simultaneously. With Host, two groups failed their initial attempts because of multiple participants starting the application and two groups because of incorrect touching order. After solving these initial difficulties, all groups were able to successfully create groups with all the methods.

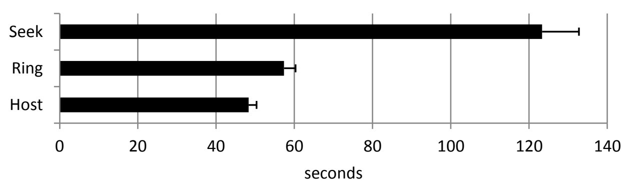

Figure 7. Completion times.

We measured the fastest completion time for the group creation task for each method in each of the sessions from the video recordings and device logs. The participants were instructed to create a group as they would in a real-life situation. If the participants clearly performed the group creation task in a non-optimal way, for example, encountered problems or started to explore different features, the moderator asked them to repeat the task until the process was completed smoothly. Fig. 7 illustrates the mean completion times, with the error bars showing the standard errors. The fastest method was Host, followed by Ring. Seek was clearly slower than the two touch-based methods. While the questionnaire results on perceived quickness show similar order, the distinctions are smaller with no statistically significant differences.

Attractiveness

Half of the participants (12/24) commented that Seek was old-fashioned and boring. *"[P23] Seek was so 90's, engineering style."* Further, many participants (10/24), especially the ones that were less technologically oriented, commented that Seek was far too technical for them. They felt that Seek had too many steps and it was too complicated to use. *"[P16] Seek is too technical. Predictable but not intuitive. Not fun to use."* Compared to Seek, the touch based methods, especially Ring, were considered to be novel, intuitive, and simple to use. *"[P20] I think Ring is very stylish. It is new… I am not a very technical person, but Ring was simple to use and understand what was happening."*

These qualitative findings are supported by the quantitative results. In NASA-TLX (Fig. 5), Seek was rated significantly higher in mental effort and frustration compared to Ring. On the AttrakDiff questionnaire (Fig. 6), both Host and Ring were rated significantly more attractive (ATT) than Seek.

Acting as a Group

During the group creation task, the participants clearly acted together as a group instead of individuals. There was rich interaction between the participants, suggesting and agreeing the next actions and confirming the results. The participants were eager to help each other, if they noticed that some other participant was experiencing problems with the system. This contributed to the high success rate of the group creation tasks. The attention of the participants was divided between their own devices, and the devices and actions of the other participants. In touch-based methods, the touching actions were clearly visible to everybody, making it easier to follow the situation already before the participants' own devices joined the group and started to provide feedback about the system status. In Seek, the participants were forced to check the status by asking verbally or by peeking on the other users' screens.

Leader-Driven vs. Peer-Based Group Creation

One of the main differences between Ring and Host was that in Host, the group-binding process was driven by one person (the leader) who did most of the work, while in Ring, all participants contributed to the group-binding process as equal peers. This had several interesting effects.

Host provided the leader with control over who could join the group. Many participants (13/24) considered that this would be an important feature in some situations. *"[P20] If there were people [around] that I didn't know so well, like at my workplace, Host would be the best [method] because I could control with whom I share."* Some participants (5/24) suggested that the leader should also be able to force participants to leave the group. Further, Host allowed one person to create a group for everybody, so that the others did not have to do anything. Some participants (9/24) commented that it was good that only one person who was able to create a group was required, for example, if some of the participants were less technologically oriented than the others, or if some of the participants were not fully able to use their devices because of some situational factors (for example, because they had children sitting on their knees).

During the evaluation sessions, the participants were very polite towards each other in selecting the leader. However, as commented by one of the participants, selecting the leader might be more challenging in real-life situations, involving complex group dynamics and cultural factors. *"[P12] How can this guy be the leader, if [another person] is the senior? Or if the oldest guy is the leader, he might not know much about technology. Or with youngsters, if there is one who is the leader of the group, how does the group creation go?"* On the other hand, Host was considered as natural in situations where there was a clear leader, for example, in official meetings.

Almost half of the participants (11/24) felt that Ring brought people more together and helped to create a greater sense of community, because everybody was equally involved in creating the group and was forced to interact with the others by touching their devices. *"[P5] Ring makes a spiritual chain between participants. It makes you feel better."* Participants compared Ring *to "[P14] passing the torch"* or *"[P12] shaking hands"*, and commented that it helped to *"[P10] break the ice"* and *"[P9] take down the barriers."* Ring was considered to be particularly suitable for informal situations where there was no strict hierarchy, for example, when meeting a group of friends.

While Host worked well for small groups with all participants located near each other, most participants (16/24) commented that it would not work for larger groups because it would be tiring for the leader to touch a large number of devices, nor longer distances because the leader could not reach all other devices without moving around. Also other factors, for example, having dinnerware on the table, might make it difficult for the leader to touch the other devices. Some participants (10/24) commented that Ring would scale better to larger groups and distances. One participant contrasted the difference between Host and Ring with distributing handouts in meetings. *"[P6] In large meetings, there is no time to give handouts to everybody one at a time. You circulate them."* On the other hand, in a large scattered group, it might be difficult to know who is the last person and should complete the group.

Touching

In Ring and Host, identification of the devices intended to participate in the group was based on touching. Almost half of the participants (11/24) commented that touching was an easy and intuitive way to add participants to the group. *"[P19] Touching to join was a clear, physical, easy, natural way to bring someone into the group."* On the other hand, some participants (9/24) commented that touching could be socially awkward, for example, in formal situations, and brought about privacy issues. *"[P13] Touching is the same as using the other person's phone myself."* In the case of group creation, however, there was a clear reason to touch the other person's phone, so it did not feel like an invasion of privacy. Many participants (10/24) spontaneously pushed their devices forward when another user approached to touch it. This might simply have been a polite gesture to make it easier for the other participant to reach the device, but it could also have indicated giving a permission to touch one's personal device. Finally, some participants (6/24) stressed that to be useful, touching should be detected fast and work very reliably.

Other Gestures

In addition to touching, Host also used gestures for two other purposes. The first gesture allowed the participants to leave the group by flipping their devices upside down. Most participants (19/24) flipped their devices by putting them upside down on the table – only a few flipped their devices in their hands. Some participants (8/24) commented that flipping was a novel, simple, and entertaining way to leave the group. On the other hand, some participants (9/24) raised concerns that it was difficult to know and remember the gesture and it was easy to do it accidentally.

The second gesture enabled the participants to move between Photo Sharing and Discovery Modes by putting their devices on the table and picking them up. Half of the participants (12/24) commented that they did not like this feature because holding the device in their hands was the natural way to use the device and allowed them to control the privacy and viewing angle of their screens and because there might not always be a table available to put the device on. *"[P21] Keeping the device on the table is not something I usually do. I usually hold the device in my hand."*

Ordering

In order to allow throwing of photos between devices, the participants had to define the order of the devices on the table. In Seek, this was done manually by the leader, while in Ring and Host, the participants were expected to touch the devices in counter-clockwise order and the order of the devices was automatically determined based on the touching order. Almost all participants (20/24) considered the requirement to touch the devices in a specific order too restrictive, difficult to remember, and unforgiving to errors. *"[P17] I did not like that you had to go in [counter-]clockwise order. Why not the other way? It should work both ways. It is difficult to remember and learn."* Instead, the participants liked the flexibility and robustness that the manual reordering provided to them. *"[P12] Being able to easily change the order would be the number one feature for me."* The participants pointed out several cases, where manual reordering would be beneficial, for example, if there was a human or technical error in the initial group creation phase, or if the participants moved or changed places. Almost half of the participants (11/24) considered the colored dots, which identified the devices on the screen inadequate, and proposed that textual names should be used in addition to the color.

Perceived Security

In Seek, security was based on six-digit authentication strings that were automatically generated by the system. The participants who wanted to join the group had to manually copy and enter the authentication string into their devices. The participants considered the authentication strings as passwords that they were familiar with in other systems. The dominant way of sharing the password was that the leader read the password aloud. Typically, the password had to be repeated many times as not all the participants were ready to enter it at the same time, or some of the participants missed parts of it. In only one of the six sessions, the participants shared the password by putting the device of the leader at the center of the table, so that everybody could read the password from the screen. However, also in this case some of the participants sitting further away from the leader had difficulties in obtaining the password because they could not clearly see the screen. Most participants (14/24) considered the passwords awkward and would have preferred some other security mechanism. *"[P18] If you need that security level, there must be a better way than [passwords]."* Some participants proposed improvements to the passwords used, for example, making the passwords shorter, using common words, or allowing the participants to define the passwords. Half of the participants (12/24) considered sharing the password verbally as a security risk as anybody in the proximity could hear it. In that sense, the passwords were thought to provide a false sense of security. *"[P14] Password is a complication without any security element."*

In Seek and Ring, security was based on physical proximity enforced by the short range of the touch detection technology. Compared to passwords, which were familiar to all, this was a new concept to the participants. Most participants (13/24) considered that touching provided adequate security for scenarios like sharing photos, provided that the detection technology works reliably and the range is not too long. *"[P22] If phones have to touch, it is quite safe. If somebody I don't know comes so close, I would be alert anyway."* This finding is also supported by the extended NASA-TLX results, which indicate no significant differences in perceived security between Seek and the touch-based methods. Still, some participants (8/24) raised concerns over unauthorized persons accessing their devices by touching, for example, when they had their devices in their pockets in a crowded bar or in a queue.

DISCUSSION

Seek vs. Touch-Based Methods

Both quantitative and qualitative results of the user evaluation show that touch-based methods provide a promising alternative to dominant scanning and password based group-binding methods. While Seek was familiar and reliable in practice, it was considered to be technical, complicated, old-fashioned, and boring. Overall, the participants clearly preferred the touch-based methods and considered them to be simple and intuitive as well as novel and enjoyable to use. The touch-based methods were also faster and they allowed the participants to better maintain awareness of the status of the group formation task as the touching actions could easily be perceived by everybody. Regarding security, touching was considered to be equally secure to passwords. However, to work well in practice, touch detection should be fast and it should work reliably only within the defined distance.

Leader-Driven vs. Peer-Based Methods

The group-binding task can be divided in different ways between the participants. The study results show that different approaches have different strengths and weaknesses. The leader-driven methods, which concentrate the task on a single participant, enable the leader to have strong control over the group and require only one person who is able to create a group. On the other hand, selecting the leader may add more complexity to the group creation process. The peer-based methods, which distribute the work between all participants, help to create a stronger sense of community and scale better to larger numbers of participants and distances. The study results indicate that there is no single optimal method, but the best method depends on the application, social situation, and physical environment. Therefore, the group-binding methods should not strongly enforce a single group creation procedure, but allow for flexibility, so that the participants could adapt the method to the particular needs of each situation.

Device Ordering

The group-binding methods also allowed the determination of the device order using two different approaches: arranging the devices manually or defining the order by touching. The study results indicate that the participants found the requirement to touch the devices in a specific order too restrictive and preferred to touch the devices in a free order and then arrange the devices manually. Again, the optimal touching order depends on social and environmental factors and the group-binding methods should allow the participants to adapt the touching order to each situation. Also, flexible touching order allows the participants to better recover from human and technical errors that may occur during group creation. A well-defined relationship between the touching order and the initial positions of the participants might still be useful for advanced users who want to optimize the group creation process for efficiency.

Supporting Self-Expression and Playfulness

We observed an overall positive mood where participants collaborated and helped each other during group creation. On top of that, we also noticed participants were often laughing, making jokes by creating funny group names, celebrating their collective successes by cheering when they had successfully created a group, and describing the touch-based methods as *"[P8] this is like some Enterprise stuff from Star Trek."* These situations bring to our attention that we are not purely dealing with connecting devices together, but that people are looking for an overall experience that allows them to express themselves and be playful. Therefore, the group-binding methods should look beyond the purely functional task of connecting devices and sharing information, and aim to also engage users on other aspects such as supporting self-expression and playfulness.

CONCLUSION

We have presented a comparative evaluation of two touch-based group-binding methods, a leader-driven method called Host and a peer-based method called Ring, against a more conventional method called Seek, which was based on scanning the available devices in the proximity and passwords for security. The results indicate that the participants strongly preferred the touch-based methods in both pragmatic and hedonic qualities as well as in overall attractiveness. In terms of perceived security, touching was considered equally secure to passwords. While Host allowed better control over the group and required only one participant to be able to form a group, Ring helped to create a greater sense of community and scaled better for larger group sizes and distances. As the optimal group-binding method depends on the social and physical environment, the binding methods should be flexible, allowing the users to adapt them to different contexts of use. For determining the order of the devices, manual arrangement was preferred over defining the order by touching.

ACKNOWLEDGMENTS

We would like to thank Juha Riippi, Iiro Vidberg, Arttu Pulli, Markus Rinne, and Mikko Tolonen for implementing the software prototypes. We would also like to thank Johan Kildal for helping with NASA-TLX, and Susan Fussell for helping with the statistical analysis.

REFERENCES

1. Ayatsuka, Y. and Rekimoto, J. tranSticks: physically manipulatable virtual connections. In *Proc. CHI '05*, 251-260.
2. Chong, M. and Gellersen, H. How users associate wireless devices. In *Proc. CHI '11*, 1909-1918.
3. Chong, M. and Gellersen, H. Usability classification for spontaneous device association. *Personal and Ubiquitous Computing*, published online (2011).
4. Hart, S. G. and Staveland, L. E. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In Hancock and Meshkati (eds.) *Human Mental Workload*, North Holland Press, 1988.
5. Hassenzahl, M. The Interplay of Beauty, Goodness, and Usability in Interactive Products. *Human-Computer Interaction* 19, 4 (2004), 319-349.
6. Hinckley, K. Synchronous gestures for multiple persons and computers. In *Proc. UIST '03*, 149-158.
7. Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P. and Smith, M. Stitching: pen gestures that span multiple displays. In *Proc. AVI '04*, 23-31
8. Holmquist, L.E., Mattern, F., Schiele, B., Alahuhta, P., Beigl, M. and Gellersen, H. Smart-Its friends: a technique for users to easily establish connections between smart artefacts. In *Proc. UbiComp '01*, 116-122
9. Holtzblatt, K., Wendell, J. B. and Wood, S. *Rapid Contextual Design*. Morgan Kaufmann, 2004.
10. Ion, I., Langheinrich, M., Kumaraguru, P. and Čapkun, S. Influence of user perception, security needs, and social factors on device pairing method choices. In *Proc. SOUPS '10*.
11. Kainda, R., Flechais, I. and Roscoe, A. Two heads are better than one: security and usability of device associations in group scenarios. In *Proc. SOUPS '10*.
12. Kray, C., Rohs, M., Hook, J. and Kratz, S. Group coordination and negotiation through spatial proximity regions around mobile devices on augmented tabletops. In *Proc. TABLETOP 2008*, 1-8.
13. Lucero, A., Holopainen, J. and Jokela, T. Pass-Them-Around: collaborative use of mobile phones for photo sharing. In *Proc. CHI '11*, 1787-1796.
14. Lucero, A., Jokela, T., Palin, A., Aaltonen, V. and Nikara, J. EasyGroups: binding mobile devices for collaborative interactions. In *CHI EA '12*, 2189-2194.
15. Mayrhofer, R. and Welch, M. A human-verifiable authentication protocol using visible laser light. In *Proc. ARES '07*, 1143-1148.
16. McCune, J.M., Perrig, A. and Reiter, M.K. Seeing-is-believing: using camera phones for human-verifiable authentication. In *Proc. SOUPS '05*, 110-124.
17. Nithyanand, R., Saxena, N., Tsudik, G. and Uzun, E. Groupthink: usability of secure group association for wireless devices. In *Proc. Ubicomp '10*, 331-340.
18. Rashid, U. and Quigley, A. Interaction techniques for binding smartphones: a desirability evaluation. In *Proc. HCD '09*, 120-128.
19. Rekimoto, J. SyncTap: synchronous user operation for spontaneous network connection. *Personal and Ubiquitous Computing* 8, 2 (2004), 126-134.
20. Rekimoto, J., Ayatsuka, Y., Kohno, M. and Oba, H. Proximal interactions: a direct manipulation technique for wireless networking. In *Proc. INTERACT '03*, 511-518.
21. Terrenghi, L., Quigley, A. and Dix, A. A taxonomy for and analysis of multi-person-display ecosystems. *Personal and Ubiquitous Computing* 13, 8 (2009), 583-598.
22. Uzun, E., Saxena, N. and Kumar, A. Pairing devices for social interactions: a comparative usability evaluation. In *Proc. CHI '11*, 2315-2324.
23. Zimmerman, T. G. Personal area networks: near-field intrabody communication. IBM Systems Journal 35, 3-4 (1996), 609-617.