Gradual Engagement between Digital Devices as a Function of Proximity: From Awareness to Progressive Reveal to Information Transfer

Nicolai Marquardt[1,2], Till Ballendat[1], Sebastian Boring[1], Saul Greenberg[1], Ken Hinckley[2]
[1]University of Calgary, 2500 University Drive NW, Calgary, Alberta, T2N 1N4, Canada
[2]Microsoft Research, One Microsoft Way, Redmond, WA 98052, USA
{nicolai.marquardt, sebastian.boring, saul.greenberg}@ucalgary.ca, till@ballendat.info, kenh@microsoft.com

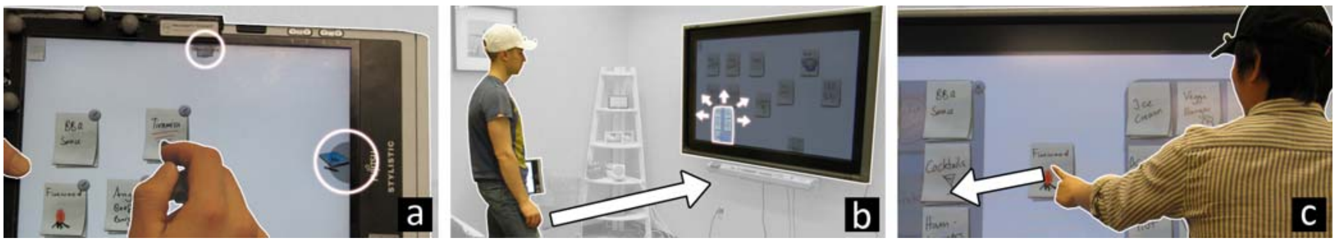

Figure 1. Gradual engagement, showing examples of (a) awareness, (b) progressive reveal, which (c) leads to information transfer

ABSTRACT

Connecting and information transfer between the increasing number of personal and shared digital devices in our environment – phones, tablets, and large surfaces – is tedious. One has to know which devices can communicate, what information they contain, and how information can be exchanged. Inspired by *Proxemic Interactions*, we introduce novel interaction techniques that allow people to naturally connect to and perform cross-device operations. Our techniques are based on the notion of *gradual engagement* between a person's handheld device and the other devices surrounding them as a function of fine-grained measures of proximity. They all provide *awareness* of device presence and connectivity, *progressive reveal* of available digital content, and interaction methods for *transferring* digital content between devices from a distance and from close proximity. They also illustrate how gradual engagement may differ when the other device seen is personal (such as a handheld) vs. semi-public (such as a large display). We illustrate our techniques within two applications that enable gradual engagement leading up to information exchange between digital devices.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces – Input devices and strategies

General terms: Design**,** Human Factors

Keywords: Proxemic interactions, proximity, handhelds, interactive surfaces, awareness, gradual engagement

INTRODUCTION

Personal mobile devices (phones, tablets…) and semi-public stationary devices (information appliances, interactive surfaces…) are an increasingly commonplace way for people to ubiquitously access digital information. Most of these devices are optimized for a seamless user experience when one uses them individually. Yet, using multiple devices in *concert* (such as for transferring information from a mobile phone to the device of a nearby person) is often tedious and requires executing complicated command sequences. This is why several projects in the area of ubiquitous computing (ubicomp) began introducing new techniques to facilitate transfer of content between nearby devices, e.g., [9][11][22]. However, significant challenges remain. People do not know *which* devices they can communicate with, *what* information they contain, and *how* information can be exchanged in a controlled manner. To mitigate these problems, we introduce the concept of *gradual engagement* between a person's handheld device and the other devices surrounding them as a function of proximity. We realize this concept as a suite of novel interaction techniques, all based on providing a seamless transition from awareness leading up to interaction.

People's natural understanding and use of personal space when they interact is called proxemics [8] in social science, and has been applied to ubicomp as proxemic interactions [7][18]. Inspired by this theory, our gradual engagement techniques leverage fine-grained proxemic measurements between people and devices (e.g., distance, orientation, identity). Specifically, engagement increases continuously across three stages as people move and orient their personal device towards other surrounding devices.

Stage 1. *Awareness of device presence and connectivity* is provided, so that a person can understand what other devices are present and whether they can connect with one's own personal device.

Stage 2. *Content awareness* is provided by *progressively revealing a device's available digital content* on other

Marquardt, N., Ballendat, T., Boring, S., Greenberg, S. and Hinckley, K. Gradual Engagement between Digital Devices as a Function of Proximity: From Awareness to Progressive Reveal to Information Transfer. Research Report 2012-1025-08, Dept. Computer Science, University of Calgary, Calgary, AB, Canada T2N 1N4

devices, so that people know what of their content can be accessed on those other devices for eventual exchange.

Stage 3. *Transferring digital content between devices* is provided via various strategies, each tailored to fit naturally within particular situations and contexts: from a distance *vs.* from close proximity; and transfer to a personal device *vs.* a semi-public device.

For example, Figure 1 illustrates one of our techniques for gradually increasing engagement. (a) The person notices device connectivity between the tablet he is holding and the large display by the appearance of icons on that display. (b) Content of his tablet becomes progressively revealed on the large display as the person approaches it. (c) Once in front of the wall display, the person can now directly manipulate the fully visible content of the tablet through direct touch on the wall display, e.g., by dragging and dropping.

Our primary contribution is the notion of using gradual engagement between digital devices as a function of proximity. In particular, we (1) provide novel concepts for using fine-grained proxemics between people and devices to facilitate transition from awareness to interaction; (2) introduce interaction techniques for sharing content appropriate at different proxemic distances and for different device categories (personal *vs.* semi-public); and (3) illustrate the application of these concepts within two example applications.

We first introduce proxemics and proxemic interactions. We then describe our awareness techniques, followed by particular interaction techniques for information transfer between devices (see accompanying video, also at grouplab.cpsc.ucalgary.ca/Publications/, which illustrates this work). We close by outlining our implementation, discussing limitations, and suggesting future work.

BACKGROUND: PROXEMICS FOR UBICOMP INTERACTIONS

Proxemics – as introduced by anthropologist Edward Hall [8] – is one of the seminal theories for describing and studying people's use and understanding of spatial relationships in everyday encounters with others. People often use changes of spatial relationships – such as distance or orientation – as an implicit form of communication. Hall's studies, for example, revealed patterns in how certain physical distances correlate to social distance when people interact: ranging from intimate (6–18"), to personal (1.5–4'), social (4–12'), and public (> 12') settings. Other observations further refined this understanding of people's use of spatiality. For example, spatial features of the environment (e.g., location of walls, doors, furniture) influence people's use of proxemics, and orientation relative to others when we communicate is another driving factor.

Despite the importance of people's use and understanding of proxemics when they interact, these fine-grained spatial relationships are only rarely considered holistically in the design of interactive systems. Inspired by early work of considering proxemics in system design [13][25], the concepts of *proxemic interaction* [1][18] are a first order approximation for applying the insights of proxemic theories to inform ubicomp interaction design. That work describes five important dimensions to consider when designing proxemic-aware ubicomp systems (i.e., defining relationships between people, devices, and objects) [7]: the *distance* and *orientation* between entities, their *movement*, the *identity* of entities, and *location* features (e.g., *fixed features* of an environment, such as the spatial layout of a room, location of doors and furniture, that give further meaning to that setting). For example, a common pattern in proxemic interaction is to interpret decreasing distance and increasing mutual orientation between a person and a device within a bounded space as an indication of gradually increasing interest of that person to interact with that device. In this paper, we advance research in such proxemic interactions even further, by introducing awareness and interaction concepts that consider how proxemics can drive people's interaction with multiple surrounding digital devices and people.

EXAMPLE APPLICATIONS

We use two example applications throughout the paper to illustrate how our various techniques leverage proxemic interaction and gradual engagement to facilitate access to digital information. First, *Proxemic Brainstorming* is an interactive digital brainstorming tool. Its users can create, change, and manage virtual sticky notes on their personal pen-enabled tablets. A large whiteboard provides a public sharing space for notes, and different techniques (explained shortly) allow temporary or permanent transfer of the digital notes between all devices. *Proxemic Photo Canvas* facilitates transfer of digital photos from a network-enabled digital camera to other devices, such as a large display or a digital photo frame.

DEVICE PRESENCE AND CONTENT AWARENESS

We begin with concepts that leverage proxemic information to provide gradual engagement via awareness information about the presence of nearby devices and their location (stage 1), and the content they contain (stage 2).

Stage 1. Awareness of Device Presence & Connectivity

While ubicomp ecologies may contain many devices, only some of them – for a variety of reasons – are likely able to connect with a user's personal device to the point that the person can do useful work between them (such as transferring content). While these devices may sense this information (e.g., via service discovery protocols), the user is often left in the dark about these opportunities for inter-device interaction.

Consequently, we implemented methods that make a person aware of whether his personal device and other nearby devices can detect each other's presence and are able to connect. The basic idea is that a person sees a visual indicator – a subtle notification – about which devices in the surrounding environment can possibly interact with his handheld device, where they can subsequently move toward a particular device to either establish that connection or to reveal further information and interaction possibilities (which would occur in stages 2 & 3, discussed shortly). This is particularly important in dynamically changing or unfamiliar environments: some devices may be hidden or disguised as a non-digital device (e.g., a digital picture frame appliance), or only some of the surrounding devices may allow connections to them (e.g., a device may not support a certain application). Information about these possible connections as well as simple ways to actually establish the connection is crucial if seamless interaction across devices is to occur.

Proxemic relationships. We use rules to determine when to trigger awareness of device presence and connectivity. By connection, we mean whether one device should connect to another device based on human dynamics *vs.* whether a device is technically capable of connecting to another. We exploit the five aforementioned *proxemic dimensions* [7] as sensed factors, where combinations of them allow us to create nuanced rules of connection behaviour.

Location informs devices if they (and the people using them) are in the same room. In almost all cases, devices present in the same room are far more relevant for interaction than the ones in another room. For example, when a person with a tablet enters a new room through the door, notifications can be triggered about other devices available in that particular room. Other devices in close proximity but in adjacent rooms (e.g., behind the walls) are not shown. In proxemics terms, doorways, walls and other boundaries are fixed features that further demark people's sense of social distance; we believe such fixed features are applicable to how devices determine possible candidates for cross-device connections.

Physical distance between devices is an essential factor we exploit for determining device connection and triggering notifications. Proxemic theory states that people naturally stand close to other people they are interested in and want to communicate with. Similarly, we believe that the distance between the user's personal device and other devices in the ecology is a natural indicator of whether a connection between the two should be signaled to the user and subsequently established. Distance measurements can also be applied as a filter that prevents too many notifications in environments with a large number of digital devices. In that case, awareness information is only shown of a limited number of devices that have the smallest distance (e.g., the five closest devices).

Movement – the change of distance over time – are indicators of increasing or decreasing interest. When we are interested in something we move closer to it, while we move away when we are less interested. We can apply this to device-device connectivity. For example, if a person holding a tablet is approaching the large display, we can interpret this as increasing interest of that person to interact with both devices, perhaps in tandem. The devices then highlight cross-device content sharing opportunities over other nearby devices.

Orientation of one device towards another is another indicator that the person wants to connect the two. This again mimics interpersonal interaction: when people interact, they orient themselves to either face the other person or stand side by side. Orientation between devices could simply be whether one device is facing towards or away from another device, or a finer measure that considers and acts upon the angle between the two. For determining cross-device connections, we focus on all devices that are either located in front or at the sides of the device. We assume that if a person wants to interact with a device located behind them, they turn around to face this device, and if they are uninterested, they face away. For example, the visual feedback shown in Fig. 1a+b would appear or fade away as the person turns towards or faces away from the display.

Identity of devices functions as a filter for possible connections. Known devices can trigger the connection notification from a larger distance, while unknown devices need to be located next to each other to establish a successful (and more socially secure) connection. This technique follows the principle that “distance implies distrust” [5], and similarly that closer proximity between devices implies trust. Identity also distinguishes classes of devices, where (for example) connectivity to another person's personal device may be dealt with differently than to a semi-public device, as each suggests different social expectations.

The combination of these five proxemics factors informs the decision about device connectivity, and the corresponding visual/auditory/tactile feedback provided, that eventually allows a user to leverage this knowledge of device presence and connectivity for further interaction.

Notifications about device presence and location. Given the above, a broad variety of notification mechanisms can inform a person about the presence of other nearby devices and opportunities for interaction: audible signals, vibrotactile haptic feedback, visual notifications, etc. Yet given the increasing number of devices in a ubicomp ecology, we opted for a visual approach, as such notifications can be displayed in a more ambient and distinguishable manner. Visuals can portray device identity and location, and – as we will shortly see – can also serve as containers showing content (stage 2) and act as portals for information exchange (stage 3).

In general, all device screens in close proximity display graphical icons representing the location of surrounding connectable devices. Each icon informs the users: where the device represented by the icon is physically located; that there is a potential connection between those devices; and that both devices can interact with each other (e.g., allowing the transfer of information between them). Icon appearance can be informative, such as a graphic that represents the nearby tablet. Or they can be augmented with

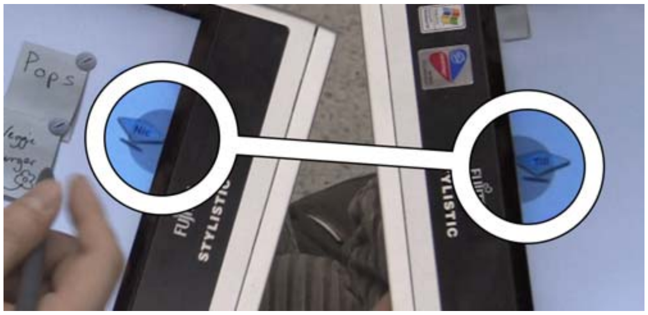

Figure 2. Icons at the edge of the screen indicate the presence of other devices in close proximity (i.e., the two tablet computers).

other information, such as the name of that device and/or its owner. Figure 2 exemplifies this in Proxemic Brainstorming: as the two people move their tablets towards each other, icons at the edge of both screens show the other devices, including the name of the device's owner.

In addition, icon locations are continuously animated around the edge to represent the directional location of the corresponding device. Figure 2 also shows this, where we see how both display's icon locations illustrate their physical spatial relationship. Figure 3 is similar, except it shows how several locations are indicated in a multi-device environment, in this case of two handhelds and a large display. Again, this helps reducing ambiguity of which icon corresponds to which device in the environment.

Because icon location is dynamic, people can further identify the mapping of device icons to actual physical devices by changing their device's distance and orientation and observing icon changes. If multiple devices are shown on a tablet's edge, for example, a person can move and/or rotate the screen and see the icons' positions updated in real-time. Naturally, the same continuous feedback applies when a person is moving closer to a cluster of devices. While approaching those devices, their corresponding icons on the tablet continuously change (e.g., by spreading apart) to reflect the new relationship between the tablet and each device. Thus, a person can move seamlessly towards the particular device desired for interaction.

Stage 2. Content Awareness and Progressive Reveal

As proximity gradually increases, we provide awareness about the available *content* on devices. Knowing what content a device offers for transfer is important information for a person to decide on further interactions. In fact, revealing content available for interaction or transfer to another device creates opportunities that invite a person to discover more about this content, eventually leading to more in-depth interactions.

Revealing content on personal vs. public devices. The information about available content we reveal on the display of other devices differs between *personal* and *public* devices.

For personal devices, we currently only provide an awareness icon of surrounding devices, but not their content. This is due primarily to size constraints: showing

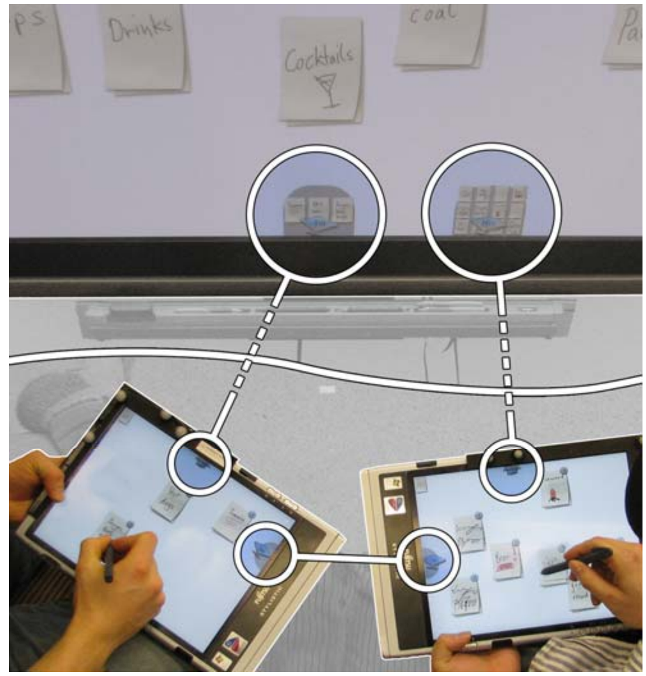

Figure 3. *Content awareness:* Proxy icons indicate the presence of nearby tablets and a large interactive display. Available content of the tablets is displayed as thumbnails atop these icons.

content on the small screens of *personal devices* may interfere with other content the user is viewing or interacting with. As we will see, we use other stage 3 methods to reveal content on personal devices during explicit information exchange.

Public devices (e.g., a wall-mounted display), however, reveals content located on one's *personal devices* people approach the display. For example, the wall display in Figure 3 shows both tablets' awareness icons at its lower edge, where each icon now contains small thumbnail images of all Proxemic Brainstorming notes on the corresponding tablet (i.e., 3 notes on the left tablet, 12 notes on the right one). Even though these thumbnails are too small to allow for full readability, they provide awareness information about the number of notes available for sharing on each of the tablets.

Proximity-dependent progressive reveal. Importantly, revealing content is not all or none. Rather, the *distance* and *orientation* between two devices can directly affect the level of detail of content awareness information shown on other devices. Our *proximity-dependent progressive reveal* technique maps the distance between devices to the amount of information shared between them. The closer two devices are, the more information is shared between them. The level of detail shown (i.e., the amount of information shared) can change either at *discrete* distance levels, or *continuously* with changes in distance. As well, the level of detail can change depending on the orientation between devices. Again, this can happen at discrete angles (e.g., facing to or away from another), or through continuous changes of the orientation (e.g., from 0 to 180 degrees). Progressive reveal is important for two reasons. First, it

Figure 4. *Proximity-dependent progressive reveal*: As a person approaches: (a) the large display first shows the presence of his digital camera, (b) then the last photo taken with the camera is, and (c) a detailed view of multiple photos is revealed, all as a function of proximity.

presents people with opportunities as they approach another device; as with ambient displays, this could mediate the move from background peripheral awareness to foreground interaction. Second, it gives them the chance to pull away, for example, if they see content about to be revealed that they would rather not make public.

The following two example scenarios illustrate both discrete and continuous approaches for progressive reveal of device content when approaching another device.

First, Fig. 4 shows how content of the *Proxemic Photo Canvas* is progressively revealed at *discrete* distance thresholds. Fig. 4a reflects Stage 1: a person holding a digital camera first sees the camera icon on the large display from afar (for illustrative purposes, the icon is shown magnified as an inset in 4a+b). Figs 4b,c reflects Stage 2. When he approaches the wall display and crosses the next distance threshold (here 2m), the most recently taken photo stored in the digital camera is shown next to the camera icon (4b). When he moves directly in front of the wall display while holding the camera close to the screen (~30cm distance), multiple photos on the camera are revealed and shown in a spiral around the camera icon (4c).

Fig. 1 illustrates how Proxemic Brainstorming *continuously* reveals content during Stage 2 – in this case multiple sticky notes located on a person's tablet – as he moves closer to the large display. The wall display shows thumbnails of all sticky notes located on the tablet above the tablet's awareness icon (Fig. 1b). While the actual text on these notes is not yet readable, the number of available notes is already visible. As the person moves closer to the wall display, the thumbnails increase in size continuously until the person stands directly in front of the display. The sticky notes are now shown at full size (1c), allowing the person to read the text of all notes stored on the tablet and to pursue Stage 3 interactions, explained shortly.

Implicit vs. Explicit Reveal. The above methods illustrate how content is revealed via a person's implicit actions. However, reveal can be complemented by explicit methods as well to fine-tune what is revealed. To illustrate, we implemented a combination of implicit and explicit progressive reveal in Proxemic Photo Canvas called *tilt-scrolling* (Fig. 5). During Stage 2, a person now sees a few of his camera's latest photos – organized in a spiral – progressively revealed as an implicit consequence of moving towards the display. To see more content (and while still distant from the display), the person can now explicitly tilt the camera leftwards or rightward to browse through the timeline of photos. Thus, the camera device becomes a physical controller.


Figure 5. *Tilt-to-scroll:* explicit reveal

STAGE 3: PROXEMIC-DEPENDENT TECHNIQUES FOR INFORMATION TRANSFER BETWEEN DEVICES

Stage 1 and 2 indicate device presence, connectivity and available content, eventually leading to Stage 3, where a person can interact with progressively revealed content. We now present a series of interaction techniques that allow for sharing and transferring content between devices.

We stress that the power of these Stage 3 techniques is that they are used in conjunction with the previous Stage 1 and 2 methods *vs.* as stand-alone techniques similar to those found in the literature. Importantly, these techniques consider proxemic relationships between devices to drive the interaction, and come into play at particular points during Stages 1 and 2. We are particularly interested in two contexts:

- whether information exchange is a single person activity (based on the proximity of a handheld to a semi-public display) or a cooperative multi-person activity (based on the proximity of at least two handheld devices).
- how they allow people to interact at different levels of proximity i.e., from a distance *vs.* within reach.

Single Person Transfer: from Personal to Public Device

First, we present a series of techniques that primarily allow a *single person* to share content from their personal device to a public display. We begin with distant-based interactions that could be performed in the early periods of progressive reveal, to within reach interactions at later periods.

Large display drag and back (from a distance) allows a person to temporarily show digital content from their personal device on a large

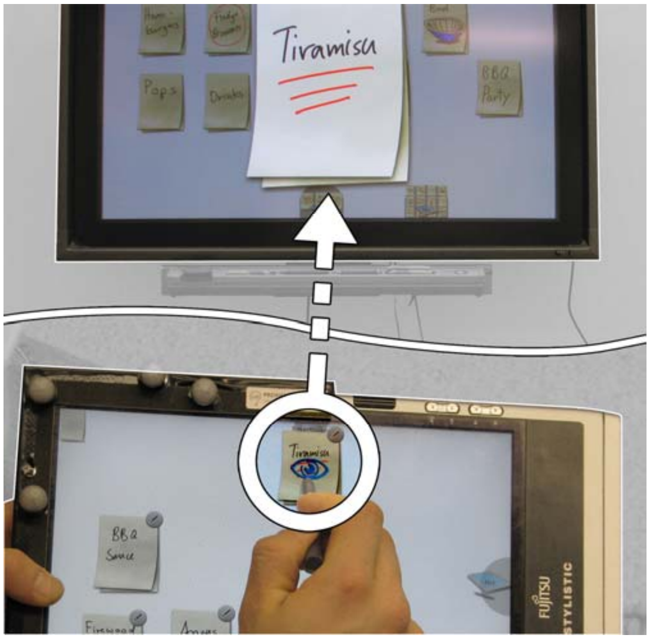

Figure 6. *Large display drag and back*: (bottom) dragging content on the wall display's awareness icon on the tablet, (top) content appears full screen on the large display.

public display. The idea is that the person owns the information, but is making it more convenient for others to view it. To share content temporarily on a large display, a person can drag content onto the awareness icon representing a nearby large screen. For example, Fig. 6 bottom shows a person dragging a note onto that icon. As they do so, a *viewing icon* appears atop the content (here: the 'eye' icon shown inside the circle of Fig. 6) indicating that one is about to share the note on that particular public display. As the person releases the note, the content appears in full screen view on the wall display (Fig. 6 top). To remove shared content, a person simply drags the content back from the device's awareness icon onto the tablet's canvas. Sharing also works for multiple people simultaneously: if others do similar actions, all shared content is shown side by side on the large display.

Point-to-pin (from a distance) lets a person copy content from the camera onto a distant public display by pointing at it and subsequently performing a throwing gesture to pin it there (see Fig. 7). The pointing ray is the extension of the stretched out arm holding the camera. Initially, a preview of the most recently taken photo is shown where that ray meets the large wall display (i.e., an explicit Stage 2 action). 'Throwing' is done by forward-accelerating the hand holding the camera, which permanently copies that photo onto the screen at that location. The technique also works on a digital picture frame, where the photo is then shown full screen in the frame.

Figure 7. (b) *Point-and-pin*

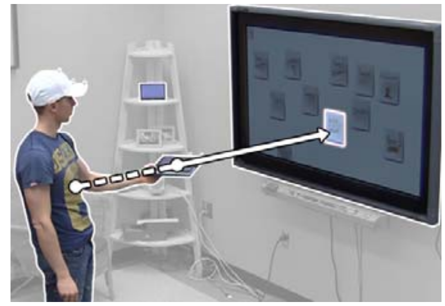

Figure 8. *Point, select and edit*

Point, select, and edit (from a distance). While content on a large display is convenient for viewing, editing may be more efficient on one's portable device. To select content for transferring back to the tablet, the tablet itself can function as a distant pointing device. A person holds the tablet away from his body and points it towards the display. The system calculates the intersection of the pointing ray (here: a line connecting the person's torso and the position of the tablet device) with the large display's surface (Fig. 8). This action highlights the note (with a colored border) that is closest to that intersection point. To transfer the note to the tablet temporarily for editing purposes, the person taps on the tablet's screen. To place back the note on the large display, the person points at a location on the display and again taps the tablet's screen to confirm.

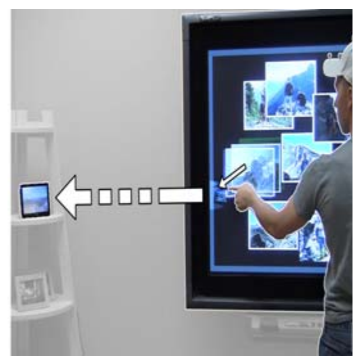

Figure 9. *Portal drag*

Portal drag to transfer (from a distance). We can also exploit the awareness icons of Stage 1 as portals to transfer information between them via drag and drop. Fig. 9 illustrates the Proxemic Photo Canvas on a large display and a picture frame appliance; their awareness icons are visible at their borders. A person is transferring content from the large display to the picture frame simply by dragging a photo onto the picture frame portal, which then shows that image in full size in the frame.

Drag in and out (close proximity). In this variation, illustrated in Fig. 10, the tablet's content is progressively revealed in Stage 2 by growing it in size directly in front of the approaching person. (The area also follows the person's

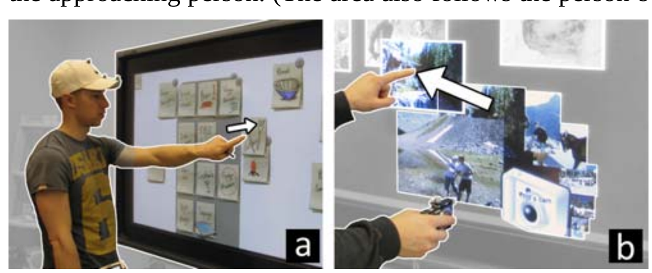

Figure 10. *Drag in and out in close proximity in both applications*

side by side movements). When within direct touch distance, this content becomes interactive, i.e., it allows that person to access his tablet's content by directly touching the large display. In particular, a person transfers content by dragging items into or out of their personal area.

Figs. 10a and b illustrate how Proxemic Brainstorming and Proxemic Photo Canvas allow one to drag notes and photos to an empty region on the screen, which moves them between the devices. While both progressively revealed their contents in visually different ways, the transfer operation is identical.

Figure 11. Point, touch and edit

Point, touch and edit (close proximity). Our next technique is similar to point, select and edit, except that it works in close proximity. In particular, the tablet itself can be used as a physical pointing device, where touching the device on the large screen will pick up or drop off information. This function of the tablet becomes active when a person stands within touch distance, and holds the tablet in a way that one of its corners points at content on the large display. Fig. 11 illustrates this with Proxemic Brainstorming. As the device moves towards the display, a projected pointer highlights the currently selected note. When the person touches a note with a corner of the tablet, the note is picked up and temporarily transferred to the tablet device for editing. After editing, a person can quickly place that note back to a given location on the large display by touching that location with a corner of the tablet.

Interaction of People-to-Devices-to-People

The next suite of techniques is tailored to multiple people collaboratively sharing content with each other through their personal devices, possibly including a large display. Unlike the single user techniques, these include coordination protocols that influence how handoffs are achieved.

Collaborative handoff. In collaborative work scenarios, people may want to pass on digital information to another person. Often, this requires tedious sequences of tasks such as sending files by email or copying and retrieving content using portable media. Our notion of a *proxemic-aware collaborative handoff* represents a simpler method for transferring content between devices. The idea is that one person starts the gesture on their personal device, and a second person continues this gesture on their personal device to complete the handover process. That is, one person cannot transfer information without cooperation from the other person. Both must also be in close proximity before these techniques are activated, and we also expect people to monitor each other's actions in a way that mediates their social protocols. Our technique is inspired

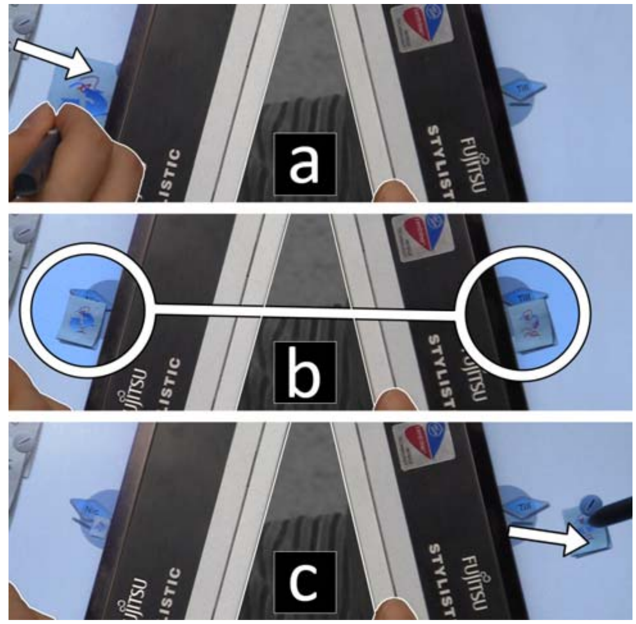

Figure 12. *Collaborative handoff* allows handover of digital information between tablets: (a) dragging content on awareness icon of other tablet, (b) content appears on second tablet, and (c) dragging content away from the icon moves the content permanently to the second device.

by *collaborative* stitching, where a stroke gesture is started by one person on one device and continued by another on another device to stitch together those display workspaces [10].

Figure 12 illustrates an example Proxemic Brainstorming exchange between two people who have moved their tablets besides each other. As before, both are aware of connection availability via progressive reveal, where in this case the awareness icon size is larger as people move closer. Similar to our previously-described 'portal drag to transfer', a person can initiate content-sharing by dragging a sticky note onto the awareness icon of the second person's tablet (Fig. 12a). What is different is that a thumbnail of the content then appears on the second tablet, so that it is temporarily visible on both screens (Fig. 12b). If the second person drags the thumbnail image from the awareness icon onto his screen (thus *continuing* the first person's drag operation), the thumbnail on the first person's tablet disappears and the content is now permanently stored on the second person's device (Fig. 12c). Through this continuation of the gesture that was started by the first person, the second person '*accepts*' the content transfer action. If the person does not accept, the transfer is not performed. As well, if the transfer has not yet been accepted (i.e., phase 2; Figure 12b), the first person can also cancel the transfer by dragging the content back onto his or her screen.

Drag between a public intermediary. Two people can use the shared screen area of the large public display as a way to hand off content. The idea is that because information on that display is public, it implicitly gives permission to both actors to exchange information between their devices.

Figure 13. *Drag between a public intermediary.* (a) person drags note out of his personal interaction area, (b) using the empty space between the interaction areas as a clipboard. (c) second person drags note into his interaction space of the tablet, and (d) the note is now moved to his tablet.

Figure 13 illustrates this. Two people are standing in direct touch distance in front of a large wall display with their tablet device in hand. Via progressive reveal, the personal content of both their devices are visible on the wall display as two interaction areas – one per person – in positions that reflect the side-by-side locations of both people (see the rectangular grey boxes containing sticky notes on the screen in Fig. 13). The large interaction areas on the screen make it easy to view and modify content.

Two different versions illustrate different ways of performing the transfer. In the handoff version, a person can drag a note to the shared public area (i.e., the regions not covered by individual interaction areas) on the large display (Fig. 13a,b), but not into the other person's area. The second person accepts that transfer by picking up this note and drag it to his own interaction area (Fig. 13c,d).

The second version does not require this handoff, relying instead on social protocol as augmented by the high visibility of all actions. Here, a person can directly move (or take) a note directly from one tablet to another by dragging it from one interaction area straight to the other.

IMPLEMENTATION

We implemented the techniques and applications using the Proximity Toolkit [17]. The toolkit wraps different hardware tracking systems (e.g., a motion capturing system, and the Kinect infrared based depth camera) and makes the proxemic information of tracked people and devices easily accessible through an event-driven API. The toolkit functions as a prototyping platform that enables the exploration of interaction techniques considering fine-grained proxemic measurements. While our sensing system is not well suited for wide deployments, we believe these tracking technologies are ideal tools that enable the exploration of proxemic-aware interaction techniques, until more practical tracking systems come to market.

We briefly explain a few essential aspects of our implementation that are in particular relevant for the design of the presented interaction techniques.

First, many of our techniques monitored people's or devices' presence in one or multiple discrete zones around other devices. The system then triggers notifications about entities entering or leaving one of these zones. We extended the proximity toolkit's circular notifications zones to support a wider variety of zones: rectangular, elliptical, and other arbitrary (defined as a polygon mesh) shaped zones (Figure 14). The applications then subscribe for changes (i.e., entities entering or leaving) in these zones.

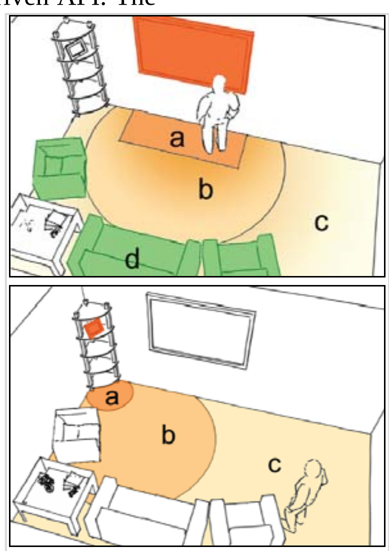

Figure 14. Discrete zones defined around (a) a large wall display and (b) a digital picture frame in the shelf.

Second, our techniques frequently use visual content (on tablets or wall display) that reacts to changes in a person's or devices' proxemic relationship; e.g., the awareness icons that change their size and displayed content depending on distance, orientation, and so on. We designed *reusable proxemic-aware GUI widgets* that facilitate making content on the screen react to one or multiple of the proxemic dimensions. These widgets are designed as containers to contain other widgets as content. The widget's proxemic-aware features are a follows.

(a) *Automatically follow the direction/orientation of any entity tracked in the space around the display*. For example, the awareness icons are always displayed on the edge of the screen closest to the entity they represent.

(b) *Directly map the size of the widget container to any proxemic measures*, e.g., the distance or orientation angles between two entities. Once the monitored entities move, the size is adjusted automatically; e.g., it grows larger when entities move closer.

(c) *React to discrete events for changing the visual appearance of the widget and its content.* Examples are: show content (e.g., when person with device enters a room), hide content (when leaving), or show different content with different level of detail (when person moves closer).

Overall, we see these widgets as a starting point of an extensible programing library of *generic* proxemics-aware widgets. These widgets extend the behaviour of the Proximity Toolkit [17], and encapsulate the common and reused behaviours of software reacting to proxemic changes of people's and devices' relationships.

RELATED WORK

The work presented in this paper primarily relates to three research domains: (1) awareness techniques for nearby devices and their content, (2) connections between such devices, and (3) cross-device transfer of digital information.

Awareness of Nearby Devices and Available Content

Visualization of available devices becomes important in ubicomp environments, as an increasing number of diverse devices are present, and these are sometimes not easily visible to a user. A few systems began exploring methods to inform a person about surrounding devices and possible connections. *ARIS* uses a map to visualize devices located within the same room [3]. Gellersen et al.'s *RELATE* Gateways [6] provide a similar visualization, but make use of sophisticated tracking systems to dynamically update the positions of all devices. In an alternative view, icons at the border of a mobile device screen represent the type and location of surrounding devices (a technique previously introduced by Rekimoto et al. [21]). A later comparative study by Kortuem et al. [14] supports the advantage of using spatial visualizations for device location.

Aside from awareness about devices, visualizing content has also been explored. The aforementioned *ARIS* shows applications running on devices located in the same room in a world-of-miniature fashion [3,23]. In *Drag-and-Pick*, content that is located in the direction of an initial drag operation (even on other devices in that direction) appears close to point of interaction [2]. However, no awareness about what interactions are supported is given.

Our work extends these notions with awareness visualizations incorporating progressive reveal, and with interaction techniques that consider fine-grained distance and orientation relationships to allow a person fluently moving from awareness to direct interaction.

Connecting Devices

Researchers investigated how connections can be established between devices – predominantly between devices in close proximity. Most of the developed systems define *one* discrete spatial region around devices, where a connection is triggered once the distance between two devices becomes smaller than that discrete threshold. Often, this distance depends on the actual sensing technology used (e.g., the limited sensing range of RFID or Bluetooth).

With Smart-its friends [12], a connection can be established once two devices sense similar values through attached sensors (e.g., accelerometers). A cross-device connection can be established by shaking a pair of devices simultaneously. Want's RFID-based technique [26] allows detecting nearby objects and devices and associating or retrieving digital information. Rekimoto et al. later [21] combined RFID and infrared communication for establishing device connectivity. For connections between devices with larger distances, pointing is commonly used. Swindells et al., for example, use an infrared-emitting pen to point at a device and initiate a connection to that device [24].

Gestures are further considered well suited to connect devices [15]. In *Stitching*, users couple devices by drawing a stroke that begins on one display and ends on another [10]. Similarly, *Synchronous Gestures* can be used to bump devices together to initiate a connection [11]. However, those techniques require very close proximity of both devices. Nevertheless, Chong et al. confirmed that proximity is one of the 'big five' categories of how users associate devices [4].

We extend existing work by making use of a broader range of fine-grained spatial relationship measures (i.e., distance, orientation and identity) to determine more subtle aspects of connectivity as a function of proximity.

Techniques for Cross-Device Content Transfer

Previous work introduced various techniques to transfer information between devices, which we categorize according to the distance of involved devices: close proximity or in spatial zones.

Numerous techniques exist that require the user to be within reach of both displays. In *Pick-and-Drop*, users pick up content on one display and place it on another through a digital pen [22]. *Corresponding Gestures* work similarly [19]: to 'pick' up content, users perform a 'selection' gesture on top of the content they want to select. They then perform the same gesture on the target display to complete the transfer. *Touch & Interact* temporarily shifts the interaction focus from a large display onto a mobile device [9]. These techniques are optimized for exchanging content of devices in close proximity or even directly touching one another [10]. Spatial relations and the presence in discrete zones have also been used to mediate the information exchanged between devices. Similar to Vogel et al.'s work [25], *Hello.Wall* [20] introduced the notion of 'distance-dependent semantics', where the distance (here: close, far, out of range) of a device an individual is using from the wall defined the interactions offered and the kind of information shown. Kray's *group coordination negotiation* [16] introduced spatial regions for interaction around mobile phones. Their scenario used these regions to negotiate exchange of information with others. Feedback about a phone's presence in any of the regions was visualized on a tabletop. Content is transferred between devices depending on how they were moved in and out of the three discrete regions.

We extend this prior work, where we contribute techniques that allow a person to move from awareness at a larger distance, to gradually revealing more detail about devices and content when approaching, to direct interaction for transferring digital information between devices when standing in either close proximity or at a distance. We also consider *continuous* changes of distance to drive the interaction, rather than just discrete zones.

In another paper, also currently submitted to UIST [27], we consider F-Formations and micro-mobility as other aspects of proximity. That paper concentrates mostly on transfer techniques, and complements the work reported here.

CONCLUSION

We believe that proxemic-aware interaction techniques can help us design future ubicomp systems that more appropriately react to people's social understanding and use of personal space. These designs can consider the fine-grained proxemic relationships between people's devices, such as distance, orientation and direction of movement, to drive interactions. We argued that gradual engagement is an essential prerequisite of such systems. This spans three stages: (1) awareness notifications about device presence, (2) progressive reveal of available content for sharing, to (3) providing a range of interaction techniques appropriate to the particular contexts as defined by distance, device type, and group engagement. Such awareness becomes even more important as ubicomp ecologies emerge with an increasing number of personal and public devices of all different form factors and capabilities.

The set of techniques we introduced is not a complete or exhaustive set, but a starting point suggesting further interaction techniques that considers people's expectations of proxemics during ubicomp system design.

ACKNOWLEDGMENTS

Both NSERC SURFNET and the NSERC/AITF/SMART Technologies Industrial Chair in Interactive Technologies provided funding via these programs.

REFERENCES

1. Ballendat, T., Marquardt, N., and Greenberg, S. Proxemic Interaction: Designing for a Proximity and Orientation-Aware Environment. *Proc. ACM ITS 2010*
2. Baudisch, P., Cutrell, E., Robbins, D., et al. Drag-and-Pop and Drag-and-Pick: techniques for accessing remote screen content on touch- and pen-operated systems. *Proc. INTERACT'03*, IFIP, IOS Press (2003).
3. Biehl, J.T. and Bailey, B.P. ARIS: an interface for application relocation in an interactive space. *Proc. GI '04*, CHCCS (2004), 107–116.
4. Chong, M.K. and Gellersen, H. How users associate wireless devices. *Proc. CHI '11*, ACM (2011).
5. Fishkin, K.P., Roy, S., and Jiang, B. Some Methods for Privacy in RFID Communication. In *Security in Ad-hoc and Sensor Networks*. 2005, 42–53.
6. Gellersen, H., Fischer, C., Guinard, D., et al. Supporting device discovery and spontaneous interaction with spatial references. *Personal Ubiquitous Comput. 13*, 4 (2009), 255–264.
7. Greenberg, S., Marquardt, N., Ballendat, T., Diaz-Marino, R., and Wang, M. Proxemic Interactions: The New Ubicomp? *ACM Interactions 18*, 1 (2011).
8. Hall, E.T. *The Hidden Dimension*. Doubleday, Garden City, N.Y, 1966.
9. Hardy, R. and Rukzio, E. Touch & interact: touch-based interaction of mobile phones with displays. *Proc. MobileHCI '08*, ACM (2008), 245–254.
10. Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P., and Smith, M. Stitching: pen gestures that span multiple displays. *Proc. AVI '04*, ACM (2004).
11. Hinckley, K. Synchronous gestures for multiple persons and computers. *Proc. UIST '03*, ACM (2003).
12. Holmquist, L., Mattern, F., Schiele, B., Alahuhta, P., Beigl, M., and Gellersen, H.-W. Smart-Its Friends: A Technique for Users to Easily Establish Connections between Smart Artefacts. *Proc. Ubicomp '01*, Springer (2001), 116.
13. Ju, W., Lee, B.A., and Klemmer, S.R. Range: exploring implicit interaction through electronic whiteboard design. *Proc. CSCW '08*, ACM (2008).
14. Kortuem, G., Kray, C., and Gellersen, H. Sensing and visualizing spatial relations of mobile devices. *Proc. UIST '05*, ACM (2005), 93–102.
15. Kray, C., Nesbitt, D., Dawson, J., and Rohs, M. User-defined gestures for connecting mobile phones, public displays, and tabletops. *Proc. MobileHCI '10*, ACM (2010), 239–248.
16. Kray, C., Rohs, M., Hook, J., and Kratz, S. Group Coordination and Negotiation through Spatial Proximity Regions around Mobile Devices on Augmented Tabletops. *Proc. Tabletop '08*, IEEE (2008), 1–8.
17. Marquardt, N., Diaz-Marino, R., Boring, S., and Greenberg, S. The Proximity Toolkit: Prototyping Proxemic Interactions in Ubiquitous Computing Ecologies. *Proc. UIST '11*, ACM (2011).
18. Marquardt, N. and Greenberg, S. Informing the Design of Proxemic Interactions. *In IEEE Pervasive Computing – Special Issue on Pervasive I/O.*, (2012).
19. Nacenta, M.A., Aliakseyeu, D., Subramanian, S., and Gutwin, C. A comparison of techniques for multi-display reaching. *Proc. CHI '05*, ACM (2005).
20. Prante, T., Röcker, C., Streitz, N., et al. Hello. Wall–Beyond Ambient Displays. *Adj. Proc. Ubicomp '03*.
21. Rekimoto, J., Ayatsuka, Y., Kohno, M., and Oba, H. Proximal Interactions: A Direct Manipulation Technique for Wireless Networking. *Proc. INTERACT '03*, (2003), 511–518.
22. Rekimoto, J. Pick-and-drop: A Direct Manipulation Technique For Multiple Computer Environments. *Proc. UIST '97*, ACM (1997), 31–39.
23. Stoakley, R., Conway, M.J., and Pausch, R. Virtual reality on a WIM: interactive worlds in miniature. *Proc. CHI '95*, ACM (1995), 265–272.
24. Swindells, C., Inkpen, K.M., Dill, J.C., and Tory, M. That one there! Pointing to establish device identity. *Proc. UIST '02*, ACM (2002), 151–160.
25. Vogel, D. and Balakrishnan, R. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. *Proc. UIST '04*, ACM (2004), 137–146.
26. Want, R., Fishkin, K.P., Gujar, A., and Harrison, B.L. Bridging Physical and Virtual Worlds with Electronic Tags. *Proc. CHI '99*, ACM (1999), 370–377.
27. Cross-Device Interaction via Micro-mobility and F-Formations. In submission to UIST 2012, Authors removed for blind review.