

Gradual Engagement: Facilitating Information Exchange between Digital Devices as a Function of Proximity

Nicolai Marquardt¹, Till Ballendat¹, Sebastian Boring¹, Saul Greenberg¹, Ken Hinckley²

¹University of Calgary,

Department of Computer Science, Calgary, AB, Canada

{nicolai.marquardt, sebastian.boring, saul}@ucalgary.ca, till@ballendat.info

²Microsoft Research

One Microsoft Way, Redmond WA 98052

kenh@microsoft.com

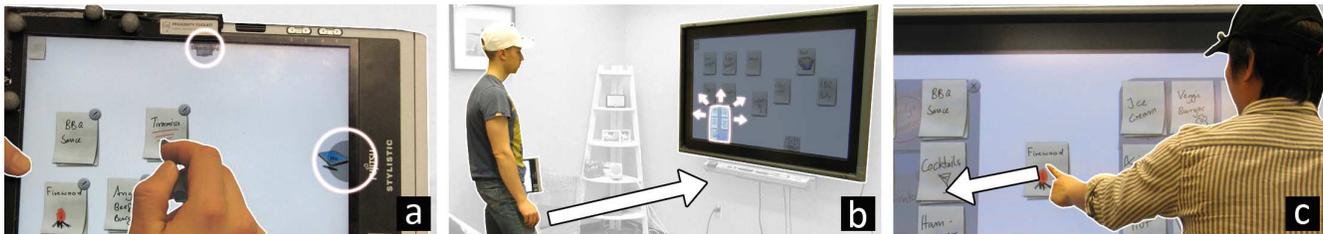


Figure 1. Gradual engagement, showing examples of (a) awareness, (b) progressive reveal, which (c) leads to information transfer

ABSTRACT

The increasing number of digital devices in our environment enriches how we interact with digital content. Yet, cross-device information transfer – which should be a common operation – is surprisingly difficult. One has to know which devices can communicate, what information they contain, and how information can be exchanged. To mitigate this problem, we formulate the *gradual engagement* design pattern that generalizes prior work in proxemic interactions and informs future system designs. The pattern describes how we can design device interfaces to gradually engage the user by disclosing connectivity and information exchange capabilities as a function of inter-device proximity. These capabilities flow across three stages: (1) *awareness* of device presence/connectivity, (2) *reveal* of exchangeable content, and (3) interaction methods for *transferring* content between devices tuned to particular distances and device capabilities. We illustrate how we can apply this pattern to design, and show how existing and novel interaction techniques for cross-device transfers can be integrated to flow across its various stages. We explore how techniques differ between personal and semi-public devices, and how the pattern supports interaction of multiple users.

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

Keywords: gradual engagement; proximity; awareness; proxemic interactions; handhelds; interactive surfaces

INTRODUCTION

Personal mobile devices (e.g., phones, tablets) and semi-public stationary devices (e.g., information appliances, in-

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

ITS'12, November 11–14, 2012, Cambridge, Massachusetts, USA.

Copyright 2012 ACM 978-1-4503-1209-7/12/11...\$15.00.

teractive surfaces) are an increasingly commonplace way for people to ubiquitously interact with 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 interaction 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., [13,16,30]. However, significant challenges remain. People do not know *which* devices can communicate with one another, *what* information they contain that is exchangeable, and *how* information can be exchanged in a controlled manner.

To mitigate these problems, as our primary contribution we formulate a design pattern called *gradual engagement*, which we then refine to ease the information transfer task. As a design pattern [5,33], its strengths lie in (1) unifying prior work in proxemic interaction, (2) synthesizing essential, generalizable interaction strategies, and (3) providing a common vocabulary for discussing design solutions. Most importantly, the pattern informs and inspires future designs, but also allows for variations of the pattern applied to different domains.

As we will see, the gradual engagement pattern describes how devices can gradually engage the user by disclosing connectivity and information exchange capabilities as a function of proximity. That is, as people move and orient their personal device towards other surrounding devices [11,24], the interface progressively moves through three stages affording gradual engagement: (1) *awareness* of device presence and connectivity, (2) *reveal* of exchangeable digital content, and (3) interaction methods for *transferring* digital content between devices tuned to particular distances and device capabilities (cf. Figures 1 & 2).

As our secondary contribution, we illustrate how we can apply this pattern to design, and demonstrate how adapta-

tions of existing interaction techniques flow across the three stages. In addition, we introduce novel interaction techniques that make use of this design pattern. In particular, we explore how gradual engagement techniques differ when the other device seen is personal (such as a handheld) vs. semi-public (such as a large display), and how gradual engagement applies to multi-user *collaborative* activities.

We begin with a brief review of proxemics and proxemic interactions, as these are the theoretical grounds behind our work. We then introduce the gradual engagement design pattern: how it emerges from prior work in proxemic interactions, and how we refine it to ease the information transfer task. We then describe particular interaction techniques, where our presentation is structured according to the three stages of the gradual engagement design pattern. The accompanying video illustrates these techniques in action.

PROXEMICS FOR UBICOMP INTERACTIONS

Proxemics – as introduced by anthropologist Edward Hall [12] – 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. 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. Proxemics mediates many aspects of social interaction. For example, it influences casual and serendipitous encounters [20], is a nuance in how people greet one another [Ch. 6 in 19], and is a major factor in how people arrange themselves for optimal small group collaboration via spatial-orientational maneuvering [Ch. 4 & 7 in 19]. Proxemics as a social construct also helps people *gradually engage* with one another: people are spatially aware of others around them, how they approach each other to signal interest, and how they arrange themselves as they engage in conversation.

Within HCI, a variety of researchers used proxemics to motivate particular ubicomp system designs (e.g., [18,35]; see related work). This evolved *proxemic interactions* [2,24], a general construct that applies the insights of proxemic theory to the holistic design of ubicomp interaction. Proxemic interactions describe five important dimensions to consider when designing proxemic-aware ubicomp systems. These dimensions operationalize the relationships between people, devices, and objects as sensed or stored variables: the *distance* between entities, their relative *orientation* to one another, their relative *movement*, their *identity*, and *location* features that give further meaning to that setting [11].

PROXIMITY AND GRADUAL ENGAGEMENT

The vast majority of interfaces are premised on the notion that a person is fully attending it, i.e., the system is designed to support *foreground* activities and tasks. However, a variety of systems also recognize that the person may not be directly attending them (i.e., it is in the *background* of their attention), where they still try to be helpful by presenting an interface that selectively informs the user of information of interest.

One class of examples includes ambient displays [22] embedded in a physical environment. The display usually presents non-critical information unobtrusively, which a person can monitor at a distance and at the periphery of their attention (thus providing basic awareness). The display often contains a way for the person to easily transition to more in-depth information exploration if the person decides to engage with it; this normally occurs by the person approaching and directly interacting with that display. That is, such displays implicitly incorporate a *binary notion* of proximity: from afar, and within interaction reach.

Proxemic interactions provide another class of examples that use a much more refined notion of proxemics. Many proxemic interaction systems commonly 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. Influential earlier work considered such gradual increasing engagement between a *person* and large interactive displays [2,18,35,36]. For example, Vogel et al. directly applied Hall’s theory to a person’s interaction with a public display [35]. They defined four discrete zones around the display that affect a person’s interaction when moving closer: from far to close, interactions range from ambient display of information, to implicit, subtle, and finally personal interaction. Similarly, Ju’s interaction techniques with a digital whiteboard remain public and implicit from a distance, and become increasingly more private and explicit when the person moves closer to that display. Ballestad et al. introduced proxemic interaction concepts illustrated with an interactive media player [2], where the device reacts to one or multiple people’s proxemics relationships by varying displayed content and supporting diverse modes of interaction. Wang et al. described a proximity-based advertising display that changes its presentation of information not only to engage people for interaction, but to try to re-attract them if they appear distracted [36].

We generalize the sequence inherent in these (and other) systems as a design pattern we call *gradual engagement*¹. The basic idea is that:

- background information supplied by the system provides awareness to the person about opportunities of potential interest when viewed at a distance;

¹ Jan Borchers describes an even more general pattern titled ‘Attract-Engage-Deliver’ [5]. The difference is that our pattern incorporates proxemics as a first-class element.

- the person can gradually act on particular opportunities by viewing and/or exploring its information in more detail simply by approaching it; and
- the person can ultimately engage in action if so desired.

This pattern is, of course, directly inspired by proxemic theory and by the systems that reflect that pattern. The pattern also characterizes what we thought was the ‘best’ of how proxemics was previously applied to ubicomp design.

APPLYING GRADUAL ENGAGEMENT TO CROSS-DEVICE INFORMATION TRANSFER

The previously mentioned systems are primarily focused on people’s interaction with large displays, where the display’s content changes as a function of a person’s distance. However, the relatively recent explosion of smart phones and other hand-held devices, as well as the more general availability of computer displays embedded in the environment, mean that people are now living in a dynamic device ecology. This begs the question of how people can interact *across* such devices, where we focus on the particular problem of how people can transfer information between them.

We refine the *gradual engagement* design pattern by considering fine-grained proxemic relationships *between multiple devices* allowing seamless transitions from awareness to information transfer. Specifically, engagement increases continuously across *three stages* as people move and orient their personal device towards other surrounding devices (Fig. 2):

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. We leverage knowledge about proxemic relationships between devices to determine when devices connect and how they notify a person about their presence and established connections.

Stage 2. Reveal of exchangeable content is provided, so that people know what of their content can be accessed on other devices for information transfer. At this stage, a fundamental technique is progressively revealing a device’s available digital content as a function of proximity.

Stage 3. Transferring digital content between devices, tuned to particular proxemic relationships and device capabilities, is provided via various strategies. Each is 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.

We illustrate the application of this gradual engagement pattern between devices as a suite of interaction techniques, all based on providing a seamless transition leading from awareness, to reveal, to interaction. The remainder of the paper will revisit each of the

three stages to introduce these techniques. First, however, we review prior work that relates to our derivation of the gradual engagement pattern.

PRIOR WORK APPLIED TO GRADUAL ENGAGEMENT

We briefly sample prior work that contributed to our derivation of particular stages of the gradual engagement pattern. Beyond the review, we also later explain how many of these prior techniques can be applied to people’s interactions across all three stages of gradual engagement. Figure 2, bottom, summarizes how these works fit within and thus contribute to the pattern.

Awareness of Device Presence and Connectivity

Most systems define a discrete spatial region around devices, where a connection is established (and information transfer possible) once the distance becomes smaller than a certain threshold. Often, this distance depends on the actual sensing technology used (e.g., sensing range of RFID or Bluetooth).

Visualization of available devices becomes important in ubicomp environments, as an increasing number of diverse devices are present. Their presence, location, and ability to connect (or not) are rarely easily visible to a user. A few systems began exploring methods to inform a person about surrounding devices and possible connections. Most commonly, a map visualizes devices located in the environment (e.g., Sentient Computing [1]) or the same room (e.g., ARIS [4]). Gellersen et al.’s *RELATE* Gateways provide a similar visualization, but make use of sophisticated tracking systems to dynamically update the positions of all devices [10]. In an alternative view, icons at the border of a mobile device screen represent the type and location of surrounding devices [10,29]. Kray’s *group coordination negotiation* introduced spatial regions for interaction around mobile phones [21]. 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. Marquardt et al. observe F-Formations (i.e., patterns of how people stand in close proximity formations) to determine when to federate devices [25].

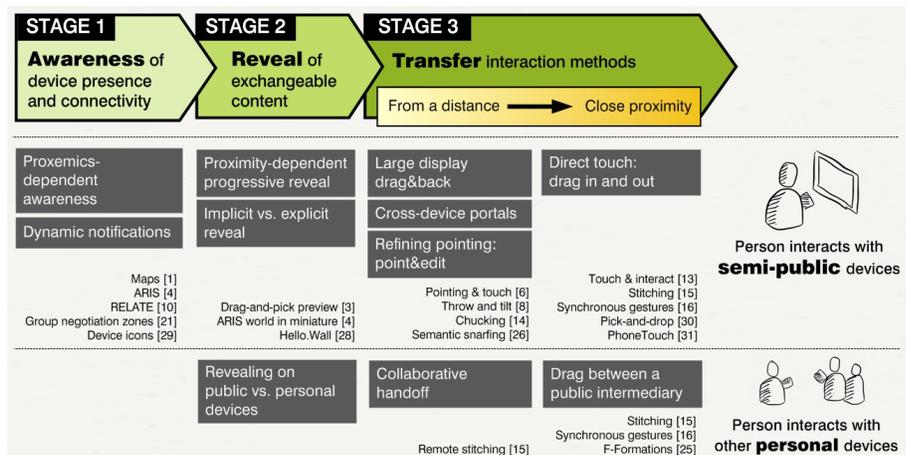


Figure 2. Three sequential stages of the gradual engagement pattern (top row) and interaction methods derived from the pattern supporting awareness and interaction in each stage: person interacting with *semi-public* devices (middle row) or *personal* devices of other people (bottom row).

Explicit Connections. Various systems allow people to manually associate two devices from a distance. This is usually done via pointing one device at the other. Swindells uses an infrared-emitting pen to point at a device to control it [32]. *Semantic snarfing* [26] also uses pointing to allow someone to take over temporary control of remote interfaces. Similarly, others have suggested ways to manually associate nearby devices that are all within reach. With *Smart-its friends* [17], a connection is established when two devices are shaken simultaneously and sense similar accelerometer values. In *Synchronous Gestures* people can bump devices [16] – including phones and interactive tabletops [31] – together to initiate a connection. In *Stitching*, users couple devices by drawing a stroke that begins on one display and ends on another [15]. Overall, Chong et al. confirmed that proximity is one of the ‘big five’ categories of how users associate devices [7].

Revealing Exchangeable Content.

Several systems visualized exchangeable content. *Hello.Wall* introduced the notion of ‘distance-dependent semantics’, where the distance (here: close, far, out of range) of a person’s device from the wall screen defined the kind of information shown on the mobile display [28]. The aforementioned *ARIS* shows applications running on devices located in the same room in a world-of-miniature fashion [4]. In *Drag-and-Pick*, content that is located in the direction of an initial drag operation appears close to the point of interaction – even on other devices in that direction [3].

Transferring Digital Content

Once connected, diverse techniques allow information transfer. For example, Want’s RFID-based technique allows detecting nearby objects and devices and associating/retrieving digital information [37]. In *Pick-and-Drop*, users pick up content on one display and place it on another with a digital pen [30]. *Touch & Interact* temporarily shift the interaction focus and content from a large display onto a mobile device [13]. Somewhat later, Marquardt et al. consider F-Formations and micro-mobility to drive device-to-device information transfer while people stand in close proximity formations [25]. Rekimoto combined near-field RFID and remote infrared communication for seamless information transfer [29]. Further examples for cross-device information transfer from a distance are: throwing gestures performed with a phone [8], touch and pointing gesture combinations [6], chucking motions towards the other device [14], or *corresponding Gestures* through cursor selections in multi-screen environments [27].

In summary, various techniques exist – most suited for particular discrete distances between devices – that fit into particular stages (but rarely all stages) of the gradual engagement pattern (see Figure 2). In the next sections, we use our design pattern to build on these earlier works. In particular, we illustrate interaction techniques that allow a person to move seamlessly 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. By extending earlier work, we also consider how particular device types can influence this interaction, e.g., personal handhelds vs. semi-public stationary devices.

RUNNING EXAMPLE: PROXEMIC BRAINSTORMING

We use an example application throughout the paper as a running example to illustrate how our various techniques leverage proxemic interaction and follow gradual engagement to facilitate access to digital information. *Proxemic Brainstorming* is a multi-user, 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. The video figure illustrates this application and the dynamics of the various methods described below. The video also demonstrates a second application, applying the pattern to facilitate transfer of digital photos from a network enabled digital camera to other devices, such as a large display or a digital photo frame.

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. Building upon [1,4,10,29], the basic idea is that a person sees a visual indicator – a subtle notification – about which devices as well as their location in the surrounding environment can possibly interact with his handheld device (e.g., icons in Figure 3). People can then 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.

Proxemics-dependent awareness. We use rules to determine when to trigger awareness of device presence and connectivity. By connection, we mean whether one device

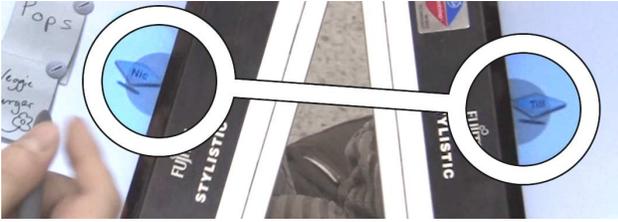


Figure 3. Icons at the edge of the screen indicate the presence and location of other devices in close proximity (2 tablet computers).

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* [11] as sensed factors: 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. Location also informs context; some locations (e.g., public vs. home spaces) would afford quite different connectivity semantics.

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 – is an indicator 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 connect and ultimately interact with both devices in tandem.

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 (at the extreme, this becomes pointing [26,32]). 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” [9], and similarly that closer proximity between devices implies trust (although this depends on location context). 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.

Dynamic 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 ubi-comp 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 (Figures 1a, 3, 4 & 5a). Each icon informs the user: where the device represented by the icon is physically located; that there is a potential connection between those devices; and that devices can interact with one other (e.g., allowing information transfer). Icon appearance can be informative, such as a graphic that represents the nearby tablet. Or they can be augmented with other information, such as the name of that device and/or its owner.

Figure 3 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 and the name of the device’s owner. Extending earlier work (e.g., [2,10,29]), icon locations are continuously animated around the edge to represent the directional location of the corre-

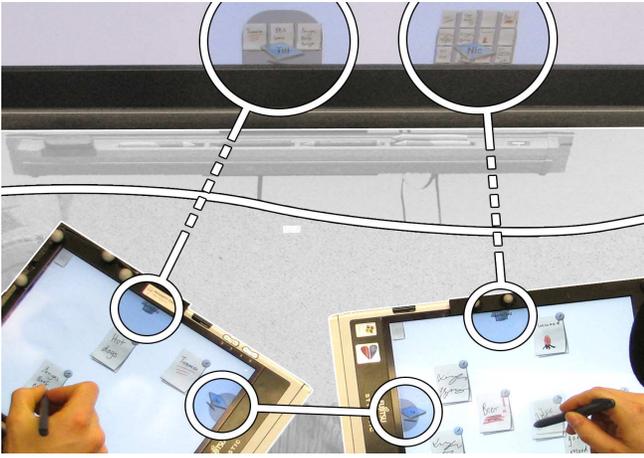


Figure 4. Content awareness: Proxy icons indicate the presence of nearby tablets and a large interactive display. Available content of the tablets is displayed as thumbnails atop these icons.

sponding device. In Figure 3, we see how both displays’ icon locations illustrate their physical spatial relationship. Figure 4 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 own 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 to reflect the new relationship between the tablet and each device. Thus, a person can move seamlessly towards and gradually engage the particular device desired for interaction.

STAGE 2: REVEAL OF EXCHANGEABLE CONTENT

As proximity increases, the gradual engagement pattern suggests that devices should reveal *content* available for exchange. 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.

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. Building upon work presented in [2], 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 three reasons. First, it 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 [35]. 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. Third, it provides implicit security: in public contexts, fine details may appear in small size, and only when a person is (say) directly in front of other device, thus masking it from passersby.

For example, Figure 5 illustrates how Proxemic Brainstorming *continuously* reveals content during Stage 2 – in this case multiple sticky notes located on people’s tablets – as they move closer to the large display. The wall display shows thumbnails of all sticky notes located on the tablets above the awareness icons (Fig. 5, right side). For the person sitting at a distance, the actual text on these notes is not yet readable, but the number of available notes is already visible. For the second person moving closer to the wall display, the thumbnails increase in size continuously (5b). For the third person standing directly in front of the display, the sticky notes are shown at full size (5c), allowing the person to read the text of all notes stored on the tablet and to pursue Stage 3 interactions, explained shortly.

Alternatively, instead of continuous growth of the awareness information, we can reveal content progressively at *discrete* distance thresholds. For example, the brainstorming application could first show only a device’s awareness icon when at a larger distance (Stage 1), switch to showing

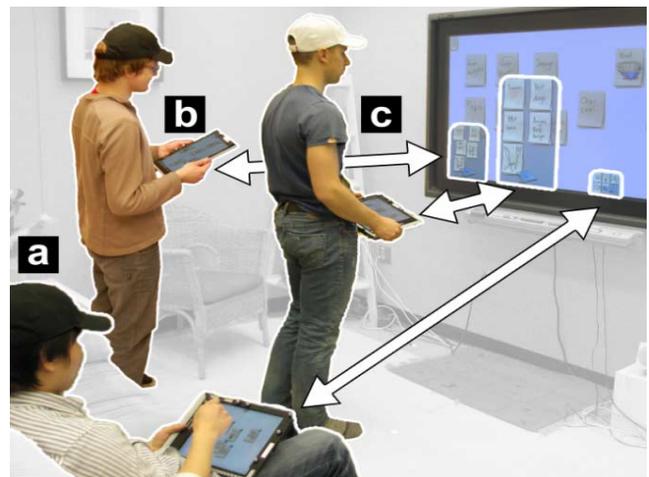


Figure 5. Proximity-dependent progressive reveal of personal device data of multiple users at different distances to the display: (a) minimal awareness of person sitting further away, (b) larger, visible content of a person moving closer, and (c) large awareness icons of person standing in front of the display.

a single thumbnail of the latest sticky note when the device is closer (Stage 2), and finally switch again to show all content once the person with the tablet stands in front of the screen. Furthermore, both continuous and discrete reveal can be combined, in a way that discrete stages trigger content changes and continuous distance changes affect the size of the currently displayed content.

Implicit vs. Explicit Reveal. The above method illustrates 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. In this case, once content becomes revealed during Stage 2 the person can perform an explicit action to reveal more content. For example, tapping on the tablet screen allows cycling through multiple sets of sticky notes. Of course, alternative forms of explicit input (e.g., hand gestures, device movement) could be considered to cause similar explicit reveal behaviours (the video figure shows explicit reveal with a tilt-to-scroll technique for revealing additional photos of a camera).

Revealing content on personal vs. public devices. The information revealed about available content on the display of other devices should differ between *personal* and *semi-public* devices. For personal devices, we currently only provide an awareness icon of surrounding devices, but not their content. This is partially due to privacy reasons, but also size constraints: showing 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.

Semi-Public devices (e.g., a wall-mounted display in a meeting room), however, reveal content located on one’s *personal devices* as one approaches the display. For example, the wall display in Figure 4 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.

STAGE 3: TECHNIQUES FOR INFORMATION TRANSFER BETWEEN DEVICES

Stage 1 and 2 indicate device presence, connectivity and available content, eventually leading to Stage 3 of the gradual engagement pattern, where a person can interact with progressively revealed content. We now present a series of novel interaction techniques (and others from related work) 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 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 *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 he does 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

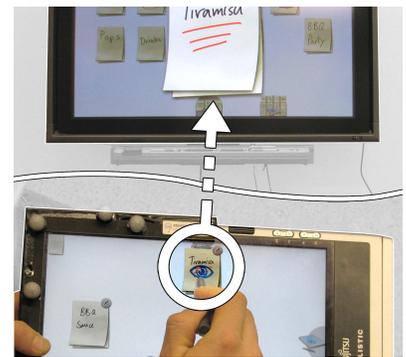


Figure 6. Large display drag and back: (bottom) dragging content on the wall display’s awareness icon on the tablet, (top) content appears full screen on the large display. Notes from multiple users are shown side by side.

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. This allows fast and easy sharing of digital content with all members of a group.

Cross-device 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. This extension of the *portals* concept [34] supports transfer methods across *multiple* devices. Fig. 7 illustrates the transfer between a large display and a small screen; their awareness icons are visible at their borders. A person is

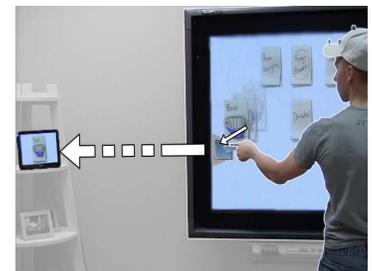


Figure 7. Cross-device portal drag

transferring content from the large display to the small screen simply by dragging a note onto the screen’s portal, which then shows that note in full size on the small screen.

Integrating and refining pointing techniques: point & edit (from a distance). Considering all stages of the gradual engagement pattern also helps us integrate and refine existing gestural interaction techniques for information transfer.

For example, we integrated the *point & edit* technique as a refinement of *semantic snarfing* [26] and *touch+air* [6] pointing gestures 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 specific proxemic relationship of person and device is triggering the pointing mode). The system calculates the intersection of the pointing ray with the large display’s surface. 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.

Direct touch: drag in and out (close proximity). In this technique, illustrated in Fig. 8, 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 side-by-side movements). When within direct touch distance to the large display, 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 between the two devices (tablet & large display) by dragging items into or out of their personal area. Fig. 8 illustrate how Proxemic Brainstorming allows one to drag notes to an empty region on the screen, which transfers them across devices.



Figure 8. Drag in and out in close proximity

Again, we can integrate refinements of existing techniques at this stage. For example, inspired by *PhoneTouch* [31], the tablet itself can now be used as a physical pointing and selection device. Touching the device on the large screen will pick up or drop off information. Considering proxemics refines this technique: the pointing 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. As the device moves towards the display, a projected pointer highlights the currently selected note (thus providing continuous feedback before touching). 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 edit-

ing, 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.

Collaborative Transfer

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 (from a distance). 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* (inspired by collaborative stitching [15]), represents a simpler method for transferring content between devices. The idea is that one person starts the gesture on his personal device, and a second person continues this gesture on his 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. We expect people to monitor each other’s actions in a way that mediates their social protocols.

Figure 9 illustrates an example of content exchange in the Proxemic Brainstorming application between two people who have moved their tablets besides each other. As before, both are aware of connection availability via progressive reveal,

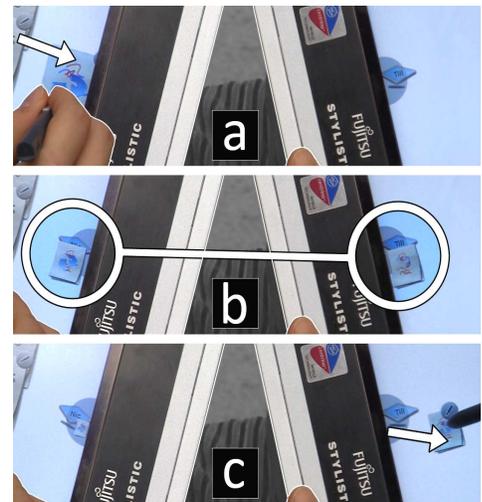


Figure 9. Collaborative handoff. (a) dragging content onto awareness icon representing the other tablet, (b) content appears on 2nd tablet, and (c) dragging content off the icon transfers it.

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. 9a). 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. 9b). 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



Figure 10. 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.

(Fig. 9c). 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 9b), the first person can cancel the transfer by dragging the content back onto his or her screen.

Drag between a public intermediary (close proximity). 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 10 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. 10). 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. 10a,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. 10c,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 move (or take) a note directly from one tablet to another by dragging it from one interaction area straight to the other.

DISCUSSION, FUTURE WORK, AND CONCLUSION

Gradual engagement, generalizability, limitations.

We believe that proxemic-aware interaction techniques following the gradual engagement pattern can help designing future ubicomp systems so that they drive the information transfer process in a way that more appropriately reacts to people’s social understanding of personal space. Besides our presented application to cross-device transfers, the generalized gradual engagement pattern can be applied to other areas, e.g., interactive advertisements [36] or games [11]. In a similar way, these applications could benefit from (1) giving awareness notifications about presence, (2) reveal of content or possible interactions, and (3) providing a range of interac-

tion techniques appropriate to the particular contexts as defined by distance, orientation, and group engagement. Overall, the gradual engagement pattern and our derived set of techniques is neither a complete or exhaustive set, nor do they handle all issues that will likely emerge. Rather, it is a starting point suggesting further exploration of patterns and interaction techniques for ubicomp system design.

Large ecologies of people and devices.

We believe the gradual engagement pattern will be more and more relevant as ubicomp ecologies emerge with an increasing number of personal and public devices featuring different form factors and capabilities. As shown, the pattern and the techniques we derived support a variety of 1-to-1 (e.g., *collaborative handoff*, *cross-portal drag*) and 1-to-many (e.g., *progressive reveal*, *large display drag and back*) collaborative settings. A major advantage of using the stages of gradual engagement is that it leads to an implicit *filtering* of choices presented to a user. For example, following the pattern prevents a system showing an overwhelming number of icons of all present devices in a large group, but instead fosters a design that only reveals the device presence between neighbours. Nevertheless, future work may extend our techniques to offer alternative 1-to-many sharing possibilities, e.g., by allowing dragging content on multiple device icons to share transfer content to all devices of a group (implemented as a cross-device group propagation technique in [25]).

Gradual engagement and privacy.

We recognize that the gradual engagement pattern for cross-device transfer can introduce privacy concerns in some situations. Thus, designs must incorporate safeguards to guarantee privacy of people’s personal information. First, as mentioned in Stage 1, location information is essential to drive sharing behaviour. For example, while it is likely to have loose sharing between devices at home, only restricted sharing might be desired in the office, and no sharing (but maybe only device awareness) in public settings. Second, implicit protection rules can be applied to prevent sharing. For example, a device in a person’s pocket stays invisible, but shares content when the person takes it out and points it towards other devices. Third, explicit actions and commands on the device can allow a person to manually stop sharing and close inter-device connections at any time (e.g., by pressing a button). This is a rich and fertile area for future work.

Pattern applied to different tracking hardware.

The principles of the gradual engagement design pattern can be applied to interactive systems using diverse *high-* and *low-fidelity* tracking systems. While the technology can limit

or enhance how the design pattern is applied in a particular situation, the pattern itself goes beyond any specific technology.

At the high-fidelity end of the spectrum of possible hardware, our implementation uses an infrared-based motion capturing system [23]. The precise tracking information it provides of people and devices' distance or orientation allowed us to explore a large part of possible kinds of interactions. Such a system, however, is not suitable for wide deployments.

At the opposite side of the spectrum, a possible low-fidelity system using sensor fusion of depth-camera streams and wireless-radio signals for distance + orientation measurements (e.g., [25]) can similarly integrate gradual engagement methods. Despite its lower tracking fidelity, it would allow for applying diverse methods across all three stages of the pattern, such as progressive reveal, cross-device portals, and/or collaborative handoff.

Many other tracking systems would support gradual engagement as well. For example, eye-tracker based systems can provide interaction awareness, reveal information, and offer interaction methods depending on people's attention through eye contact. Likewise, a system using GPS based positioning and digital compass data could apply the pattern in a larger-scale outdoor deployment. Overall, the gradual engagement pattern has the potential of being applied to many other proxemic-aware systems with diverse tracking capabilities along this low-/high-fidelity spectrum.

REFERENCES

1. Addlesee, M., Curwen, R., Hodges, S., et al. Implementing a Sentient Computing System. *Computer* 34, 8 (2001), 50–56.
2. Ballendat, T., Marquardt, N., and Greenberg, S. Proxemic Interaction: Designing for a Proximity and Orientation-Aware Environment. *Proc. of ITS*, ACM (2010).
3. 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*, IOS (2003).
4. Biehl, J.T. and Bailey, B.P. ARIS: an interface for application relocation in an interactive space. *Proc. of GI*, CHCCS (2004).
5. Borchers, J. *A Pattern Approach to Interaction Design*. Wiley, 2001.
6. Bragdon, A., DeLine, R., Hinckley, K., and Morris, M.R. Code space: touch + air gesture hybrid interactions for supporting developer meetings. *Proc. of ITS*, ACM (2011).
7. Chong, M.K. and Gellersen, H. How users associate wireless devices. *Proc. of CHI*, ACM (2011), 1909–1918.
8. Dachselt, R., Buchholz, R. Natural throw and tilt interaction between mobile phones and distant displays. *CHI EA*, ACM (2009).
9. Fishkin, K.P., Roy, S., and Jiang, B. Some Methods for Privacy in RFID Communication. *Proc. of SASN*, ACM (2005).
10. Gellersen, H., Fischer, C., Guinard, D., et al. Supporting device discovery and spontaneous interaction with spatial references. *PUC*, 13, 4 (2009).
11. Greenberg, S., Marquardt, N., Ballendat, T., Diaz-Marino, R., and Wang, M. Proxemic Interactions: The New Ubicomp? *ACM Interactions* 18, 1 (2011), 42–50.
12. Hall, E.T. *The Hidden Dimension*. Doubleday, 1966.
13. Hardy, R., Rukzio, E. Touch & interact: touch-based interaction of mobile phones with displays. *Proc. CHI*, ACM (2008).
14. Hassan, N., Rahman, M.M., Irani, P., and Graham, P. Chucking: A One-Handed Document Sharing Technique. *Proc. of INTERACT*, Springer (2009).
15. Hinckley, K., Ramos, G., Guimbretière, F., Baudisch, P., and Smith, M. Stitching: pen gestures that span multiple displays. *Proc. of AVI*, ACM (2004), 23–31.
16. Hinckley, K. Synchronous gestures for multiple persons and computers. *Proc. of UIST*, ACM (2003), 149–158.
17. Holmquist, L., et al. Smart-Its Friends: A Technique for Users to Easily Establish Connections between Smart Artefacts. *Proc. of Ubicomp*, Springer (2001), 116.
18. Ju, W., Lee, B.A., and Klemmer, S.R. Range: exploring implicit interaction through electronic whiteboard design. *Proc. CSCW*, ACM (2008).
19. Kendon, A. *Conducting Interaction: Patterns of Behavior in Focused Encounters*. Cambridge University Press, 1990.
20. Kraut, R., Egido, C., Galegher, J. Patterns of contact and communication in scientific research collaboration. *Proc. CSCW*, ACM (1988).
21. 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*, IEEE (2008).
22. Mankoff, J., Dey, A.K., Hsieh, G., Kientz, J., Lederer, S., and Ames, M. Heuristic evaluation of ambient displays. *Proc. of CHI*, ACM (2003), 169–176.
23. Marquardt, N., Diaz-Marino, R., Boring, S., and Greenberg, S. The Proximity Toolkit: Prototyping Proxemic Interactions in Ubiquitous Computing Ecologies. *Proc. UIST*, ACM (2011).
24. Marquardt, N. and Greenberg, S. Informing the Design of Proxemic Interactions. In *IEEE Pervasive Computing – Special Issue on Pervasive I/O*, (2012), 14–23.
25. Marquardt, N., Hinckley, K., and Greenberg, S. Cross-Device Interaction via Micro-mobility and F-formations. *Proc. of UIST*, ACM (to appear, 2012).
26. Myers, B.A., et al. Interacting at a Distance Using Semantic Snarfing. *Proc. of Ubicomp*, Springer (2001), 305–314.
27. Nacenta, M.A., Aliakseyeu, D., Subramanian, S., and Gutwin, C. A comparison of techniques for multi-display reaching. *Proc. of CHI*, ACM (2005), 371–380.
28. Prante, T., Röcker, C., Streitz, N., et al. Hello. Wall–Beyond Ambient Displays. *Adjunct Proceedings of Ubicomp*, (2003).
29. Rekimoto, J., Ayatsuka, Y., Kohno, M., and Oba, H. Proximal Interactions: A Direct Manipulation Technique for Wireless Networking. *Proc. INTERACT*, (2003).
30. Rekimoto, J. Pick-and-drop: A Direct Manipulation Technique For Multiple Computer Environments. *Proc. UIST*, ACM (1997).
31. Schmidt, D., Chehimi, F., Rukzio, E., and Gellersen, H. PhoneTouch: a technique for direct phone interaction on surfaces. *Proc. of UIST*, ACM (2010), 13–16.
32. Swindells, C., Inkpen, K.M., Dill, J.C., and Tory, M. That one there! Pointing to establish device identity. *Proc. of UIST*, ACM (2002), 151–160.
33. Tidwell, J. *Designing Interfaces: Patterns for Effective Interaction Design*. O'Reilly Media, Inc., 2005.
34. Voelker, S., Weiss, M., Wacharamanatham, C., and Borchers, J. Dynamic portals: a lightweight metaphor for fast object transfer on interactive surfaces. *Proc. of ITS*, ACM (2011).
35. Vogel, D., Balakrishnan, R. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. *Proc. UIST*, ACM (2004).
36. Wang, M., Boring, S., and Greenberg, S. A Public Advertising Display that Captures and Preserves the Attention of a Passerby. *Proc. of Pervasive Displays*, ACM (2012).
37. Want, R., Fishkin, K.P., Gujar, A., and Harrison, B.L. Bridging Physical and Virtual Worlds with Electronic Tags. *Proc. of CHI*, ACM (1999), 370–377.