

SURFNET /

Designing Digital Surface Applications

Chapter Extract:

Scott, S.D., Besacier, G., McClelland, P., Tournet, J., Goyal, N., Cento, F. (2016). Cross-Device Content Transfer in Table-Centric Multi-Surface Environments. In Surfnet: Designing Digital Surface Applications. (Maurer, F., Ed.). University of Calgary, Calgary, Canada, SurfNet, pp. 53-79. ISBN 978-0-88953-388-2- (PDF)

Frank Maurer, (Ed.)



© All Authors as listed

This book is licensed under a CREATIVE COMMONS License,
available for download at <http://dspace.ucalgary.ca/handle/1880/50450>.

ISBN 978-0-88953-387-5 (Softcover)

ISBN 978-0-88953-388-2 (PDF)

Published by:



University of Calgary
2500 University Drive, NW
Calgary, AB Canada
T2N 1N4

This book has been published using funds from the Natural
Sciences and Engineering Research Council of Canada
(NSERC), Strategic Network Grants program.



NSERC
CRSNG



Cross-Device Content Transfer in Table-Centric Multi-Surface Environments

Stacey D. Scott, Guillaume Besacier, Phillip McClelland, Julie Tournet, Nippun Goyal, and Frank Cento

Introduction

There has been increasing interest in the surface computing community to use small, personal surfaces, such as tablets or smartphones, in conjunction with large surfaces, such as interactive walls and digital tabletops. Combining personal and large surfaces into a functional multi-surface environment (MSE) introduces new design challenges. For example, effective mechanisms are needed for transferring content across different surfaces to allow the most flexible use of the available personal and large surfaces. Significant cross-device transfer research exists in the Human-Computer Interaction (HCI) and Computer-Supported Cooperative Work (CSCW) fields, particularly in the area of multi-device environments (MDEs) (Rekimoto and Saitoh, 1999; Nacenta et al., 2005; Nacenta et al., 2009; Wallace et al., 2009; Wallace, 2011). This research has yielded many useful cross-device transfer techniques (see Nacenta et al. (2009) for a review). Yet, most of these techniques rely on mouse-based, or otherwise device-aided, input capability that is unavailable in touch-based MSEs. For example, a popular cross-device transfer technique is Rekimoto's (1997) PICK-AND-DROP (P&D) technique, which relies on a digital pen to transfer content from one display to another.

To address this limitation, we conducted a series of three studies to systematically investigate how existing cross-device transfer techniques could be applied or adapted for use in touch-based MSEs. These studies focused on cross-device transfer in a tabletop-centric MSE (T-MSE) context, where a small group of people, each with an individual multitouch tablet, were engaged in a joint activity around a multitouch digital tabletop. The first study examined how two popular cross-device transfer techniques (a "virtual portals" technique (explained below) and the aforementioned P&D technique) could be applied (or adapted) to a "current" T-MSE setup. In this T-MSE, the digital tabletop was unable to distinguish between different users interacting with the tabletop—a limitation of most current multitouch digital tabletops. It therefore posed unique challenges for

cross-device transfer during multi-user interactions. The second and third studies continued the investigation of the P&D technique, further evolving its design adaptation in each subsequent study to better optimize its use for T-MSEs and the specific application task context. The latter two studies focused on a “future” T-MSE set-up that was able to differentiate between users interacting with the tabletop. This capability built on new above-the-surface sensing methods from SurfNet (Genest et al., 2013) and the broader surface computing research community (Hilliges et al., 2009; Pyryeskin et al., 2012; Haubner et al., 2013).

In the remainder of the chapter, we provide an overview of existing cross-device transfer mechanisms, and discuss their limitations for touch-based MSEs. Next, we overview the DOMINION game as the application context for the three studies. We then overview our study methodology, which remained relatively fixed across the three studies. Next, we report each study. Full, detailed versions of Studies 1 and 2 have previously appeared in HCI literature (Scott et al., 2014a; Scott et al., 2014b); thus, only select findings are included in their respective study sections. Study 3 is a previously unpublished follow-up study that investigated design limitations of the P&D adaptation explored in Study 2. Finally, we reflect on insights learned from these investigations and their implications for cross-device transfer in T-MSEs.

Cross-Device Transfer in Multi-Surface Environments

(Components of the background presented here were also reported, in full or in part, in earlier publications on Study 1 (Scott et al. 2014a), Study 2 (Scott et al. 2014b).)

Cross-device transfer is an active area of research in MSEs, and the broader area of multi-device environments. Also, to address reach and ergonomic issues related to dragging digital objects over a large distance, single-surface object transfer techniques have been developed that minimize the need for long drag-and-drop actions. This section overviews these single-surface transfer mechanisms first, followed by the mechanisms used to move content across multiple devices. As all three studies explored the Pick-and-Drop (P&D) technique, this mechanism, and its applicability to touch-based T-MSEs, is discussed in detail.

Object Transfer across Large Surfaces (Within-Device Transfer). Using direct-touch interaction to drag digital content across a large surface has several known ergonomic issues, including fingertip discomfort due to friction and arm and finger fatigue. Moreover, some locations are difficult to reach. Therefore, drag-and-drop extensions have been developed for moving content across large surfaces, including techniques that move an object onto a distant object (e.g. a folder) or location (Baudisch et al., 2003; Hascoët, 2003; Collomb et al., 2005; Collomb and Hascoët, 2008; Doeweling and Glaubitt, 2010). Techniques have also been developed that leverage the physicality of direct-touch surfaces, such as tossing or flicking interaction gestures that use pseudo-physics to “propel” objects to distant

locations (Scott et al., 2005; Weber et al., 2008; Wilson et al., 2008). The aforementioned P&D technique has also been used to transfer objects from one location to another on pen-based interactive wall and tabletop surfaces (Haller et al., 2010). Further, P&D has been shown to be more efficient than drag-and-drop in these contexts (Rekimoto, 1998). Another approach is to move objects from one surface location to another by using “virtual portals,” where an object placed on a portal (typically a virtual interface container or widget) in one location then appears on an associated portal in another location (Besacier et al., 2007; Voelker et al., 2011). The above single surface transfer techniques, especially those designed for direct-touch environments, provide useful inspiration for touch-based cross-device transfer.

Object Transfer across Multiple Devices (Cross-Device Transfer). Existing cross-device transfer techniques broadly fall into three main categories: moving content across contiguous virtual workspaces; moving content via a virtual portal; and moving content via a physical proxy.

Contiguous virtual workspace techniques are based on the physical configuration of displays in the environment. In this approach, displays are connected to a common software architecture that maintains awareness of the physical configuration of the displays (static or dynamic configurations are possible). The display configuration information is then used to provide a contiguous virtual workspace across devices. Thus, moving an object off the edge of one display moves it to the nearest edge of the adjacent display (Rekimoto and Saitoh, 1999; Streitz et al., 2001; Johanson et al., 2002; Hinckley et al., 2004). For example, in PointRight (Johanson et al., 2002), several large screen displays and an interactive tabletop share a single mouse pointer. A static adjacency map, based on the room topology, determines where the pointer moves when it leaves the edge of a screen. In Stitching (Hinckley et al., 2004), an ad-hoc adjacency map is created, with the system inferring the user’s intention to join two adjacent displays when a “stitch” gesture is drawn, starting on one display and ending on a second display. This map can then be used to move digital artefacts between connected tablet computers. Marquardt et al. (2012) propose a similar tablet-to-tablet transfer capability between adjacent tablets, but instead of using a connection gesture they establish the initial ad-hoc connection by tilting one tablet towards the other.

A disadvantage of the contiguous virtual workspace approach for transferring digital objects between a tabletop and a personal surface is the asymmetric size of the displays. The large edges of the tabletop do not map well to the small edges of a tablet or smartphone. The virtual portals technique mentioned above can be used to resolve this issue by providing a dedicated portal area on each device for transferring content (Hinckley et al., 2004; Bachl et al., 2011; Fei et al., 2013). We examined a virtual portals method called BRIDGES in Study 1. The previous two cross-device transfer approaches require people to drag the transferred object to and

from the virtual portal (or display edge) from its origin and to its destination. This can introduce the aforementioned ergonomic issue related to long-distance touch-based dragging. Physical proxy techniques address this intermediary interaction step issue by using a physical object to manage the transfer. They allow for collection and placement of the transferred object directly from its origin to its destination on the respective displays by taking advantage of the three-dimensional space around the displays. This approach involves binding a digital object to a physical object and then moving the physical object to the target display. This typically requires a system-recognized object to facilitate the binding/unbinding process, such as a digital pen (Rekimoto, 1997; Baudisch et al., 2003; Haller et al., 2010; Scott et al., 2014a) or “puck” (Kobayashi et al., 2008). For example, P&D (Rekimoto, 1997) allows someone to “pick up” a digital object at its original location using a digital pen and “drop” it directly at the destination location using the digital pen. This technique evokes the commonly used drag-and-drop concept, and bears strong similarity to the familiar action of lifting and relocating a physical object.

Given the more direct origin-to-destination interaction process, physical proxy techniques like P&D are highly desirable in T-MSEs. They reduce intermediary drag actions across a large tabletop surface, and so, provide more efficient interaction and avoid the ergonomic issue of long distance dragging. Thus, we were highly interested in using P&D in our T-MSE applications. However, the touch-based interaction and the multi-user nature of T-MSEs introduced difficulties for applying P&D in this context; we discuss these issues further below.

Applying PICK-AND-DROP to Touch-based, Multi-User T-MSEs. In touch-based surfaces, no digital pen (or other readily available physical object) is available to serve as the proxy for P&D transfer. In our research, we address this by using the user’s hand as the physical proxy between the tabletop and a personal tablet. This allows someone to “pick-up” the object using a menu or gesture on the tabletop, move their hand to their tablet and then “drop” the object by touching the tablet (and vice-versa). However, in a collaborative T-MSE, multiple people may wish to simultaneously transfer content between various devices. In this situation, the system needs to associate the correct picks with the correct drops, which is only possible if the system knows who is doing what in the environment.

Because people often bring and, exclusively use, their own personal devices in a group setting, a reasonable design strategy in a T-MSE context is to associate a specific user with a specific personal surface (e.g., a tablet, smartphone), and to assume that all interactions with that device are made by that person (i.e., the device “owner”). Using this strategy, we can then assume that all picks or drops on a given personal device are performed by the device owner. Knowing who is doing what on the shared tabletop is more challenging. Indeed, most existing tabletop systems cannot distinguish between different users. Thus, automatically associating picks or drops with

a given person is more difficult, and requires some design adaptation of the P&D technique or additional user-identification system capabilities.

In Study 1, we addressed this issue by providing dedicated “personal territories” along the tabletop edge in front of each group member. Any picks or drops conducted in these territories were associated with the “owning” user, enabling simultaneous, multi-user P&D transfers. In Studies 2 and 3, we addressed this issue by augmenting our tabletop with user-identification capabilities, as detailed in the Study Methodology section below.

Research Approach

Figure 1 summarizes the overall research approach used across the three studies, including the study research questions, the cross-device transfer techniques included in the studies, and the T-MSE environments used in the studies. The figure shows the progression from Study 1’s comparison of two existing cross-device transfer approaches (Bridges virtual portals vs. P&D physical proxy) to Study 2 and 3’s investigation of successive design refinements of a single cross-device transfer approach (P&D) to improve its usability in T-MSEs. Each successive study focused on addressing interaction issues revealed by the previous study. The following section details the specific study methodology that was used in the studies.

	Research Question	Cross-Device Transfer Methods		Tabletop Environment
Study 1	How well do existing transfer methods support T-MSEs?	BRIDGES (virtual portals)	(TERR.-ADAPTED) PICK-AND-DROP (physical proxy)	“Today’s Tabletop” (no user-identification) + Personal Tables
Study 2	Does SURFACE GHOSTS feedback improve effectiveness of P&D in T-MSEs?	PICK-AND-DROP W/ SURFACE GHOST feedback	PICK-AND-DROP W/O SURFACE GHOST feedback	“Tomorrow’s Tabletop” (with user-identification) + Personal Tables
Study 3	Does adding tablet feedback improve transfer awareness during P&D in T-MSEs?	P&D W/ SURFACE GHOST + TABLE feedback	P&D W/ SURFACE GHOST only feedback	“Tomorrow’s Tabletop” (with user-identification) + Personal Tables

Figure 1. Overview of studies conducted to investigate cross-device transfer in T-MSEs.

Study Methodology

(Components of the methodology presented here were also reported, in full or in part, in the earlier publications on Study 1 (Scott et al. 2014a) and Study 2 (Scott et al. 2014b).

All studies utilized a mixed-methods research methodology that involved quantitative and qualitative study measures. All studies were conducted in the same controlled human-computer interaction laboratory environment at the University of Waterloo.

Participants. In all studies, participants were recruited both from the University of Waterloo student and staff population and from local board game stores’ clientele through email lists, social media sites, and posters. To promote natural group behaviour, participants were required to sign-

up with one or two friends, depending on the study, and to have previous experience with the commercial version of the DOMINION game. Table 1 summarizes the participant details for each study.

Study	Number of Participants	Gender of Participants	Age of Participants	Group size	Number of Groups
Study 1	28	23 Male, 5 Female	20-44 (M=27, SD=6.5)	2 people	14
Study 2	18	16 Male, 2 Female	20-38 (M=26, SD=5)	3 people	6
Study 3	24	19 Male, 9 Female	18-59 (M=30, SD=9)	2 people	12

Table 1. Participant details for each study.

Experimental Design. Studies 1 and 2 each included only one independent variable in a single factor (Study 1: transfer technique, Study 2: visual feedback) within-subjects study design, with three levels for each factor in each study. Study 3 included two independent variables in a two-factor, 2 (tablet feedback) x 2 (tabletop feedback), mixed methods design where the tablet feedback was a within-subjects factor and tabletop feedback was a between-subjects factor. This more complex study design is further detailed in the main Study 3 section. Table 2 summarizes the study conditions utilized for each study.

		Within-Subject Factors			Between-Subject Factors	
Study 1	Factor:	<i>Cross-device transfer technique</i>			n/a	
	Factor Levels:	TERR.-ADAPTED PICK-AND-DROP (TA-P&D) w/ IMPLICIT CONTROL	TA-P&D w/ EXPLICIT CONTROL	BRIDGES		
Study 2	Factor:	<i>Tabletop feedback during P&D transfer</i>			n/a	
	Factor Levels:	SURFACE GHOSTS w/ IMPLICIT OWNERSHIP	SURFACE GHOSTS w/ EXPLICIT OWNERSHIP	No Feedback		
Study 3	Factor:	<i>Tablet feedback during P&D transfer</i>			<i>Tabletop feedback during P&D Transfer</i>	
	Factor Levels:	TABLET BRIDGE visualization	No Feedback (No BRIDGE)	n/a	SURFACE GHOSTS w/ IMPLICIT OWN.	SURFACE GHOSTS w/ EXPLICIT OWN.

Table 2. Summary of the main experimental design details used in each study, separated by the respective within- and between-subjects factors.

Experimental Task. The DOMINION Game. DOMINION is a 2-4 player medieval themed card game, in which each player builds their own personal deck to utilize during game play by “buying” cards from a bank of shared card decks. Game play in DOMINION typically occurs on a turn-by-turn basis, though players can take some actions during other players’ turns. In a typical turn, a player draws a minimum of five cards from their deck, and then makes several card-based actions (e.g. revealing (i.e. “playing”) one or more cards to “buy” resources, “attacking” other players (i.e., forcing them to discard cards), or discarding unused cards). Players monitor their opponent’s game actions and may alter their game strategy in response to an opponent’s actions.

To facilitate investigation of cross-device transfer in this game, a custom digital tabletop software application of the DOMINION game was developed that incorporated the use of multiple, portable tablets to provide each player a private digital space (Figure 2). In this digital DOMINION game, cards can be freely moved and rotated using direct touch manipulation.

When two cards are moved to the same position, they are automatically stacked into a deck of cards. A card may be drawn from a deck of cards by touching and dragging the top card, while the whole deck can be moved by dragging its border.




Figure 2. DOMINION digital tabletop system. The Personal Play Area denotes the personal territories used in the BRIDGES condition in Study 1. These colour-delimited areas were omitted in the PICK-AND-DROP conditions in all three studies (from Scott et al. 2014a).

Cards and decks can be flipped via a contextual pie menu invoked by tapping on a card or deck. Decks can also be shuffled with this menu. In both cases, a short animation confirms the action. In the study conditions reported in this chapter that involve PICK-AND-DROP (P&D) transfer, card “picks” were also performed via this menu.

Equipment and Setting. Studies 1 and 2 utilized a custom-built infrared laser light plane (LLP) multitouch digital tabletop with a surface size of 90x130 cm and projected display of 1280x800 pixel resolution (see Figure 2). Study 3 utilized an upgraded custom-built multitouch tabletop incorporating a 4K (3840x2160 pixel) resolution 55-inch flat-panel LED display fitted with a PQLabs infrared multitouch frame. In Study 1, participant pairs sat at adjacent sides of the tabletop. In Study 2, the 3-person participant groups sat at three adjacent sides of the tabletop, with the middle player seated at the long side of the rectangular table. In Study 3, participant pairs sat facing each other at the long sides of the tabletop. This change in seating arrangement from Study 1 was made due to the wide screen configuration of the upgraded table and, consequently, the larger disparity in the length between the long and short sides.

In all studies, participants were each provided a 7-inch Galaxy Tab tablet computer. Tablets were preconfigured to be associated with the player's position at the table to facilitate the cross-device transfer methods under study. Separate laptops were set-up on nearby desks for administration of the study questionnaires. Study questionnaires were administered through the SurveyMonkey® (<http://www.surveymonkey.com>) online data collection service. In all studies, the DOMINION tabletop software application used TUIO multitouch events. In Studies 1 and 2, an infrared camera under the table and the open-source toolkit Community Core Vision (CCV) ([59](http://</p>
</div>
<div data-bbox=)

ccv.nuigroup.com/) were used to process touch. In Study 3, the PQLabs input frame natively produced TUIO data. Finally, in Studies 2 and 3, user identification of tabletop touches and above-the-table arm movements were obtained using a Microsoft Kinect mounted 1.5m above the digital table and an adapted version of the KinectArm toolkit (Genest et al., 2013), as described in Scott et al., 2014b.

Procedure. Participants performed the main study activities together in a group of 2 (Studies 1 and 3) or 3 (Study 2), but completed written forms and questionnaires individually. Upon arriving, participants first completed informed consent forms and a background questionnaire that gathered demographic information and their prior game play experience. They were then given a short demonstration of the experimental hardware systems. Each participant group played three games in a row, one for each study condition. The order of presentation of the three conditions was counterbalanced. In addition, three different sets of ten previously selected Dominion cards were used for the banks of purchasable cards, always presented in the same order to avoid interfering with the counterbalancing of the conditions. Learning effects related to card sets were not anticipated, as all players had previous experience with DOMINION.

Before the first condition, players were given a brief demonstration of the system. In Study 1, each cross-device transfer method was also demonstrated before each condition. In Studies 2 and 3, the P&D technique was only demonstrated at the beginning of the study. Most groups also took 4-5 minutes at the start of each game to read aloud the description of each available card in the bank for the session. After each condition, players completed a post-trial questionnaire about that condition. After the final game and post-trial questionnaire were completed, participants either completed a post-experiment questionnaire (Studies 1 and 2) and/or interview with the researchers (Studies 2 and 3). Finally, participants were thanked and paid for their participation. All three studies were approved by the university's institutional ethics review process.

Data Collection and Analysis. In all studies, quantitative and qualitative data were collected and analyzed. Participant interactions with the tabletop and tablets were captured in computer log files. Video data (with audio) and observer notes captured participants' verbal and non-verbal behaviour during the sessions. Background, post-trial, and post-experiment questionnaires included closed- and open-ended questions. All post-trial feedback questions utilized a 7-point Likert-style rating scale to capture participant perceptions and experiences in each condition.

Different qualitative analysis approaches were used across the three studies, characterizing the diminishing exploratory nature, and increasingly hypothesis driven goals of each successive study. In Study 1, the video data and open-ended participant responses underwent an extensive qualitative analysis, including open coding to reveal interaction and communication

patterns, as well as incidents of confusion or frustration and development of flow diagrams to represent emergent interaction strategies (Beyer and Holtzblatt, 1998), to better understand the advantages and disadvantages of each studied cross-device transfer technique (BRIDGES and TA-P&D). Details of full analysis is reported in McClelland (2013); only relevant themes and insights are reported in this chapter. In Studies 2 and 3, the video and open-ended participant responses were reviewed for patterns and emergent themes to provide context and deeper understanding of the quantitative results.

The Likert scale data from the post-trial questionnaires were analyzed using Repeated-Measures Analysis of Variance (RM-ANOVA). To account for the non-independence of group member responses, group was used as a dependent factor by using seating position at the table as the additional repeated measures factor. Thus a 3 (Condition) x 3(or2) (Seating Position) RM-ANOVA was conducted. As seating position was not expected (and was not found) to significantly impact the study measures of interest (e.g. awareness of cards being transferred, awareness of cards being transferred by a partner), we only report the main effects related to Condition in this chapter. An alpha-value of $\alpha=.05$ was used to determine statistical significance.

Study 1

The goal of Study 1 was to explore the potential of existing cross-device transfer approaches for supporting transfer in a T-MSE. Of the three main approaches discussed above, the contiguous virtual workspace approach was ruled out due to the previously mentioned display size disparity issue in T-MSE settings that can introduce confusion about where objects should be placed or will appear during transfer on the different sized devices. As mentioned, the virtual portals approach resolves this issue by bounding interaction to visible containers in the interface that indicate where object transfer can occur. As the physical proxy approach uses point-to-point transfer, rather than moving objects via the display borders, the display size disparity does not impact its use. Thus, we chose to include a virtual portals technique and a physical proxy technique in the study.

BRIDGES Interaction Design. For the virtual portals technique, we implemented a version called BRIDGES, in which a visible container widget was provided along the tabletop edge in front of each user (called the TABLETOP BRIDGE), and along the top edge of each personal tablet (called the TABLET BRIDGE). For the purpose of the study, the location of the BRIDGES were fixed, and the virtual connection between each user's TABLETOP BRIDGE and their TABLET BRIDGE was established during study set-up and fixed throughout the study. This restriction was deemed appropriate due to the nature of the experimental task—a "sit down" card game. In use cases where users are expected to move around the tabletop, the T-MSE could be augmented with proximity or user-tracking sensors to flexibly allow users' TABLETOP BRIDGES to follow them around the

environment, similar to the proximity-based virtual portals technique used by Fei et al. (2013).

In the context of the Dominion game, when a card was transferred to either the TABLET or TABLETOP BRIDGE the top half of the card would appear on the TABLETOP BRIDGE and the bottom half of the card would appear on the TABLET BRIDGE. Once on the BRIDGE, the card can be moved onto the target device by dragging it off the corresponding BRIDGE, moved back to the originating device by dragging it off the originating BRIDGE, or simply left on the BRIDGES.

Territory-Adapted-Pick-and-Drop (TA-P&D) Interaction Design. For the physical proxy method, we adapted the P&D method to the “current” T-MSE constraints discussed above in a version called TERRITORY-ADAPTED-P&D (TA-P&D). In TA-P&D, the T-MSE was divided into different spatial territories. A personal territory was provided along the tabletop edge in front of each user, a shared territory covered the rest of the tabletop workspace, and a private territory was provided on each person’s personal surface. Each user’s personal territory was virtually connected to their personal tablet (private territory), and this connection remained fixed throughout the study. A “pick” conducted in a user’s personal territory was associated with that user, which allowed them to subsequently “drop” the transferred object on their personal tablet. Similarly, picking up an object from their personal tablet allowed them to drop the object onto their personal territory, without interfering with others’ tabletop interactions.

Within the context of the DOMINION game, tabletop picks were enabled via a context menu that could be opened by tapping on a card (or deck of cards). (While the use of a context menu for initiating the “pick” action was originally due to technique limitations in implementing a “pick-up” grab gesture in our original hardware and software, it turned out that this approach later allowed for in-game efficiencies that were very popular and often requested by our players, such as multi-card pickup menu options, that would have been very difficult to achieve using gesture interaction.) Successive taps on the menu allowed for multiple cards to be picked up and then transferred together to a different location. Cards being transferred could then be dropped either back on the tabletop by tapping in the user’s personal territory or dropped onto the user’s tablet. Dropping the cards on the tablet required a “swipe-down” gesture from the top of the tablet screen (i.e. a downwards drag action) to avoid interference with card manipulation actions. For convenience, if the tablet interface was empty, the user could tap anywhere on the tablet interface to drop transferred cards. A “swipe-up” gesture on the tablet (i.e. an upwards drag action) initiated a “pick” action from the tablet. Several cards could be transferred together performing multiple successive pick actions on the tablet before tapping on the tabletop.




Figure 3 (left). BRIDGES cross-device transfer (Scott et al. 2014a); Figure 4 (right). TERRITORY-ADAPTED-PICK-AND-DROP (TA-P&D) cross-device transfer (Scott et al. 2014a).

Two control mechanisms were implemented to allow a user to temporarily perform pick and drop actions in the shared territory on the tabletop. An **IMPLICIT CONTROL** mechanism allowed a user to touch and hold any empty spot in that person's personal territory to extend it to also cover the shared territory, allowing them to temporarily pick or drop cards directly in the shared territory. An **EXPLICIT CONTROL** mechanism allowed a user to place a digital token labelled, "I Control the Centre" in their personal territory to extend their territory to cover the shared territory, and to allow them to pick or drop cards directly in the shared territory.

Post-Transfer State. An important design consideration when implementing a cross-device transfer mechanism within the context of a card game application like DOMINION is the two-sided nature of the game cards (i.e. each card has a back and front side). The simplest approach—to retain a card's face-up/down state at its originating point when it is transferred—would introduce significant interaction overhead post-transfer. For instance, most cards in the tabletop decks are initially face-down to preserve the secrecy of the card's value. Yet, users are likely to want all cards on their personal tablet to be face-up, as this space is private from others' view (unless they chose to disclose the tablet contents). Thus, the common game action of moving five cards from tabletop decks to one's personal tablet would require significant amounts of tedious turn-over actions after these transfers. Consequently in both the BRIDGES and TA-P&D techniques, cards transferred to a personal tablet were automatically turned face-up, regardless of their originating face-up/down state.

In transfers to the tabletop, the face-up/down state varied by technique. With TA-P&D, a card dropped onto an existing deck was transferred with the face-up/down state of the deck (all cards in a deck had the same face-

up/down state), while a card dropped on an empty workspace area was transferred with a face-up value, to facilitate “playing” the card. With BRIDGES, cards were always transferred from the tablet to the tabletop face-up since the most common player action after such transfer was to “play” a card by revealing its value to other players. An exception to this was if a player had left face-down cards on the BRIDGES after transferring them from the tabletop (we refer to this behaviour later as employing a “partial-transfer” strategy), they could then return it face-down to the tabletop.

Summary of Main Study Findings. Analysis of data from the three study conditions (BRIDGES, TA-P&D (EXPLICIT CONTROL) and TA-P&D (IMPLICIT CONTROL)) revealed that, in general, all conditions sufficiently supported card transfers, as evidenced by the, on average, 322 transfers that occurred per game across the study. The results also revealed a lack of clear preference for transfer method across players. Reported preferences differed drastically between groups, and even between players within groups. For example, one player commented that having the BRIDGES widgets “partly on both screens was beautiful and very helpful”, while another player reported that, in the BRIDGES method, having cards appear “in two places [on both the tablet and tabletop] was a little unwieldy”. Similarly conflicting comments were made about the TA-P&D method: One player reported that “Pick up is a much better mechanic [than BRIDGES]”, while another commented that “Picking up cards was NOT intuitive”.

Question	Pick-and-Drop Implicit Control	Pick-and-Drop Explicit Control	Bridges	RM- ANOVA
	Mean (SD)	Mean (SD)	Mean (SD)	
I had fun playing the game.	5.6 (1.4)	5.9 (1.0)	5.9 (0.9)	F(2,26)=1.00, p=.38
When the other player took action, I always understood their motivations for doing so.	5.4 (1.4)	5.4 (1.8)	5.4 (1.6)	F(2,26)=0.02, p=.98
When taking my turn, I was always aware of my play options.	6.1 (1.2)	6.3 (1.0)	6.3 (0.7)	F(2,26)=0.40, p=.68
I was always aware of the other player’s actions.	4.3 (1.6)	4.6 (1.9)	4.4 (1.8)	F(2,26)=0.41, p=.67

Table 3. Average participant ratings on enjoyment and awareness-related post-condition survey questions from Study 1 (1=strongly disagree, 7=strongly agree).

The RM-ANOVA analysis of the post-condition questionnaires similarly revealed no consistent player preference or perceived utility for any single transfer method. Participant ratings were generally positive on enjoyment and awareness-related measures (with mean ratings of 5.4 to 6.3 out of 7), with no significant differences across conditions (see Table 3).

The qualitative analysis shed light on the lack of clear preference between transfer techniques. It revealed that the effectiveness of a given transfer technique was player- and context-dependent. Preliminary analyses revealed that players in the two TA-P&D conditions rarely, if ever, used either the Explicit or Implicit Control methods for picking and dropping cards directly in the shared territory (i.e. most picks/drops were performed in the players’ personal territories). Thus, both TA-P&D conditions were aggregated into a

single TA-P&D condition for the in-depth qualitative analysis. This analysis revealed several key benefits and limitations of each method that impacted their use: the required cognitive and physical effort, and the ability of the method to maintain the privacy and secrecy of transferred data.

Some players found the TA-P&D transfer method more cognitively demanding than the BRIDGES method since the TA-P&D method required players to mentally keep track of which card(s) they had picked up, and were currently holding, during the transfer process. Once a player picked up a card it would disappear—“in the ether”, as reported by one participant—and was no longer visible on either the tabletop or tablet until the corresponding drop action occurred. While both the tablet and tabletop interfaces provided visual feedback in response to pick/drop actions, such as a short animation on the tabletop after a pick occurred, and the hand-of-cards being rearranged on the tablet after a pick/drop action, these interface changes appeared to be too subtle, or were sometimes occluded from the player’s view. In contrast, cards were always visible on the BRIDGES widgets during the transfer process, eliminating any mental burden from players regarding the state of the cards. Consequently, players reported that it was “easier to keep track of cards” with the BRIDGES method.

Despite its cognitive simplicity, BRIDGES required more physical effort than the point-to-point TA-P&D method. In BRIDGES, players had to drag cards across the tabletop to/from the TABLETOP BRIDGE and to drag cards on/off the TABLET BRIDGE during each transfer. Also, multi-card transfers required multiple drag actions to/from the respective BRIDGES. Thus, some players found transferring cards with BRIDGES to be quite tedious, as evidenced by the player comment, “The hand zone [BRIDGES] was super annoying... It just added more clicks to the game.” In contrast, TA-P&D allowed for multiple cards to be picked up at once and then transferred (and dropped) together.

BRIDGES was also found to be less privacy-preserving than TA-P&D. As mentioned in the Post-Transfer State section, all cards transferred from the tablet to the tabletop in BRIDGES arrived face-up on the TABLETOP BRIDGE to simplify post-transfer game actions, which commonly involved “playing” a card (i.e. revealing its value to opponents). However, at the end of each player turn, players discarded unplayed cards onto the player’s discard deck, typically located in their personal territory. In highly competitive games, revealing the value of discarded cards could reveal a player’s game strategy to observant opponents, potentially reducing a player’s competitive advantage. The “partial-transfer” strategy described early was adopted by some high-competitive players to help preserve card secrecy with the BRIDGES method, but this strategy had limitations that made it unusable for non-expert players (see Scott et al. (2014a) for details). In contrast, the TA-P&D method used the drop context to determine the face-up/down state of transferred cards. Thus, the secrecy of the card values dropped onto a face-down deck, such as the discard deck, would be preserved.

In Summary, while the BRIDGES method provided simple, straightforward usability that provided persistent feedback of the transferred cards, it also was less physically efficient and did not preserve the privacy of transferred objects as well as the TERRITORY-ADAPTED PICK-AND-DROP transfer method. As privacy is often an important goal of providing personal surfaces in a multi-surface environment, and efficiency of an interaction method is always an important usability goal for interaction techniques, we chose to investigate the TA-P&D method further in subsequent studies. More specifically, these follow-up studies focused on reducing the cognitive effort required to use this method for content transfer.

Improving P&D Transfer with SURFACE GHOSTs Visual Feedback

A common HCI approach for helping people understand ongoing changes in a computer system is to provide persistent visual feedback related to changes in system state (Smallman and St. John, 2003; Scott et al., 2006; Chang et al., 2014). Study 1 revealed that the brief visual feedback provided after a card was picked up on either the tabletop or tablet was insufficient. During the actual transfer stage, no visual feedback was provided to indicate that cards were being “held” by the user. Thus, if someone became distracted after picking up a card—for instance, by an opponent’s game play actions or an ongoing conversation—they might forget they were holding a card and hence be surprised when the card appeared in the interface when subsequently touching the tabletop or their tablet.

Changing a virtual object’s visual appearance has been previously used to indicate changes in object state. For example, in Rekimoto’s (1997) original P&D implementation, when the digital pen hovered over the target display (within millimetres), the transferred object was displayed with a virtual shadow cast underneath it. This object-with-shadow representation would follow the hovering pen around in the interface until the object was dropped on the display, and then the shadow would disappear, leaving the active object. Similarly, “shadow” or “silhouette” object representations have been used to indicate objects being copied across adjacent tablet devices (Hinckley et al., 2004) and objects being held above the tabletop in a 3-dimensional tabletop workspace (Hilliges et al., 2009). Based on this prior work, we hypothesized that showing a similar visual representation of transferred cards in the interface during the transfer process may help reduce the cognitive effort associated with using our touch-based P&D transfer method. We also felt that providing feedback on who was transferring which cards would further reduce any user confusion in our multi-user setting. So, we designed the SURFACE GHOST object representation to provide visual feedback of cards being transferred with our touch-based P&D transfer method.

In Study 1, players tended to position their tablets directly along the tabletop edge. Thus, cross-device transfer interaction occurred largely over the tabletop surface. Therefore, we hypothesized that displaying visual feedback of transferred objects on the tabletop as the objects are

carried over the tabletop surface should provide (sufficiently) persistent visual feedback during transfer. Accordingly, the SURFACE GHOST visual feedback was designed to appear in the tabletop interface underneath the “owning” user’s hand as it traveled across the tabletop surface between the originating pick location and the target drop location. SURFACE GHOSTs were displayed as semi-transparent, greyscale versions of transferred objects. When multiple objects were being transferred at once, they were stacked together and a counter displayed the total number of transferred objects. Figure 5 illustrates the SURFACE GHOST visual designs for single-object (c) and multi-object (d) transfers in a digital card game.

To accommodate concurrent multi-user card transfers, the SURFACE GHOST design also conveyed ownership of the transferred object(s) through a number of static and dynamic design features. The basic SURFACE GHOST design provided several implicit indications of ownership: upon pick up the SURFACE GHOST object would “fly” (via a brief animation) toward its owner, the SURFACE GHOST object was oriented toward its owner, and it was displayed in real-time beneath the owner’s hand as their hand moved across the tabletop surface. As we were unsure how apparent such ownership information needed to be in DOMINION game setting, we developed two versions of the SURFACE GHOST design. The IMPLICIT OWNERSHIP version provided the above ownership information along with another, still subtle, indication of ownership; a large dark arrow attached to the bottom of the SURFACE GHOST object that “pointed” to the owning user (Figure 5, c and d). The EXPLICIT OWNERSHIP version replaced the black arrow with a more visually salient representation of the owner; a semi-transparent white silhouette of the owner’s arm displayed on the tabletop beneath the user’s physical arm. The SURFACE GHOST object was positioned at the arm silhouette’s hand (Figure 6), indicating that that user was “holding” the card.

To implement either of these SURFACE GHOST designs, it was necessary to move beyond a “current” T-MSE set-up to a “future” T-MSE set-up that provided multi-user identification and above-the-surface tracking. In this enhanced environment, the system was able to keep track of who was transferring which cards. Thus, it was no longer necessary to divide the tabletop into personal and shared territories to facilitate simultaneous multi-user P&D transfers. Hence, in Studies 2 and 3, users could perform pick or drop actions at any location on the tabletop interface. So, we dropped the “Territory-Adapted” aspect of our P&D implementation, and refer to the technique as simply P&D transfer when discussing the method used in Studies 2 and 3 rather than TA-P&D. As user-tracking was limited to the area above the tabletop surface, any pick or drop actions on a tablet were still assumed to belong to the “owning” user.

Study 2

The goal of Study 2 was to determine whether the SURFACE GHOSTs visual feedback reduced the confusion people experienced during the P&D

transfer process, and improved their awareness of cards being transferred. We were also interested in learning whether the SURFACE GHOSTs feedback improved people’s awareness of when other players were transferring cards during the game, thereby improving their collaborative awareness. Given the multi-user nature of our task environment, another goal was to determine how the two different ownership designs (IMPLICIT vs. EXPLICIT) might impact people’s awareness of transferred objects, and overall transfer performance.




Figure 5 (left). SURFACE GHOSTs in a card game context: (a) a normal card, (b) a deck of cards, (c) a SURFACE GHOST (with IMPLICIT OWNERSHIP feedback) of one card being transferred by the Left Player, and (d) a SURFACE GHOST of multiple cards being transferred by the Bottom Player (from (Scott et al 2014b)).
Figure 6 (right). SURFACE GHOSTs with EXPLICIT OWNERSHIP in a tabletop card game context (from (Scott et al 2014b)).

Following the methodology described above, groups of three participants completed three DOMINION game play sessions using the P&D transfer method with the three different visual feedback conditions: SURFACE GHOSTs with IMPLICIT OWNERSHIP (IMPLICITSG), SURFACE GHOSTs with EXPLICIT OWNERSHIP (EXPLICITSG), and a control condition no feedback (NF).

Summary of Main Study Findings. Similar to Study 1, players performed a significant amount of P&D transfers during the study. A total of 4455 P&D transfers occurred across all game sessions. Similar to Study 1, participants’ preferences were evenly split across the three conditions (6 preferred IMPLICITSG, 6 preferred EXPLICITSG, 6 preferred the control (NF)). Despite the fact that a third of the participants preferred the control (NF) condition, the RM-ANOVA analysis of the post-condition questionnaires revealed that both SURFACE GHOST conditions significantly increased reported awareness of transferred cards compared to the control (NF) condition for tabletop-to-tablet transfers. Yet, the analysis revealed that the SURFACE GHOST feedback did not provide the same awareness benefits for card transfers in the opposite direction (tablet-to-tabletop transfers). No differences were found in reported awareness levels between the two SURFACE GHOST conditions in either transfer direction. Similarly, no differences were found in reported awareness levels of card transfers performed by others at the table across all conditions. Table 4 summarizes the reported awareness levels across conditions and the RM-ANOVA results.

Question	No Feedback Mean (SD)	Surface Ghost w/ Implicit Ownership Mean (SD)	Surface Ghost w/ Explicit Ownership Mean (SD)	
<i>I was always aware...</i>				RM- ANOVA
<i>...when I had a card in my hand when moving from the tabletop to my tablet</i>	4.4 (1.9)	5.7 (0.8)	5.9 (1.2)	F(2,10)=8.16, <i>p</i> =.008* *Sig. Contrast: NF vs. SGwImp: <i>p</i> =.028 *Sig. Contrast: NF vs. SGwExp: <i>p</i> =.024
<i>...of how many cards I had in my hand when moving from the tabletop to my tablet</i>	4.6 (1.9)	5.8 (1.2)	5.9 (1.1)	F(2,10)=8.71, <i>p</i> =.006* *Sig. Contrast: NF vs. SGwImp: <i>p</i> =.026 *Sig. Contrast: NF vs. SGwExp: <i>p</i> =.027
<i>...when I had a card in my hand when moving from my tablet to the tabletop</i>	5.5 (1.5)	5.7 (1.0)	5.8 (1.0)	F(2,10)=0.56, <i>p</i> =.59
<i>...of how many cards I had in my hand when moving from my tablet to the tabletop</i>	5.4 (1.2)	6.1 (0.8)	6.0 (1.2)	F(2,10)=8.38, <i>p</i> =.007* *Sig. Contrast: NF vs. SGwImp: <i>p</i> =.007
<i>...of when my partner had cards in transit</i>	3.8 (1.9)	4.4 (1.9)	4.3 (1.8)	F(2,10)=1.56, <i>p</i> =.26

*significant at $\alpha=.05$.

Table 4. Average participant ratings on awareness-related post-condition survey questions Study 2 (1=strongly disagree, 7=strongly agree).

Our results show that both SURFACE GHOST designs were more effective at promoting awareness of transferred objects during transfers originating on the tabletop than transfers originating on the tablet. The qualitative data analysis provided insights on this asymmetric awareness benefit of SURFACE GHOSTs by revealing how participants used this visual feedback. SURFACE GHOSTs were found to support three main aspects of P&D transfer: confirming that a pick or drop worked, keeping track of how many cards were picked up, and confirming that picked up cards went to the right player.

Confirming that an intended pick or drop action succeeded was the most prevalent use of the SURFACE GHOST feedback. Players frequently used the local animation of the SURFACE GHOST object “flying” from the card’s original location toward the owning user to confirm picks. Also, players commonly shifted their hand and wrist positions during pick actions to facilitate viewing the SURFACE GHOST object located under their palm (which was more robustly tracked than their fingertip) or arm silhouette during this pick confirmation process. Similarly, players often double-checked that the SURFACE GHOST feedback disappeared after a drop operation. In the control (NF) condition, the lack of feedback often resulted in participants redoing a whole sequence of actions.

Players also made extensive use of the counter provided in the multi-object SURFACE GHOST design to track how many cards they had picked up during multi-card. In the control (NF) condition, players relied on counters attached to each deck to determine how many cards they had picked up, by tracking how much the number decremented after each pick up. This method was more cognitively demanding, as revealed by Study 1. In contrast, the SURFACE GHOST counter provided the information directly, without mental calculation, and was available if players missed the original pickup actions.

The third main use of SURFACE GHOSTs was to confirm that cards on the tabletop were picked up by the right person. Due to technical limitations of above-the-table tracking, the system’s user identification was occasionally

incorrect when players were interacting in close proximity. When this occurred during pickup, the card(s) would be associated with the wrong user. As part of the pick confirmation behaviour described above, players commonly relied on the local animation of the SURFACE GHOST object to confirm the correct user association. Figure 7 illustrates an example where the SURFACE GHOST object animation helped participants to detect an incorrect association, during simultaneous proximal interactions. If this animation was missed, the various forms of persistent and dynamic feedback provided by SURFACE GHOSTS was also useful: the dynamic movement of the SURFACE GHOST object following a user's hand, and in particular, in the case of EXPLICITSG, the arm silhouette, was reported to be particularly useful in diagnosing inaccurate user identification.




Figure 7. Players using SURFACE GHOST animation to recognize that a picked card went to the wrong player: a) Left Player waits to interact near Right Player's hand b) Right Player's menu (highlight) is oriented toward Left Player due to inaccurate user identification, yet Right Player does not appear to notice, c) Right Player picks up card, and d) the Surface Ghost (highlight) flies toward Left Player's hand. Right Player says, "I am under the impression that you might have my cards". From (Scott et al 2014b)).

The qualitative analysis also revealed that participants unexpectedly appropriated the P&D technique for transferring cards between different tabletop locations, rather than using drag-and-drop transfer. All participants in all conditions exhibited this behaviour, even though they were only shown how to use P&D for cross-device transfers. They spontaneously, often accidentally, discovered this possibility during game play. Distance did not seem to be a main factor for triggering within-tabletop transfers: the same players were observed using drag-and-drop to transfer cards over long distances, and using P&D to transfer cards over very short distances.

In summary, we found that SURFACE GHOSTS feedback successfully promoted transfer awareness during tabletop to the tablet transfers, but was less effective during tablet-to-tabletop transfers. The lack of improved awareness during transfers originating on the tablet was likely caused by the lack of SURFACE GHOSTS feedback during tablet pick and drop actions, due to the positioning of the tablets outside the active tabletop area. The study also revealed that both IMPLICIT and EXPLICIT OWNERSHIP design variations provided sufficient ownership information in most transfer situations, yet the arm silhouettes provided by the EXPLICIT OWNERSHIP design provided better support for coping with common technical issues encountered on multi-touch surfaces, minimizing frustration and improving the overall user experience. The final study focused on increasing transfer awareness during tablet-to-tabletop transfers, thereby improving the overall

cross-device transfer experience.

Improving Awareness during Tablet-to-Tabletop P&D Transfers

The fact that the SURFACE GHOSTS feedback was unavailable during pick operations on the tablet was only a minor issue when transferring a single card: the SURFACE GHOSTS feedback would appear as soon as the user’s hand was over the tabletop, and so, visual feedback was available almost immediately after the pick operation. However, during a multi-card pick up sequence that required the user to make repeated pick operations (recall, a tablet pick operation involved dragging a card upwards across the top edge of the tablet using a swipe-up gesture); each successive pick operation would bring the user’s hand repeatedly back over the tablet surface (and away from the tabletop surface). Thus, this interaction sequence delayed the appearance of the SURFACE GHOST feedback until the final card had been picked up. Consequently, the user had to rely on (sometimes subtle) changes in the arrangement of cards in the tablet interface to confirm the success of the pick operation, which was easy to miss if the tablet contained a number of visually similar cards.

To address the ineffective feedback on the tablet, we considered various design solutions. We first considered a variant of SURFACE GHOSTS on the tablet, but found it had several drawbacks. The first issue was technical: tracking a user’s hand above a tablet—especially when players moved their tablet—was highly challenging and not feasible in our tracking environment. Second, there was limited screen real-estate to display a useful SURFACE GHOST object or arm silhouette. Also, it would likely be obscured from the user’s view by their physical hand, or positioned off the display. Thus, we wanted to provide a device-appropriate feedback mechanism that would serve the same purpose as SURFACE GHOSTS on the tabletop: convey which cards, and how many cards were currently being held by the user. A consistent feedback from Study 1 was that the visual feedback provided by the BRIDGES mechanism provided high levels of transfer awareness. Also, the location of the TABLET BRIDGE coincided with the swipe-up and swipe-down gestures for tablet pick and drop actions. Thus for Study 3, we included a modified version of the TABLET BRIDGE visualization (without the BRIDGE transfer functionality). Unlike the split card visualization in Study 1, in Study 3 we displayed miniature versions of entire cards along the top edge of the tablet during transfer (see Figure 8).



Figure 8. The modified TABLET BRIDGE visualization. When cards are dropped on tablet, miniature cards disappear from TABLET BRIDGE and appear full size in main tablet interface below.

We also made improvements to the SURFACE GHOSTS tabletop feedback and to the overall P&D interaction process to better support the DOMINION task environment. First, we fixed an interaction bug revealed during Study

2 (detailed in Scott et al. (2014b)) that interfered with touch actions over the arm silhouette in the EXPLICITSG condition. We also displayed a second counter on the lower left corner of the SURFACE GHOST multi-card visualization to improve visibility of the counter. Finally, we added an option to pick up 5 cards at once to the context menu to facilitate this frequent Dominion game action. Figure 9 shows the updated SURFACE GHOST designs on the tabletop used in Study 3.




Figure 9. The updated SURFACE GHOST visual feedback and tabletop environment: EXPLICITSG (left), and IMPLICITSG (right).

Study 3

The primary goal of Study 3 was to determine whether the combination of SURFACE GHOSTS feedback on the tabletop and TABLET BRIDGES feedback on the tablets improved player’s overall awareness during P&D transfer, in both transfer directions. A secondary goal of Study 3 was to determine whether our software improvements resolved transfer performance issues observed in the EXPLICITSG condition in Study 2.

To reflect these primary and secondary goals, we modified the study method used in Studies 1 and 3. To address our primary goal of comparing the effectiveness of adding the TABLET BRIDGE feedback, we included a tablet feedback factor with two levels: BRIDGE and NO BRIDGE conditions. To address the secondary goal of assessing the timing performance of the modified EXPLICITSG design, we included a tabletop feedback factor: EXPLICITSG and IMPLICITSG conditions. Due to practical concerns involved with playing full-length DOMINION games in each study condition, we chose to use a mixed within-subjects (tablet feedback) and between-subject (tabletop feedback) experimental design, rather than a fully crossed, within-subjects design, to minimize participant fatigue. Also, as Study 3’s main measures related to the tablet feedback factor focused on player’s perceived awareness of their own transferred cards, for practical issues, we utilized a participant group of size two (similar to Study 1).

Each group completed three DOMINION game play sessions using the P&D transfer method under three different visual feedback conditions. All groups experienced the EXPLICITSG tabletop feedback both with BRIDGE tablet feedback (EXPLICITSG+B) and with NO BRIDGE tablet feedback (EXPLICITSG+NB), and experienced either the IMPLICITSG tabletop feedback with BRIDGE tablet feedback (IMPLICITSG+B) or with NO BRIDGE tablet feedback (implicitSG+NB). Thus, each group only played one condition with the IMPLICITSG feedback on the tabletop (with or without

the TABLET BRIDGE visualization). Our data analysis of the awareness metrics only included data from the EXPLICITSG conditions to enable more statistically robust repeated-measures analysis of questionnaire responses, while our data analysis of the transfer timing metrics utilized both within- and between-subjects analyses across conditions, as described below, due to the more numerous occurrences of card transfers available from the interaction logs.

Study Findings. The data analysis revealed that participants had a strong positive reaction to the addition of the TABLET BRIDGE visualization. Twenty-two out of 24 participants preferred having the BRIDGE feedback on the tablet (18 preferred EXPLICITSG+B; 4 preferred IMPLICITSG+B), while the remaining two preferred the NO BRIDGE conditions (1 preferred IMPLICITSG+NB, 1 preferred EXPLICITSG+NB). According to participant's post-experiment interview comments, the two preferences for the NO BRIDGE condition was influenced by a minor interaction difference between the BRIDGE and NO BRIDGE conditions: the "tap anywhere to drop" convenience feature when the tablet was empty was missing in the BRIDGE conditions due to inherited functionality from the Study 1 BRIDGES transfer method (unfortunately not identified during pilot testing). However, the lack of this feature was not mentioned by most participants, who appeared to prefer using the swipe-down drop gesture. For the remaining few who also commented on this missing feature, their overall preference for the BRIDGE condition appeared to be strongly influenced by the high level of transfer awareness it provided.

The data analysis also revealed that providing the TABLET BRIDGE feedback significantly improved participants reported awareness of transferred cards, in both transfer directions. Also, analysis of the transfer timing data found no differences between EXPLICITSG and IMPLICITSG conditions, suggesting that our software modifications addressed the transfer time performance issues related to the EXPLICITSG design uncovered in Study 2. As the timing investigation was included to validate our software implementation improvements rather than our transfer method interaction concept, timing results are not included here, but are detailed in an online technical report (Scott et al., 2015). We expand on the transfer awareness results below.

Perceived Awareness of Transferred Cards. The RM-ANOVA analysis of the post-condition questionnaire responses from the two EXPLICITSG conditions revealed the BRIDGE (EXPLICITSG+B) condition significantly increased reported transfer awareness compared to the NO BRIDGE (EXPLICITSG+NB) condition for both tabletop-to-tablet and tablet-to-tabletop transfers. Table 5 summarizes the reported transfer awareness data and RM-ANOVA results. (Comparing two conditions would normally call for a t-test statistic, but recall from the Methodology section that tabletop position was also included as a main between-subjects factor in all RM-ANOVA analyses across all studies to account for the effect of group. No effect of tabletop position or interaction across main factors was found.)

Question	SURFACE GHOST (Exp) w/o TABLET BRIDGE	SURFACE GHOST (Exp) w/ TABLET BRIDGE	RM- ANOVA
	Mean (SD)	Mean (SD)	
<i>I was always aware...</i>			
...when I had a card in my hand when moving from the tabletop to my tablet	4.7 (1.7)	5.8 (1.0)	F(1,11)=9.44, $p=.011^*$
...of how many cards I had in my hand when moving from the tabletop to my tablet	4.8 (1.6)	5.7 (1.3)	F(1,11)=6.29, $p=.029^*$
...when I had a card in my hand when moving from my tablet to the tabletop	5.1 (1.5)	6.0 (1.0)	F(1,11)=10.65, $p=.008^*$
...of how many cards I had in my hand when moving from my tablet to the tabletop	4.8 (1.9)	5.7 (1.3)	F(1,11)=8.93, $p=.012^*$

*significant at $\alpha=.05$.

Table 5. Average ratings on awareness-related post-condition survey questions (1=strongly disagree, 7=strongly agree).

These results supported our expectation that the BRIDGE condition would better promote transfer awareness than the NO BRIDGE condition for tablet-to-tabletop transfers. Yet, they contracted our expectation that the BRIDGE and NO BRIDGE conditions would provide similar support for transfer awareness for tabletop-to-tablet transfers, given the effectiveness of the SURFACE GHOSTS feedback alone to support transfers in this direction in Study 2. Thus, the BRIDGE condition appeared to effectively promote transfer awareness in both transfer directions. This result was confirmed by the many positive comments participants made regarding the utility of the TABLET BRIDGE in response to the open-ended survey question, "What feature of the tabletop/tablet assisted the game play?", including: "The visualization of cards at the top of the tablet greatly improved my awareness of when I had cards in transit." (P15 EXPLICITSG+B); "You could see the cards on the tablet that were in transit." (P5 EXPLICITSG+B); "The cards appearing on the tablet when in transit was helpful" (P24 EXPLICITSG+B); and, "Not seeing the cards in transit on the tablet was a hindrance." (P11 IMPLICITSG+NB).

Review of the interview, open-ended questionnaire responses, and video data also provided insights on the unexpected positive influence of the TABLET BRIDGE feedback on transfers originating from the tabletop. Participants reported extensive use of the TABLET BRIDGE feedback, when available, during tabletop-to-tablet transfers, as illustrated by the following comments: "The little bar on the tablet at the top to show what cards you took to the tablet [assisted the game play]." (P22 EXPLICITSG+B questionnaire); "Sometimes you thought you picked up 5 cards when really you hadn't, and hav[ing] that additional feedback on the tablet was nice." (G7 interview); and, "In the second game they [cards on top of the tablet screen] disappeared...It was much more clear what you were transferring from the table to your tablet when you had them up at the top." (G1 interview).

The video data revealed several specific benefits of the TABLET BRIDGE feedback during tabletop-to-tablet transfers. During DOMINION game play, players make extensive use of the "personal territory" near them in the tabletop interface. Unlike Study 1 where personal territories were explicitly delimited in the interface, in Studies 2 and 3, players implicitly established these territories, similar to common tabletop usage in other contexts (Scott

and Carpendale, 2010). The consequence of this territorial behaviour is that pick and drop actions often occur near the tabletop edge, commonly causing the SURFACE GHOST visual feedback to be displayed partially outside the interface. Due to poor touch detection near the tabletop edge on the tabletop system used in Studies 1 and 2, the active game play area in the Dominion tabletop application stopped a few centimeters from the edge. However, since the projected area covered the whole surface, the SURFACE GHOSTS object and arm silhouette visual feedback continued to be displayed in the edge area. The upgraded tabletop used in Study 3 provided improved touch detection across the whole surface. So, the active play area was extended directly to the tabletop edge to facilitate easier player's access to game content. An unintended consequence of this change was that the SURFACE GHOSTS feedback was sometimes unavailable during pick/drop actions near the table edge. Participants used the TABLET BRIDGE feedback, when available, to overcome this issue.

The TABLET BRIDGE feedback also helped compensate for the positioning lag of the SURFACE GHOST and arm silhouette caused by necessary image smoothing performed on the imperfect Kinect tracking data. Once participants became familiar with the P&D transfer mechanism, they could perform card transfers very quickly. Thus, sometimes a transfer was almost (or completely) finished before the SURFACE GHOST feedback would appear. In contrast, the TABLET BRIDGE was immediately, and persistently, available throughout the transfer process. Additionally, the new option to pick up 5 cards at once from a tabletop deck was used extensively. This substantially reduced the need for one-by-one multi-card pick-ups, which, in turn, reduced participants' use of the pick-up counter on the SURFACE GHOST multi-card visualization.

Finally, the TABLET BRIDGE feedback also helped participants cope with hardware input errors, such as errors in touch or gesture detection on the tabletop and tablet devices or errors in user tracking on the tabletop. Participants found the additional visual feedback on the tablet helpful for detecting and managing these issues, as illustrated by the comments, "The slight finicky-ness [of the tabletop touch detection] was still a problem, but was helped by the display of cards being transferred at the top of the tablet ." (P23 EXPLICITSG+B questionnaire) and "[I] Felt the sensor wasn't working as well as the first game [a BRIDGE condition]. This could have been due to having less feedback when I picked up a card. Would have been nice to know how many cards were in transition." (P13 EXPLICITSG+NB questionnaire).

Summary. The study found that providing both TABLET BRIDGE feedback on the tablet and SURFACE GHOST feedback on the tabletop improved transfer awareness for both tablet-to-tabletop and tabletop-to-tablet transfers, thereby improving the overall utility of our T-MSE P&D transfer technique. The immediate and persistent feedback provided by the TABLET BRIDGE feedback helped compensate for several technical and usability

issues of the SURFACE GHOST mechanism.

Discussion

Our three studies provided significant insights on supporting cross-device transfer in T-MSE settings. The studies also highlighted how point-to-point cross-device transfer techniques like P&D can be appropriated for within-surface transfers to help ameliorate usability issues related to dragging objects, especially across long-distances, on devices with imperfect touch input technologies (e.g. dropped objects due to lost or jittery input). We discuss these lessons learned below.

Make Object State Apparent through Entire Transfer Process. The results of Study 1 uncovered the need for visual feedback during P&D transfer. However, Studies 2 and 3 highlighted the specific need for feedback during the pick and drop actions of the three-phase P&D process (pick, transfer, drop). The limited visual feedback available on the tablet during pick operations in Study 2 hindered participants' perceived awareness for transfers originating on the tablet. Introducing the TABLET BRIDGE visualization (without associated BRIDGES portal functionality) in Study 3 provided persistent feedback during the entire P&D process: users could immediately see each picked card added to the row of miniature cards displayed on the TABLET BRIDGE, and see them disappear when cards were dropped on the target device. For tabletop-to-tablet transfers in Study 3, players could utilize either the SURFACE GHOST feedback on the tabletop or the TABLET BRIDGE feedback on the tablet to learn the state of cards involved in the transfer process, providing redundant feedback (when the SURFACE GHOST feedback was available on the tabletop). The BRIDGES transfer method from Study 1 provided similarly redundant feedback throughout the entire transfer process. Both the TABLE and TABLET BRIDGES displayed all cards being transferred (across a pair of devices), and at no time did cards disappear from view—they were either on the tabletop/tablet as full-size active cards, or they were visible on the TABLETOP/TABLET BRIDGES transfer portals. Not surprisingly then, Study 1 participants consistently reported high levels of transfer awareness in the BRIDGES condition.

Consider Efficiency at All Stages of Transfer: Beginning, Middle, and End. While the BRIDGES transfer method provided excellent awareness of transferred objects, it was also found to be extremely tedious to use in the DOMINION task context, which required frequent object transfers. The fact that each transfer operation required interaction to/from the intermediary BRIDGES containers added additional interaction steps to the overall transfer process. Participants found this to be especially effortful when performing multi-card transfers, of which there were many during the DOMINION games.

The point-to-point nature of P&D transfer allowed for more efficient transfer, especially as our implementation allowed for multiple cards to be

picked-up at the originating location and transferred at once. However, the frequent need in DOMINION to pick-up multiple (most often 5) cards each turn, introduced room for improved efficiency at the beginning of a multi-card transfer process. Indeed, the “pick up 5 cards” option added to the tabletop menu in Study 3 was highly appreciated, and utilized, by players. Allowing aggregated card transfer in the BRIDGES transfer method may be similarly useful for improving its efficiency, for instance, by allowing a deck of cards to be placed on the BRIDGES. This approach raises the design issue of whether the aggregated content (e.g. 5 cards) should be shown separately or in aggregated form on the BRIDGES containers. In Study 3, the TABLET BRIDGES visualization used the former approach: all transferred cards were displayed separately. Using this “show all” approach, users could then remove individual items “from the BRIDGE” on the target device, or could be given a mechanism (a gesture or button) to allow items to be removed together. Displaying an aggregated view would only allow for an all-at-once end-of-transfer action, and may also reduce some of the positive awareness benefits of the BRIDGES method.

P&D transfer outperformed BRIDGES for end-of-transfer efficiencies as multiple cards being transferred at once would all drop at the target location. The “tap to drop” convenience feature on the tablet (available when the tablet was empty) also improved the drop efficiency of P&D transfer over the “swipe-down to drop” interaction, as it was more forgiving due to the bigger interaction target of the whole tablet screen (vs. the top edge for the swipe-down action) and to the more robust touch detection in the central area of the tablets used in the studies. As mentioned above, end-of-transfer interaction, especially on the tablet, could be improved by providing a mechanism to allow all transferred items to be moved off the BRIDGE at once. This should be done in a task- and device-relevant way, for instance, in the DOMINION game, the TABLET BRIDGE could be augmented with a button located to one side that, when pressed, incorporated all content on the BRIDGE into the hand of cards on the tablet. This would be fairly simple, as there was only one possible destination for cards fully-transferred to the tablet. In contrast, automatically offloading the TABLETOP BRIDGE would be more complex on the tabletop, as the intended destination may be less clear. Here, a specific drag gesture (e.g. a 2-finger drag) that allows players to manually move the entire contents of the BRIDGE to the intended location may be more appropriate.

In Study 3, it was anecdotally observed that some participants misinterpreted the TABLET BRIDGE visualization to mean that cards picked up on the tabletop were automatically transferred to the tablet. This misperception was actually a commonly suggested improvement across the three studies, and one we have received from others during public demonstrations of our system. This approach would resolve many efficiency issues discussed above. However, the approach assumes that players always intend to move cards to their tablet. Yet, our studies revealed frequent use of tabletop-to-tabletop transfers, thereby introducing complexities for inferring when

cards should be transferred to one's tablet rather than be moved elsewhere on the tabletop. Nonetheless, the approach warrants further investigation as it has the potential to greatly improve the efficiency of tabletop-to-tablet transfers.

Consider Post-Transfer State, Utilize Context if Available. Another limitation of the BRIDGES method is its inability to infer the target location, and hence intended purpose, during tablet-to-tabletop transfers. Consequently, the same post-transfer state was applied to each transferred card: Cards were always transferred face-up onto the TABLETOP BRIDGE to facilitate the common "reveal a card" action. However, this design decision was not universally appreciated. The inability to control the post-transfer card state with BRIDGES prompted highly competitive players to adopt a "partial transfer" strategy, in which they left drawn cards sitting on the BRIDGES. This allowed them to keep cards face-down on the tabletop at the cost of not being able to fully view, or manipulate, cards on the tablet. These players strongly preferred the context-dependent manner of determining the post-transfer state used by the TA-P&D (and P&D) transfer method: Cards transferred to the tabletop took the face-up/down state of any deck/card they were dropped onto, or were placed face-up if dropped onto an empty area. This design decision was driven by the application task (i.e. the DOMINION game) and an early analysis of common game actions (and associated player intentions).

In the DOMINION game, the possible states of transferred objects were relatively limited: cards and card decks were the only application objects, card size and orientation were fixed on both the tabletop and tablet, and cards were either face-up or face-down. (In Study 1, orientation of cards (and decks) on the tabletop was automatically determined by whether they were located in a personal territory or the shared territory. In Studies 2 and 3, cards (and decks) were automatically (orthogonally) oriented toward the table side of the "owning" user after P&D transfers or drag actions.) However, in other task contexts, the possible object states that should be considered after transfer will vary, and may include, for instance, the scale (size) and orientation of content objects, or whether they are separate or aggregated, and for multi-dimensional objects, what side (or sides) is displayed. The size disparity between a large surface and smaller personal surface may play a factor. For instance, if a document that is currently being viewed on a smartphone display is transferred to a shared tabletop, it may be useful to display a larger portion of the document on the larger tabletop display than was visible on the smaller smartphone. Ultimately, if post-transfer state is determined automatically by the system, it should select a task- and device-appropriate state that best facilitates people's intended task activities. The selected state should optimize the overall efficiency of the transfer process by minimizing any necessary interactions to achieve a desirable post-transfer object state. Any contextual information available about the intended target location, transfer direction, task phase, etc. may be helpful in inferring a reasonable post-transfer state.

Consider Within-Surface Transfer on Large Surfaces. The studies revealed the common use of P&D transfer to move cards from one tabletop location to another. Almost all participants performed such tabletop-to-tabletop transfers. Analysis of the interaction logs for Study 2 showed no consistent pattern of participants' use of P&D transfer compared to drag actions related to the move distance: P&D transfers appeared to be as equally likely to use for short-distance tabletop moves as for long-distance moves. The video data revealed several possible motivations for choosing P&D over drag to move a card/deck on the tabletop. First, participants often appeared reluctant to drag cards/decks directly over other cards/decks, possibly due to uncertainty over the consequence of such actions (i.e. the deck/card may be disturbed). Thus, they sometimes dragged cards in a wide path around other tabletop content, or simply used P&D transfer to go above the tabletop content. Second, the imperfect touch detection on the tabletop sometimes caused the touch input to fail and cards to drop onto other content. One such instance in Study 3 prompted the user comment, "the deck just swallowed my cards". This type of input errors, unfortunately all too common in existing large-surface hardware, creates significant frustration for users. Long-distance drags are particularly vulnerable to lost-touch situations. The fact that the P&D transfer method required minimal touch interaction on the tabletop provided a reasonable coping strategy for moving content, especially across a long distance, giving the tabletop's imperfect touch detection.

Conclusions

Our studies investigating cross-device transfer demonstrated how the existing cross-device transfer methods virtual portals (BRIDGES) and physical proxy (PICK-AND-DROP) can be applied to both "current" and "future" table-centric multi-surface environments. The studies revealed both methods, with our adaptations optimized for touch-based devices, effectively supported the significant amount of transfer required by the experimental task (the DOMINION card game). They also revealed several key interaction design requirements for cross-device transfer, including the need for persistent feedback throughout the entire transfer process, the need for efficient multi-object transfer, the need to preserve privacy and content secrecy throughout the transfer process when desired, and the need to consider post-transfer object state.

While our studies revealed many useful insights, further study is warranted in a number of directions. First, occasionally in our studies, players wished to transfer cards directly from one tablet to another when "giving a card" to another player. Moreover, one can imagine other task contexts, particularly during more cooperative group activities, where people might want to exchange task content directly from one tablet to another. Future design extensions should consider this functionality. Similarly, other design extensions might include the ability to "share" tabletop content on someone else's tablet to allow more cooperative transfer patterns between available surface devices.



SURFNET / Designing Digital Surface Applications

is a compendium of research findings from a Canadian research network that integrated innovative research in two critical areas –software engineering (SE) and human-computer interaction (HCI)– to identify critical requirements, design new engineering processes, and build new tools for surface-based application development. Funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) from 2009 to 2015, SurfNet’s research clustered around three themes: Humanizing the Digital Interface, Improving Software Time to Market and Building Infrastructure for Digital Surfaces. Research was driven by the needs of four application areas: Planning, Monitoring and Control Environments; Learning, Gaming, New Media and Digital Homes; Software Team Rooms; and Health Technologies.