FlexiGroups: Binding Mobile Devices for Collaborative Interactions in Medium-Sized Groups with Device Touch

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

ABSTRACT
We present a touch-based method for binding mobile devices for collaborative interactions in a group of collocated users. The method is highly flexible, enabling a broad range of different group formation strategies. We report an evaluation of the method in medium-sized groups of six users. When forming a group, the participants primarily followed viral patterns where they opportunistically added other participants to the group without advance planning. The participants also suggested a number of more systematic patterns, which required the group to agree on a common strategy but then provided a clear procedure to follow. The flexibility of the method allowed the participants to adapt it to the changing needs of the situation and to recover from errors and technical problems. Overall, device binding in medium-sized groups was found to be a highly collaborative group activity and the binding methods should pay special attention to supporting groupwork and social interactions.

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
While mobile devices have traditionally been very personal devices targeted at individual use, over the last years there has been a growing interest in systems that combine several mobile devices together to create broader ecosystems of interaction [24]. Such ecosystems allow groups of collocated users to engage in rich collaborative activities and shared experiences with their devices. Potential application scenarios include presenting and collaboratively editing documents in business meetings, sharing photographs and videos within groups of friends in a café, or playing multi-player games with other family members in the living room.

But before a group of collocated users can engage in collaborative interactions with their mobile devices, they must first join their devices together into a multi-device ecosystem. This is a complex procedure with several steps. The necessary system and application software needs to be initiated on all devices. The devices must discover 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. Typically, short-range radio technologies, such as WLAN or Bluetooth, are used to transmit data between devices. This process of setting up the ecosystem is generally known as device binding or ecosystem binding [24] (also known as device association, pairing, or coupling [4]). 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 users cannot see the wireless connections between the devices, they cannot be sure that they are really connecting to the other devices intended to. Therefore, the binding process should also provide sufficient cues and security, so that the users can ensure that the right devices are connected.

In this paper, we study establishing an ecosystem of mobile devices to support collaborative interactions within medium-sized groups of collocated users. While the problem of a single user pairing two devices has been extensively studied in prior research, more complex scenarios involving multiple users, especially more than four users, have received little attention in prior research. We present a touch-based group-binding method called FlexiGroups that builds on earlier research by Jokela and Lucero [13] and Lucero, et al. [16]. The method is highly flexible, enabling the users to apply a broad range of different group formation strategies. We also present a laboratory evaluation of FlexiGroups in a realistic photo sharing application context with four groups of six users. The evaluation results indicate that the method was generally found easy and intuitive to use. We analyze the different group creation strategies and patterns used by the

MobileHCI 2014, September 23–26, 2014, Toronto, ON, Canada.

http://dx.doi.org/10.1145/2628363.2628376

participants during the evaluation. The participants primarily followed viral patterns where they opportunistically added other participants to the group without advance planning. The participants also suggested a number of more systematic patterns, which required the group to agree on a common strategy but then provided a clear procedure to follow. The flexibility of the method allowed the participants to adapt it to the changing needs of the situation and to recover from errors and technical problems. Overall, device binding in medium-sized groups was found to be a highly collaborative group activity and the binding methods should pay special attention to supporting groupwork and social interactions.

The rest of this paper is structured as follows. First, we provide an overview of the related work. We then give a detailed description of the FlexiGroups binding method and the evaluation procedure. Finally, we present the results of the evaluation, followed by discussion and conclusions.

RELATED WORK

The problem of ad hoc device binding has been thoroughly 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 [20]. These methods vary in terms of device hardware requirements, amount of user involvement, and level of provided security.

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 presenting a list of the found devices to the user for selecting the other device to bind with. The connections are authenticated using short strings that the user is expected to copy or compare between devices. The authentication strings can be represented as numbers, words, graphical images, audio signals, or gestures 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 [22] or touching both devices [26] simultaneously, shaking the devices together [10], or bumping the devices together [8]. Further, device binding can be based on continuous gestures spanning from one device display to another [9]. Methods based on spatial alignment of the devices include pointing, for example, with laser light [18], touching [23], or placing the devices in close proximity of each other [15]. It is also possible to bind devices with various auxiliary devices such as tokens [1] or cameras [19]. Some of the proposed methods cover only device identification or authentication, while others combine both identification and authentication into a single user action.

Binding methods are not only means for connecting devices – they have strong social and emotional aspects. In real-life situations, people do not always pick 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 [12, 21]. Users are willing to take security risks to comply with social norms [12].

The majority of earlier research has focused on scenarios of a single user binding two devices with each other (for example, binding a headset with a mobile device, or a mobile computer with a wireless access point). 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 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 [14]. 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 [25, 5].

While numerous methods and technologies have been proposed for group association, Chong and Gellersen [2] present an interesting study on what people would spontaneously do to associate a group of devices. In their study, groups of four users were asked to suggest and rate techniques for binding together different combinations of low-fidelity acrylic prototypes of various mobile and fixed devices. Device touch based methods were found to be among the most frequently proposed methods, and were also considered popular and easy to use.

The group creation task can be divided in different ways between the members of the group [13, 14]. Leader-driven methods, which concentrate the task on a single participant, allow strong control over the group and require only one participant to be able to form a group. Peer-based methods, which distribute the work between all members, help to create a stronger sense of community and scale better to larger group sizes and distances. Further, group association can be seen as a one-step procedure of binding all devices with a single action, or as a sequence of pairwise associations [2, 24].

Finally, Chong and Gellersen [4] present a framework that sums up and categorizes the different factors that influence the usability of spontaneous device association. They identify technology, user interaction, and application context as the three most important criteria.

STUDY

Objectives

In this study, we were primarily interested in two research questions:

First, earlier studies have suggested that device-binding methods should be flexible, allowing people to adopt different group creation strategies in different situations [13, 12]. Still, many of the methods tested in earlier studies have been very specific, enforcing a detailed procedure that has to be followed exactly. We wanted to test a more flexible method that would give users more freedom to adapt it to different situations. On the other hand, giving the users more options might potentially be confusing to them and provide additional possibilities for errors.

Second, we wanted to study group binding in a medium-sized group of six users. Earlier studies on device binding have focused either on individual users, pairs, or small groups of up to four participants. As the size of the group increases, a much wider variety of different approaches and strategies becomes possible. The overall process also becomes much more parallel. We wanted to better understand the different possible approaches and their strengths and weaknesses. We were also interested in group behavior, communication, and collaboration between users during the group-formation task.

FlexiGroups

In order to study these research questions, we designed a group-binding method called FlexiGroups. The method builds on the results of an earlier comparison of three group-binding methods by Jokela and Lucero [13]. While the FlexiGroups method itself is generic and can be used in many different applications, we decided to study the group-binding method in the context of a multi-user photo sharing application to create a more realistic setting for the evaluation. The application was a simplified version of Pass-Them-Around [16]. It allowed the users to browse their own photo collections stored in their devices and also supported spatial interactions of sharing photos by throwing them from one device to another.

Figure 1 illustrates the FlexiGroups group-binding method. To create a new group, one of the persons first starts the FlexiGroups application on their device by tapping the application icon in the Application Grid (Fig. 1.a). The application starts in the Add Device view (Fig. 1.b). Visual feedback on the screen instructs the person to hold the device in portrait mode and to touch another device to add it to the group. When the person moves their device close to another device, the device detects the new device (see Fig. 2). Visual, auditory, and haptic feedback is provided to indicate that the other device has been detected and to instruct the person to hold the device still while the new device is added to the group. When the new device has been added to the group, the person can continue adding more devices by touching them following the same procedure. When the person does not want to add any more persons to the group, they should press the “Done” button to enter the Tabletop Overview view (Fig. 1.d).

Note that the person who is added to the group does not have to start the application manually from the Application Grid – their device can remain idle (Fig. 1.c). When another

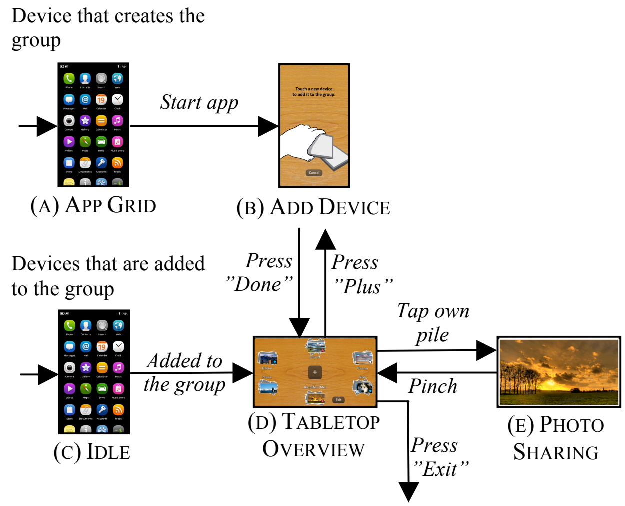

Figure 1. FlexiGroups group-binding method.

person touches the device and adds it to the group, the application is automatically launched and starts in the Tabletop Overview (Fig. 1.d). The Tabletop Overview (see Fig. 3) shows all the persons who have been added to the group. Each person is represented as a pile of photos with a textual name defined by the person next to the pile. If the person wants to add more persons to the group, they can press the “Plus” button at the center of the screen to enter the Add Device view (Fig. 1.b). The person can then add new devices by touching them in the same way as described in the previous paragraph. Any member of the group is allowed to add new devices and several persons can add new devices simultaneously.

FlexiGroups also supports defining the positions of the devices relative to each other in order to enable spatial interactions such as throwing photos between devices. This ordering phase is optional and can be omitted in applications that do not require the order of the devices to be defined. Alternatively, a similar mechanism can also be used to define other kinds of roles within the group, for example, to divide the group into two competing teams in a game application and to select captains for both teams. The ordering mechanism in FlexiGroups works as follows. The persons appear in the Tabletop Overview (see Fig. 3) in the order they are added to the group. If the order of the devices on the screen is different from the order of the devices in the real world, any member of the group can correct it by dragging the devices to the right positions on the screen. Only one person can change the order of the devices at a time. When one person is dragging a device to a new position, the other devices’ screens are locked and grayed out to indicate that another person is reordering the devices.

When the order of the devices is correct on the screen, the people can enter the Photo Sharing view (Fig. 1.e) by tapping their piles of photos on the Tabletop Overview and start sharing photos by throwing them between devices. By pinching to zoom out in the Photo Sharing view, the person

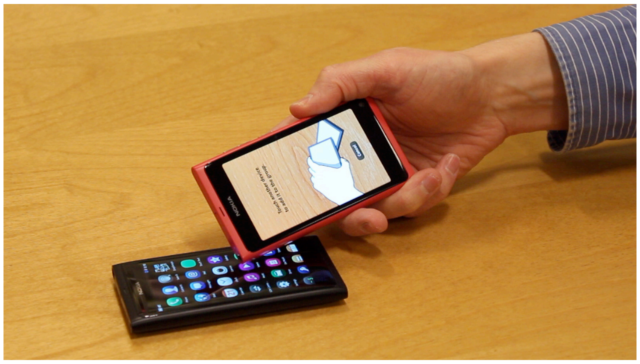

Figure 2. Touching another device to add it to the group.

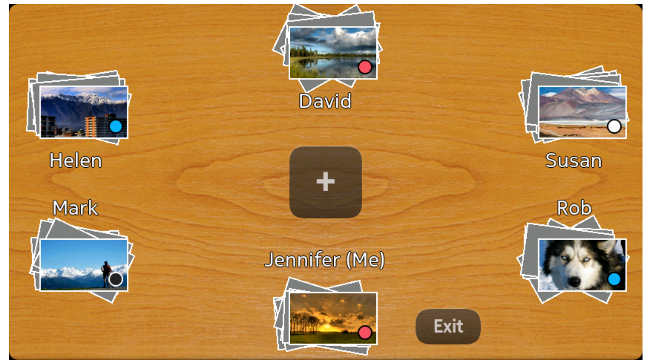

Figure 3. Tabletop Overview.

can return to the Tabletop Overview at any time to check the current members of the group, to add new members to the group, or to change the order of the devices. To leave the group, the person presses the "Exit" button in the Tabletop Overview. The group continues to run on the other devices until the last member exits the group.

Prototype Implementation

We implemented a prototype of the FlexiGroups binding method on Nokia N9[1] mobile devices running the MeeGo operating system. The prototype was built as a native C++ application on top of the Qt 4.7 software framework. QML and Qt Quick with OpenGL ES hardware acceleration were used to implement a smooth animated user interface. Device touch interactions were detected with a radio technology, which was able to detect other devices at distances closer than 20 cm (8 inches) in approximately five seconds based on wireless signal strength. While the technology generally worked reliably, there were occasionally delays in detecting the other devices in the proximity and detections of devices further away. Detailed connectivity and initialization information was then sent to the discovered device over Bluetooth. A server, which was listening to a pre-defined Bluetooth socket on the discovered device, received the information and started the FlexiGroups application. The application then established an ad hoc WLAN network and connected to the group according to the connectivity information it had received. An ad hoc WLAN network was used for communication between the devices in order to allow the application to be used anywhere independent of the available network infrastructure. All communication was handled directly between the devices without a server backend. The prototype was fully functional with real network communication, except for the security protocols, which were only simulated in the user interface.

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

Participants

We recruited a total of 24 participants for the evaluation by posting an advertisement on a local mailing list. The participants were recruited in groups of two to four people, as we wanted that each participant would know some of the other participants to make the situation more natural and comfortable to them. We assigned the participants into four evaluation sessions of six people each in the order they registered for the study, so that every participant knew at least one other participant, but no participant knew all the others in the same session. Six of the participants were female and 18 male. The ages of the participants varied between 25 and 41 years ($M=32.8$, $SD=4.7$). One participant was left-handed and 23 right-handed. The participants represented a variety of different professions, with six participants having a software engineering background, 13 having other technical background (for example, mechanical engineering), and five having a non-technical background (for example, teaching or photography). The participants were fairly advanced users of information technology – on a scale between 1-7 (1=novice, 7=expert), they evaluated their IT skills above average ($M=5.5$, $SD=1.2$). All participants were experienced smartphone users, but only two of them had used a Nokia N9 device before the study.

Procedure

We organized a total of four evaluation sessions. The sessions were arranged in our usability laboratory. In each session, there were six participants and a moderator present. Figure 4 shows the evaluation setup. The participants were sitting around a round table with a radius of 120 cm (48 inches). Each participant was provided with a Nokia N9 mobile device with the FlexiGroups application pre-installed. While all the devices were of the same model, we used devices of different colors to make them easier to differentiate: there was one black, one white, two cyan, and two magenta devices. The average duration of the evaluation sessions was 60 minutes.

As the participants arrived in the laboratory, the moderator guided them to their seats around the table and asked them to fill in the background information and consent forms.

The moderator then introduced the participants to collaborative multi-device applications and demonstrated the idea with the photo sharing application. The participants could try the application with their own devices and practice throwing photos to each other. The moderator explained that before the participants could share photos like this, they first had to bind their devices together into a group, and that the purpose of the evaluation was to test a method for this task. The moderator then continued by describing the detailed evaluation procedure. The moderator also told the participants that the method to be tested was based on device touch interactions and demonstrated how to touch another device. The participants could practice touching with their own devices until they felt comfortable doing it. We considered practicing touch interactions necessary as while many participants were aware of device touch as an interaction method, few had practical experience using it.

To start the actual evaluation, the moderator played a video that demonstrated the FlexiGroups group-binding method. We used instructions recorded on video to ensure that all groups received uniform guidance on how to use the method. The instructions demonstrated the operations that an individual user could do – they did not show how a group of users should use the method together. The moderator then asked the participants to set up a group using the method that was just demonstrated to them, so that they could start sharing photos between devices. The moderator observed the situation and only intervened if the participants encountered obvious technical problems with their devices that prevented them from proceeding. When the participants had successfully created a group and could throw photos between devices, the moderator asked the participants to exit the application and create a new group so that a different person would start the group creation. Overall, the participants created three or four groups during each evaluation session.

When the participants had tested the method several times, the moderator asked them to fill in two validated questionnaires: AttrakDiff [7], which measures the attractiveness of interactive products, and NASA-TLX [6], which measures the subjective workload experience when performing a task. To gain a broader understanding of the tested method, we extended the NASA-TLX with four additional scales: Learnability, Quickness, Security, and Overall Preference. After the participants had completed the questionnaires, the moderator interviewed them about their experiences with the FlexiGroups method. The interview was semi-structured and covered general feedback about the strengths and weaknesses of the tested method as well as specific topics such as different group formation strategies, device touching orders, and arrangement of the devices into the correct positions. The average duration of the interviews was 20 minutes. To close the evaluation session, the moderator thanked the participants and gave each participant a small reward to compensate them for their time.

Figure 4. Evaluation setup.

The group creation tasks were recorded with two video cameras: the first was placed on a tripod pointing towards the table at an angle (Fig. 4) and the other was mounted in the ceiling providing a top view of the table surface (Fig. 8 and Fig. 9). The interviews were recorded with a single video camera. Two researchers independently watched the video recordings and wrote notes about their observations. They also drew diagrams that recorded the sequences of participant actions in every group creation attempt. The same two researchers then collaboratively analyzed the data and built an Affinity Diagram [11] in a series of interpretation sessions. Each researcher independently 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. A quantitative analysis of the AttrakDiff and NASA-TLX questionnaires was done separately.

RESULTS

General

The evaluation produced overall positive results. All four evaluation groups succeeded in all of their attempts to bind their devices together and form a group. While the groups encountered some problems and made some mistakes, especially in their initial attempts to use FlexiGroups, the robustness of the method allowed them to recover and continue, and to successfully complete the group creation task. Most participants (18/24) commented that FlexiGroups was generally easy and intuitive to use and it was also easy to learn.

These qualitative results are supported by the AttrakDiff questionnaire [7] results, which are illustrated in Figure 5. The main bars indicate the means for each product dimension, while the error bars indicate standard errors. Pragmatic quality (PQ) refers to the product's ability to

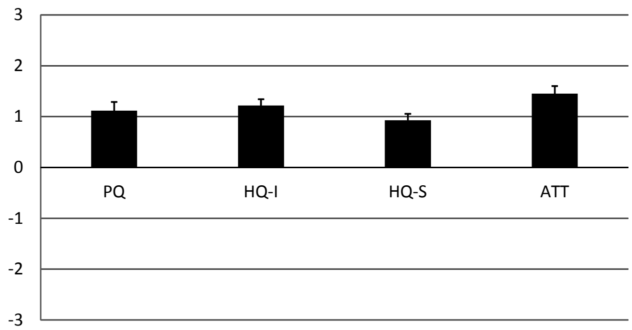

Figure 5. AttrakDiff results (higher is better).

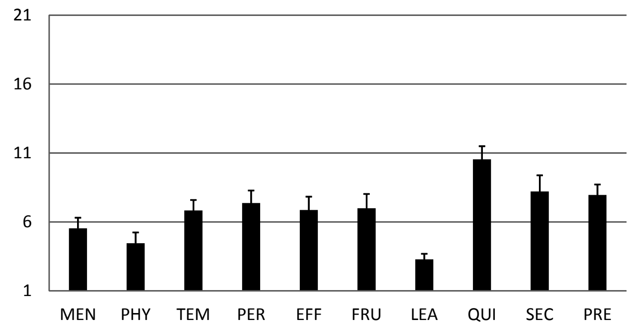

Figure 6. Extended NASA-TLX results (lower is better).

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. Both pragmatic and hedonic dimensions of AttrakDiff as well as overall attractiveness show positive and well-balanced results.

The extended NASA-TLX questionnaire [6] results are illustrated in Figure 6. The main bars indicate the means for each subscale, while the error bars indicate standard errors. The original six subscales of NASA-TLX presented on the left show an overall positive trend. The four subscales that we added for the purposes of this study are shown on the right. The low Learnability (LEA) score indicates that the method was considered very easy to learn. The method was also considered relatively secure (SEC). The Overall Preference (PRE) score further confirms the generally positive attitude towards the tested method.

However, several participants (7/24) commented that the group formation should have been faster. *"[P14] I think it is a bit slow process. It should be somehow faster. It takes a long time to set up."* The relatively high subjective Quickness (QUI) score of the extended NASA-TLX results supports these qualitative comments. We measured the fastest group creation time for each group using the video recordings. The average time to set up a six-person group was 111 seconds.

We can identify three main phases in forming a group: initiating the group creation, adding the participants to the group, and arranging the participants in the correct order for spatial interactions. We will next discuss each of these phases in more detail.

Initiating the Group

Every time the FlexiGroups application is started from the application grid, a new group is created. Therefore, when the participants wanted to create a new group, they had to agree who of them would start the application, while the others had to wait until they were added to the group. Initially, this proved to be challenging for the participants as they were used to the common practice of each person first starting the application on their device. *"[P10] I found it confusing that by just starting the application, I started my own group."* In every evaluation group, several participants started the application in parallel during their first attempt to form a group. However, the visualization of the devices in the Tabletop Overview allowed the participants to quickly realize that there were several parallel groups. The participants solved the problem by agreeing that those people in the smaller groups should exit the application and be then added to the main group by the other participants. One of the participants saw it as an interesting opportunity that there could be several parallel subgroups within a larger group. *"[P21] If you are in a bar or restaurant and there are a lot of people in the table, you are not going to be talking to everyone. ... For me, it is quite natural that you have different groups."*

As the participants understood that only one of them should start the application, they quickly developed practices to agree verbally, or with gestures, who would start the application. These practices included announcing that one would start, asking for permission to start, and suggesting that another participant should start the application. However, half of the participants (12/24) proposed that it should be possible to start several groups in parallel and then merge the groups together. This would enable building a group from bottom up so that anybody could start and also would make the group creation faster. A few participants (4/24) commented that there should be a security mechanism to confirm that both groups really want to merge. Three participants also expressed more general concerns that somebody could add them to a group that they did not wish to join by touching their devices and suggested there should always be a confirmation before a device joins to a group.

Touching Patterns

FlexiGroups gave the participants a lot of freedom regarding the overall approach on how to form a group and the order in which the individual participants were added to

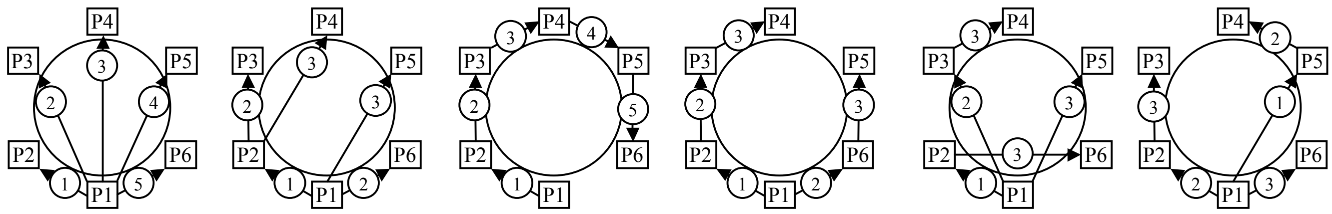

a) LEADER b) CO-LEADERS c) RING d) TWO-WAY RING e) VIRAL f) VIRAL

Figure 7. Different group creation patterns. Boxes P1-P6 represent the participants. Arrows show the touch actions between the participants, while the numbered circles indicate the order of the touch actions.

the group. A single participant could add all participants to the group, or several participants could participate in setting up the group. The participants could be added to the group one at a time, or multiple participants could be added in parallel. During the evaluation tasks and the interviews, the participants suggested a wide variety of different patterns for building a group. Figure 7 illustrates different patterns suggested by the participants.

Systematic Patterns

In systematic patterns, the participants followed a well-defined procedure in setting up a group.

Leader (Fig. 7.a). The participant P1 (the leader) creates a new group and adds all the other participants one after another. The leader may proceed in clockwise or counter-clockwise order around the group.

Co-Leaders (Fig. 7.b). The participant P1 (the leader) creates a new group and then selects another participant P2 as a co-leader and adds the co-leader to the group. The two co-leaders P1 and P2 then add the other participants one after another proceeding in opposite directions around the group. While similar to the Leader pattern, having several co-leaders scales better to larger groups.

Ring (Fig. 7.c). The participant P1 creates a new group and adds another participant P2 next to them. The participant P2 then adds the next participant P3, who continues by adding P4. This way the ring proceeds around the group. The ring can be formed in clockwise or counter-clockwise order.

Two-Way Ring (Fig. 7.d). The participant P1 creates a new group and adds another participant P2 next to them. P1 and P2 then continue by adding the participants next to them like in Ring but in opposite directions. The Two-Way Ring pattern is similar to Ring, but more efficient as it proceeds in both directions around the group.

Viral Pattern

While the participants suggested many systematic patterns, in practice their behavior was usually more random. *"[P14] Somebody initiated the group creation and then it started to spread around the table."* The participants opportunistically selected which device to touch next. *"[P3] It was very random. ... Going left and right. It didn't have any strict form."* This way the group membership spread like a viral infection, or fire, across the group from one participant to another. Figures 7.e and 7.f show examples of real viral patterns employed by the participants during the evaluation.

It was common to connect to one's own neighbors first and then add other participants further away if needed. Participants who were not yet members of the group also requested the participants who were already in the group to add them to the group. If some participant had problems in detecting another device, the other participants were eager to help and add the new device with their own devices. Sometimes the participants intentionally touched and "infected" another participant on the opposite side of the table (for example, touch action #1 in Fig. 7.f) to make the group membership spread faster. It was common that several participants were touching and adding new devices simultaneously (see Fig. 8).

Alternative Patterns

In addition to the main patterns described above, individual participants suggested a range of alternative approaches for group creation. Interesting alternatives include patterns where several devices were touched simultaneously to make the group creation more efficient. In one variation, the participant who creates the group puts their device at the center of the table. The other participants who want to be added to the group then put their devices next to it. In another variation, all devices are collected next to each other on the table and the participant who creates the group then touches all of them in one action. While these patterns

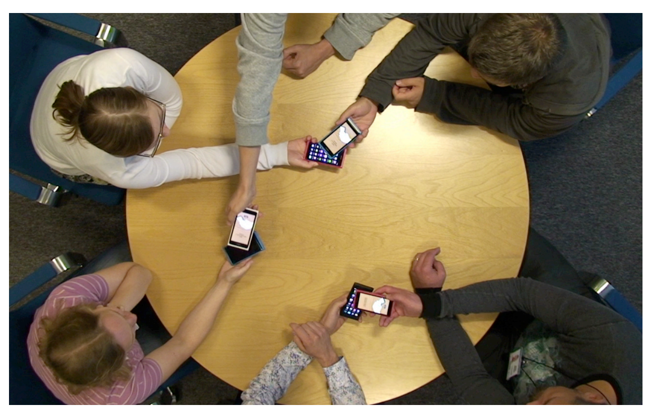

Figure 8. Three persons touching and adding devices in parallel.

would have been possible with the tested prototype, they were not used in practice by any group during the evaluation.

Device Ordering

While the phase of touching and adding devices to the group was characterized by highly parallel activity with several participants adding other people simultaneously (see Fig. 8), acting either independently or in small sub-teams, the phase of device ordering required the participants to more closely co-operate and coordinate their actions within the whole group. The transition from the adding phase to the ordering phase provided an important synchronization point in the group formation task.

Device ordering was done collaboratively with different participants taking different roles (see Fig. 9). While some participants actually moved devices on the screen, others checked whether the order was correct and provided suggestions on the needed changes verbally and with hand gestures. In most cases, several participants moved devices on the screen – typically two to four participants were involved. This created problems as several participants tried to move the devices simultaneously. While the screens of the other devices were grayed out and locked when one of the participants started to move a device around, there was a small delay in locking the screens, which sometimes allowed another participant to start moving until their screen was grayed out. Many participants (13/24) complained that they felt the device ordering phase was confusing. Several participants (7/24) also commented that the ordering phased was unnecessary and should be avoided. The system should have been able to define the device positions automatically instead.

Physical Device Handling

Half of the participants (12/24) made a natural and spontaneous gesture of pushing their devices forward when they wanted to be added to the group. Most commonly, a participant would move their device towards another participant asking them to touch the device and add it to the group. Alternatively, a participant could push their device towards the center of the table asking any of the other participants to add it to the group.

While the participants commonly touched each other's devices with their own devices to add them to the group, there was a high barrier of taking another participant's device into one's hands or manipulating it with one's fingers. This was true even if the participants did not use their own personal devices but devices we had given to them for the purpose of the evaluation. In all the sessions, there were only a few cases where a participant touched another participant's device with their hands. In those few cases, the reason was usually to help another participant who had technical problems with their device.

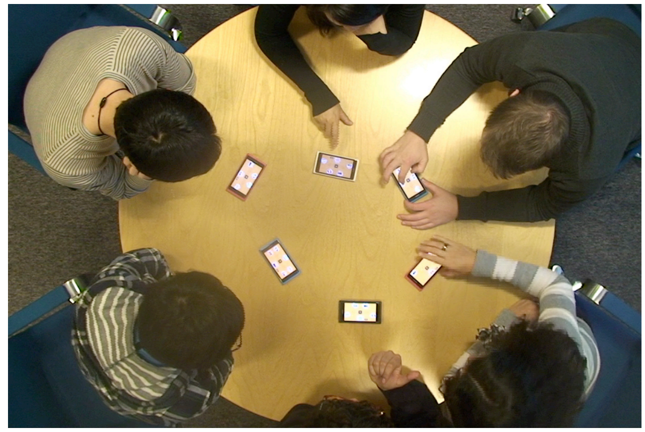

Figure 9. Ordering devices.

We encouraged the participants to keep their devices on the table as that would allow everybody to see the screens of the other participants and help create a common awareness of the group status. Keeping the devices on the table also made the proximity detection technology work faster and more reliably. Still, the participants often chose to hold their devices in their hands. In 27% of the touch actions recorded on video, both participants kept the devices in their hands, making the touch action resemble a handshake (see Fig. 8). In one of the sessions, all participants held the devices in their hands also during the reordering phase. In the other sessions, two participants liked to fiddle or toy with their devices while waiting to be added to the group.

Collaboration

We observed a high level of communication and collaboration between the participants when they were forming a group. Participants were very eager to help each other if some participant encountered problems with the application. For example, if a participant's device could not detect and add another device to the group, another participant would use their device to help and add the participant. Helping others mostly occurred spontaneously when a participant noticed that another participant could not complete some action or was doing something incorrectly. Only rarely did the participants explicitly ask for help.

A major challenge for the evaluation groups was creating and maintaining a common understanding of the overall task status as there were six persons involved and many actions were taking place in parallel. *"[P23] I cannot keep an eye on what every other person is doing."* Several techniques were used to accomplish common awareness of the task status. Verbal communication and coordination between the participants played a major role, including the participants announcing intentions to do some actions, providing feedback on other participants' actions, instructing others to take some actions, asking and confirming facts, and stating the common group status aloud. Another important technique was to observe the other participants, including both real-world actions taken by the others as well as the status information shown on the their device screens. When adding participants to the group, everybody could easily perceive the touch actions. The devices also provided audio feedback when another device was detected and added to the group. While the audio feedback was useful and could also be observed by all the participants, the participants had problems in identifying the source of the audio when several participants were touching devices in parallel. In the ordering phase, the participants had to rely more on information on the device screens, and peeking at and comparing information between other participants' screens was common. Figure 9 shows an example of a situation where all participants have placed their devices in a close formation at the center of the table to make the coordination easier.

DISCUSSION

As observed in our study and in earlier studies [13, 14], device association in large groups is a highly collaborative group activity. Considering the design of binding methods for groups, while good usability is definitely important, people as a group can help each other and are capable of overcoming and solving together most usability and technical problems they encounter. However, in larger groups, the main challenges are related to groupwork and social interactions within the group: making decisions and agreeing on a common strategy, coordinating and synchronizing actions, and keeping track of the others and the overall task status. An important consideration is also keeping everybody engaged in the process, as people easily get bored or distracted when they cannot do anything but wait for others to complete the group formation.

In the evaluation sessions, people most commonly and naturally employed Viral patterns (see Fig. 7.e and 7.f) where the group membership spread like a contagion from one person to another [24]. These required the least advance planning within the group and provided most flexibility. The Viral patterns are efficient and keep everybody involved also in large groups due to the high level of parallel activity. The possibility of forming a group bottom-up by first forming smaller groups and then merging them together into a larger group, which was suggested by some of the study participants, is interesting and should be studied further.

The participants also suggested a variety of more systematic patterns. These patterns require that the participants are aware of and agree on a common strategy to form a group. While this requires some initial planning, the systematic pattern then defines a precise sequence of actions that each participant should follow, reducing the need for coordination between the participants later. Further, the Leader pattern (see Fig. 7.a) enables the leader to have strong control over the group [13], which may be important, for example, in situations where there are unknown people present. The Co-Leaders pattern (see Fig. 7.b) provides an extension of the Leader pattern, which scales better to larger groups.

While the participants were able to successfully arrange the devices in the right order for spatial interactions in all group creation attempts, the majority of them found the ordering phase confusing. While this can be partially attributed to implementation issues, such as delays in locking the screens of the other devices when one device was used for rearranging, the ordering technique needs to be improved. Ideally, device ordering should be automatic, removing the need for manual ordering completely. However, in real-life applications this may be difficult to achieve as it may require special tracking equipment or dedicated infrastructure that may not be widely available. As suggested in an earlier study by Jokela and Lucero [13], device ordering would probably work better when done by one person. This person could be, for example, the person who creates the group. Alternatively, one of the participants could reserve the role for a longer period of time and arrange all the people in the right places. If strictly followed, the systematic patterns might enable defining the device order based on the touch order, but this may be in many cases too restrictive [13].

When designing binding methods for groups of users, it is important to consider robustness in real-life conditions. While many methods can work well in theory or with mock-ups, in reality, multi-user multi-device applications are complex distributed systems. As multiple devices are involved, there is an increased risk of technical issues: the devices may fail to detect each other, the software may crash, and the network connections may be broken. Also, all persons may not be aware of the procedure they should follow, or they may be unable to do so, for example, because they arrive late or they are occupied with other tasks such as incoming telephone calls. Therefore, the methods should not expect an exact procedure to be followed. The methods should be flexible and robust, allowing people to adapt them to the changing needs of the situation and to recover from failures.

FUTURE WORK

Our experiment, like all the other experiments with group-binding methods we are aware of, was done in a usability laboratory under ideal conditions. Also, the use case was defined by the researchers and given to the participants. In real life, various contextual and situational factors influence the group creation process. Therefore, to gain a deeper understanding of group binding in realistic situations and tasks, we believe it would be important to study group-binding methods also in more realistic settings and over extended periods of time with longitudinal field trials.

CONCLUSION

We have presented FlexiGroups, a touch-based method to bind mobile devices for collaborative interactions within a group of collocated users. The method is highly flexible, enabling a broad range of different group formation strategies. In a laboratory evaluation with four groups of six users, the method was found to be generally intuitive and easy to use and learn. When forming a group, the participants primarily followed viral patterns where they opportunistically added other participants to the group without advance planning. The participants also suggested a number of more systematic patterns, which required the group to agree on a common strategy but then provided a clear procedure to follow. The flexibility of the method allowed the participants to adapt it to the changing needs of the situation and to recover from errors and technical problems. Overall, device binding in medium-sized groups was found to be a highly collaborative group activity and the binding methods should pay special attention to supporting groupwork and social interactions.

ACKNOWLEDGMENTS
We would like to thank Arto Palin, Juha Riippi, Iiro Vidberg, Arttu Pulli, Markus Rinne, and Mikko Tolonen for implementing the software prototype.

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 groups of users associate wireless devices. In *Proc. CHI '13*, 1559-1568.
3. Chong, M. and Gellersen, H. How users associate wireless devices. In *Proc. CHI '11*, 1909-1918.
4. Chong, M. and Gellersen, H. Usability classification for spontaneous device association. *Personal and Ubiquitous Computing* 16, 1 (2012), 77-89.
5. Hang, A., von Zezschwitz, E., De Luca, A. and Hussman, H. Too much information!: user attitudes towards smartphone sharing. In *Proc. NordiCHI '12,* 284-287.
6. 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.
7. Hassenzahl, M. The Interplay of Beauty, Goodness, and Usability in Interactive Products. *Human-Computer Interaction* 19, 4 (2004), 319-349.
8. Hinckley, K. Synchronous gestures for multiple persons and computers. In *Proc. UIST '03*, 149-158.
9. Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P. and Smith, M. Stitching: pen gestures that span multiple displays. In *Proc. AVI '04*, 23-31.
10. 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.
11. Holtzblatt, K., Wendell, J. B. and Wood, S. *Rapid Contextual Design*. Morgan Kaufmann, 2004.
12. 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*.
13. Jokela, T. and Lucero, A. A comparative evaluation of touch-based methods to bind mobile devices for collaborative interactions. In *Proc. CHI '13*, 3355-3364.
14. 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*.
15. 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.
16. Lucero, A., Holopainen, J. and Jokela, T. Pass-Them-Around: collaborative use of mobile phones for photo sharing. In *Proc. CHI '11*, 1787-1796.
17. Lucero, A., Jokela, T., Palin, A., Aaltonen, V. and Nikara, J. EasyGroups: binding mobile devices for collaborative interactions. In *CHI EA '12*, 2189-2194.
18. Mayrhofer, R. and Welch, M. A human-verifiable authentication protocol using visible laser light. In *Proc. ARES '07*, 1143-1148.
19. 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.
20. Nithyanand, R., Saxena, N., Tsudik, G. and Uzun, E. Groupthink: usability of secure group association for wireless devices. In *Proc. Ubicomp '10*, 331-340.
21. Rashid, U. and Quigley, A. Interaction techniques for binding smartphones: a desirability evaluation. In *Proc. HCD '09*, 120-128.
22. Rekimoto, J. SyncTap: synchronous user operation for spontaneous network connection. *Personal and Ubiquitous Computing* 8, 2 (2004), 126-134.
23. Rekimoto, J., Ayatsuka, Y., Kohno, M. and Oba, H. Proximal interactions: a direct manipulation technique for wireless networking. In *Proc. INTERACT '03*, 511-518.
24. 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.
25. Uzun, E., Saxena, N. and Kumar, A. Pairing devices for social interactions: a comparative usability evaluation. In *Proc. CHI '11*, 2315-2324.
26. Zimmerman, T. G. Personal area networks: near-field intrabody communication. *IBM Systems Journal* 35, 3-4 (1996), 609-617.