

Distributed Tabletops: Supporting Remote and Mixed-Presence Tabletop Collaboration

Philip Tuddenham and Peter Robinson
University of Cambridge Computer Laboratory
15 JJ Thomson Avenue, Cambridge CB3 0FD, UK
{firstname}.{lastname}@cl.cam.ac.uk

Abstract

Mixed-presence tabletop interfaces aim to support collaboration between remote groups. However, it is unclear why tabletop interaction techniques should be important for mixed-presence or remote collaboration, and recent projects in this area differ as to which elements of tabletop interaction they choose to support. In this paper we discuss the benefits of tabletop interaction for mixed-presence and remote collaboration. In particular, we wish to support the natural tabletop awareness mechanisms of territoriality, orientation and consequential communication. We derive design guidelines for such systems and present Distributed Tabletops, a novel system that can be customised to investigate various mixed-presence tasks. Our early observations of Distributed Tabletops in use validate our design guidelines.

1. Introduction

Interactive tabletop interfaces have emerged as an effective tool for co-located collaboration over digital artifacts. Collaborators sit around a horizontal multi-touch surface that displays digital artifacts such as photos, documents and web pages. They can then interact simultaneously to move, reorient and manipulate the artifacts, as they might with paper artifacts on a conventional tabletop.

Several recent projects have investigated the possibility of linking together geographically-separated tabletop displays in order to create a shared workspace for remote collaboration [1, 5, 9, 3]. Each of the geographically-separated collaborators sits at his or her own tabletop display. The displays are then linked, perhaps via the Internet, so that these remote collaborators can each interact with the shared artifacts and see each other's actions, as if co-located around the same tabletop.

Previous efforts to support remote collaboration over

digital artifacts had tended to use conventional monitor- and mouse-based interaction to create shared workspaces, but this approach encountered well-documented problems [4, 6, 7]. For instance, on a conventional monitor there is often insufficient space for collaborators to work in different parts of the workspace without losing an awareness of each other's actions. Large horizontal displays and tabletop interaction techniques might allow these problems to be overcome, so that remote collaborators can communicate as though sat around the same table. Furthermore, by using multi-touch surfaces at each remote site, we can also support collaboration between geographically-separated groups of collaborators, namely mixed-presence collaboration [23].

The recent remote tabletop collaboration projects have set out to investigate these issues but, surprisingly, some important questions remain unanswered:

- Why might tabletop interaction, so effective for co-located collaborators, be similarly appropriate for remote and mixed-presence collaboration?
- Which elements of co-located tabletop collaboration are important for remote and mixed presence tabletop collaboration? Some recent remote tabletop collaboration projects differ on this issue and, for instance, very few systems allow remote collaborators to reorient digital artifacts as they can in the majority of co-located tabletop systems.
- What additional factors must be considered in the design of tabletop systems for mixed-presence collaboration, as opposed to just remote collaboration?

In the first half of this paper we review literature (Section 2) and discuss the above questions in the context of prior work (Section 3). In particular, we believe that supporting natural tabletop awareness mechanisms of orientation, territoriality and consequential communication will greatly benefit remote and mixed-presence tabletop collaboration, and we suggest design guidelines for achieving these goals.

We then present *Distributed Tabletops*, a novel system that we have created, following these guidelines, to support remote and mixed-presence collaboration over digital artifacts (Section 4). By using direct input mechanisms, preserving a co-located-style seating arrangement among distributed collaborators, allowing them to move and orient artifacts, and projecting remote arm embodiments, Distributed Tabletops supports the natural awareness mechanisms posed above. We report our early experiences with Distributed Tabletops and discuss the design trade-offs in such systems (Section 5).

2. Background

Systems for synchronous collaboration (i.e. where all participants collaborate at the same time) traditionally support either co-located collaboration or remote collaboration. Mixed-presence groupware is a recent addition to this dichotomy [23] and must support both. This paper draws on work in each of these areas.

2.1. Co-located Collaboration

Single Display Groupware (SDG) [21] describes systems that allow co-located collaborators, each with their own input device, to interact simultaneously using a shared display. Tabletop interfaces are a form of SDG that use a large horizontal display and multi-user direct input mechanisms such as styluses or a multi-touch surface. Early systems [e.g. 22] took inspiration both from augmented paper, such as Wellner’s DigitalDesk [30], and studies of co-located tabletop collaboration [26]. Recent research focuses on studies of human behaviour [e.g. 12, 19, 17] that can inform tabletop interaction techniques [e.g. 13].

2.2. Remote Collaboration using Monitor/Mouse Systems

In these systems, each remote collaborator uses their own conventional desktop computer. The computers are then linked to create a shared workspace for remote collaboration. Such systems tend to be characterized by a lack of awareness among collaborators about each other’s actions in the workspace, with collaboration suffering as a result [6, 7, 4]. The problem is particularly acute in systems in which each collaborator can manipulate their view of the workspace independently of others, for example to scroll to a different region of the workspace. Some of these awareness problems seem rooted in the combination of small display and impoverished input and remote embodiments that characterize these systems. Another area of research has taken an alternative approach using large-format displays.

2.3. Remote and Mixed-Presence Collaboration on Large Displays

In large-format remote collaboration systems, each remote collaborator uses their own large display, either horizontal or vertical, usually with direct input mechanisms (styluses or touch). The displays are then linked together to provide a shared workspace for remote collaboration.

The earliest such systems were shared drawing surfaces, such as VideoWhiteboard [27]. Remote collaborators could interact simultaneously to sketch, while shadows of their arms and bodies were projected onto the remote workspace to allow gesture and to maintain awareness and a sense of presence. Clearboard [10] additionally allowed collaborators to make eye contact “through” the drawing surface. Tang et al. [23, 24] used a shared drawing surface to investigate remote embodiments in mixed-presence collaboration.

Other projects investigated remote collaboration over tangible artifacts using a similar approach. DoubleDigitalDesk [30] was an early such system for remote collaboration over paper documents on desks. A camera mounted above each surface captured an image of the tangible artifacts and the user’s arms; this image was then projected onto the remote collaborator’s surface. More recent systems use similar techniques for collaboration over paper documents [14], board games and sketching [33], single physical objects [31], and collaborative physical tasks [11].

In these tangible remote collaboration systems, a given artifact is only ever tangible for one collaborator; other collaborators see mere video projections and cannot interact directly. This effect is unavoidable, even desirable, in tasks such as remote bomb disposal, but can be problematic in other domains. Brave et al. [2] present an alternative approach whereby each remote collaborator has a complete set of tangible artifacts. When the system senses that a user has moved an artifact, local actuators move the corresponding artifact on the remote collaborators’ displays. In practice, however, this approach constrains collaboration.

This has led to systems that use purely digital artifacts to investigate large format displays for remote collaboration, for tasks other than sketching, and without the asymmetry problem of remote tangible systems. Escritoire [1], RemoteDT [5], ViCAT/TIDL [9] and the very recent DigiTable work [3] all address technical challenges in this area. They present each remote collaborator with a shared workspace of movable interactive digital artifacts on a large horizontal display, rather like co-located tabletop interfaces. For instance, Escritoire is designed for remote collaboration over images that can be arranged using styluses, while RemoteDT and TIDL both support legacy applications in mixed-presence collaboration using, respectively, a multitouch surface and multiple mice.

These remote tabletop projects take inspiration from co-

located tabletop research, but are selective about the elements that they adopt. For instance, orientation serves key roles in co-located tabletop collaboration, as we shall discuss later, and yet only DigiTable allows collaborators to reorient digital artifacts. We were also surprised to find that none of these works justify their decisions about supporting orientation. Similarly, all these systems allow remote collaborators to share the same virtual “seat” at the table, without explaining this decision, and yet this situation would never occur in co-located tabletop collaboration. We shall further discuss these projects later in the paper.

3. Designing for Mixed-Presence and Remote Tabletop Collaboration

As we have seen, recent attempts to create mixed-presence tabletop collaboration systems have differed as to which characteristics of co-located tabletop collaboration they choose to support, and why these might be relevant to remote collaboration. We begin by enumerating the salient characteristics of co-located tabletop collaboration over digital artifacts in order to discuss the relevance of each to mixed-presence and remote collaboration.

3.1. Co-located Tabletop Interaction

We draw on observational studies of tabletop collaboration [26, 12, 19], requirements for tabletop collaboration [18] and various tabletop interfaces [e.g. 20, 13]. We aimed to avoid focusing on particular tasks by choosing characteristics of collaboration around digital artifacts, regardless of what the artifacts actually represent. We enumerate the characteristics as follows:

- A large horizontal display surface.
- Collaborators sit in different positions around the edge of the surface.
- Direct input mechanisms (stylus or touch).
- Digital artifacts can be moved and freely reoriented.
- Collaborators can see each other’s arms.
- Simultaneous interaction by multiple collaborators.
- Collaborators can see each other’s bodies and faces.
- Collaborators can talk to each other.

Allowing collaborators to talk to each other and to interact simultaneously is clearly important for any kind of collaboration, and does not warrant discussion.

Supporting adequately the ability of remote collaborators and mixed-presence collaborators to see each other’s faces and postures is technically challenging. Work in this area has only recently shown promising results [16]. Though desirable, it is outside the scope of this research.

The remaining characteristics (the first five) are essential for three natural tabletop awareness mechanisms that are central to co-located tabletop collaboration: territoriality; orientation of artifacts; and consequential communication. We now discuss in turn the relevance of each of these mechanisms to remote tabletop collaboration.

3.2. Territoriality

Tse et al. [28] studied spatial partitioning in SDG applications using conventional monitor/mouse interaction and found that collaborators used spatial partitioning to avoid interfering with each other’s actions. Scott et al. [19] observed that co-located collaborators naturally partition the space on a tabletop to serve different roles:

- A personal territory is the area directly in front of a collaborator. It allows them to reserve a particular area of the table for themselves. They can then use it as a place to reserve artifacts for themselves and to conduct individual work as part of the group task. People also monitor the progress of work in the personal territories of their collaborators.
- Group territory occupies space that isn’t considered personal territory. It is used for the main group task and also to transfer artifacts, for example to signal availability by depositing them there.
- A storage territory is an area in which collaborators store and organize artifacts into related groups or piles.

Territoriality can be considered a natural awareness mechanism that allows participants to reserve resources for themselves and to transition between individual and group work during a collaborative task. This is particularly important in *mixed focus tasks*, which require people to transition frequently between individual and group work to complete the task [6, 7]. By contrast, monitor/mouse remote collaboration systems perform badly in these mixed-focus tasks because it is difficult for collaborators to accomplish individual work while also maintaining an awareness of each others actions. For instance, for many tasks, the screen is not large enough to allow two collaborators to work side by side on different parts of the workspace and so collaborators must scroll their views of the workspace independently, losing an awareness of each other’s actions [6, 7].

The large physical size of remote and mixed-presence tabletops may go some way to solving these problems, but we believe that if such systems are truly to permit both individual and group work in the manner of co-located collaboration then they must support the corresponding awareness mechanisms, such as territoriality.

These remote and mixed-presence systems will therefore need to be sufficiently large to incorporate territories. Different collaborators will have to sit at different positions

around the edge of the “virtual table” so that they each have space for a personal territory in the shared workspace. This “virtual seating arrangement” must mimic co-located seating arrangements and be preserved across the different connected tabletops if we are to promote a sense of spatial arrangement of people. Direct input mechanisms would ensure that even remote collaborators would have to reach across the table if they are to access artifacts in another collaborators’ personal territory. Finally, all collaborators must be able to see each other’s arm movements on the table in order to emphasise the seating arrangement and therefore the locations of the personal territories. For remote and mixed-presence collaboration, this could be accomplished by projecting arm shadows that extend from each collaborator’s seat as a remote embodiment in a similar style to the systems reviewed earlier. Other remote embodiments, such as the telepointers, do not extend from the seat and hence do not emphasise the seating arrangement as we would desire.

3.3. Orientation of Artifacts

Kruger et al. [12] investigated the roles of artifact orientation in tabletop collaboration:

- **Comprehension.** People often orient artifacts to be most readable for themselves.
- **Coordination.** The orientation of an artifact indicates its ownership and availability. When people orient an artifact towards themselves it suggests that they have personal use of it. Collaborators can also establish personal and group territories by appropriately orienting artifacts in these spaces.
- **Communication.** Orienting an artifact to another person indicates that the artifact and accompanying talk and gestures are directed at that person, whereas orienting an artifact to yourself signals that you are doing personal work.

Hauber et al. [8] briefly discuss a remote tabletop system that allowed remote collaborators to position and orient digital photos as part of a wider study of presence. They report the use of orientation for comprehension, but, perhaps because of their task, not for coordination or communication.

Orientation is another natural awareness mechanism that allows participants to reserve resources for themselves and to transition between individual and group work; roles which, as we noted previously, should not be overlooked in the design of remote collaboration systems. Orientation also allows collaborators to establish an audience for their utterances, aiding turn-taking, a noted problem of group-to-group collaboration systems.

If remote and mixed-presence collaboration systems are to support orientation in the manner of co-located tabletops, then once again we need a consistent “virtual seating

arrangement” among the connected tabletops so that artifacts that are oriented to one collaborator are not oriented to any other collaborators, co-located or remote. Collaborators must be able to reorient digital artifacts in the workspace, perhaps using a technique like Rotate ‘N’ Translate [13]. Use of direct input mechanisms and arm shadows would aid awareness of the seating arrangement and of who is orienting to whom.

3.4. Mixed-Presence Embodiments

If we are to support mixed-presence collaboration, as opposed to just remote collaboration, then we must address the associated human factors. In particular, naïve use of embodiments like telepointers leads to a disparity in the conversation dynamic whereby a user is much more likely to interact with their co-located collaborators than with their remote collaborators, with a negative effect on collaboration. Tang et al. [23, 24] observed the use of a mixed-presence whiteboard and show that richer embodiments, such as arm shadows, mitigate this effect. They review the roles played by physical bodies in collaboration and suggest that remote embodiments for mixed-presence collaboration should:

- Be controlled by direct input mechanisms and allow remote collaborators to interpret current actions and the actions that led up to them. By contrast, impoverished indirectly-controlled embodiments, such as telepointers, do not adequately convey the awareness information that is unintentionally communicated in co-located collaboration, such as arm position (“consequential communication”).
- Allow remote collaborators to interpret gestures by capturing and rendering fine-grained movements and postures.
- Appear in the workspace in order to convey gestures as they relate to the workspace.
- Be visible not only to remote collaborators but also provide local feedback so that we might infer how our actions are interpreted by remote collaborators.

If we are to support mixed-presence collaboration without experiencing a conversation disparity then we must support rich arm shadow embodiments that follow these design guidelines, rather than impoverished telepointers.

3.5. Design Guidelines

Based on this analysis, we believe that by supporting natural tabletop awareness mechanisms like territoriality, orientation and consequential communication, remote and mixed-presence tabletops can provide effective support for both individual and group work. This addresses some of the




Figure 1: Mixed presence tabletop collaboration over digital artifacts using Distributed Tabletops. In this application, participants interact simultaneously to move and reorient the words to create poetry. Visible in the photos are arm shadows, personal territories, and artifacts at arbitrary orientations.

shortcomings of conventional monitor/mouse remote collaboration.

We have discussed ways in which remote or mixed-presence tabletops can support these mechanisms and we now summarise this with a series of design guidelines for such systems:

1. Large horizontal display surfaces.
2. Consistent “virtual seating arrangement” among the connected tables.
3. Direct input mechanisms (stylus or touch).
4. Digital artifacts can be moved and reoriented using a technique such as Rotate 'N' Translate [13].
5. Collaborators are represented remotely by arm shadow embodiments with local feedback.
6. Simultaneous interaction by multiple collaborators.
7. Collaborators can talk to each other.

Of the prior work, *Escritoire* [1], *RemoteDT* [5] and *Vi-CAT/TIDL* [9] do not support a consistent “virtual seating arrangement”, arm shadow embodiments or reorientation. The very recent *DigiTable* work [3] does support these guidelines, but also has a feature that allows remote collaborators to violate the consistent seating arrangement. The rationale for the features we have highlighted is not discussed, beyond a need for remote embodiment, and the system is not evaluated, so we do not know whether it supports the awareness mechanisms we have drawn upon.

We now put these guidelines into practice to create *Distributed Tabletops*.

4. Distributed Tabletops

Distributed Tabletops is a prototype system that we have created, following our design guidelines, to investigate mixed-presence and remote tabletop collaboration over digital artifacts. In this section we present an overview, showing how the system addresses each of the design principles. We then briefly explain the implementation and describe our early observations of the system in use.

Figure 1 illustrates our *Distributed Tabletops* system. We connect two or more geographically-separated large horizontal displays, with multiple collaborators sat around each. Our system then links the displays so that they all display exactly the same contents at all times, creating a shared workspace for collaboration. Each collaborator has a stylus and can interact simultaneously with digital artifacts that are displayed. We use a reusable toolkit in order to support various tasks and artifacts, such as virtual puzzle pieces, virtual pages of text, or virtual spreadsheets. Unlike prior work in this area, participants sit at different locations around the edge of the “virtual table” and can use their styluses to move and reorient components using the Rotate 'N' Translate [13] technique, as they might at a co-located tabletop.

We use arm shadows as remote embodiments. A video camera mounted above each tabletop captures collaborators’ arm gestures in the workspace. The system identifies the arm outlines, which are then displayed as translucent “shadows” on the other tables. The arm shadows are displayed in the correct place in the workspace and appear at the correct time with respect to artifact manipulations in the workspace. Collaborators also see their own arm shadows, to provide local feedback. All participants can gesture and

interact with the workspace simultaneously.

The system supports a variety of artifacts and we chose to test a “magnetic poetry” task, in which collaborators create poetry by moving and reorienting words that appear as artifacts on the table. We chose this fairly simplistic task for a variety of reasons: interaction is constrained to just moving and reorienting, so there is little scope for problems; collaborators can carry out both individual work and group work in the context of the task; and there are many small artifacts, so it is easier to observe use of orientation and the roles of territory. Nevertheless, the natural awareness issues identified earlier are not only relevant to this simplistic task but also generalize to real-world tasks.

4.1. Implementation

Much of the necessary software infrastructure is encapsulated in our T3 software [29], which is freely available for academic research. T3 is a Java toolkit designed to allow two tabletop displays to be linked together in the manner described above. Each display is controlled by a local computer running the T3 client software, which performs the rendering and controls the Bluetooth Anoto styluses used by the local participants. All the clients connect to a central T3 server, which runs the actual application. All the clients receive update messages from the server about artifact location, size, orientation and contents, and send back user input messages detailing stylus position, etc.

T3 applications create rectangular artifacts which collaborators can move, orient and interact with using their styluses. T3 supports rapid prototyping by allowing creation of artifacts containing legacy Java Swing components such as buttons, web-browsers and spreadsheets. In order to support responsive rendering of oriented components, the T3 clients store the image of each artifact as an OpenGL texture, which can then be rendered into the framebuffer at the correct location and orientation using hardware-acceleration provided by commodity graphics cards.

On top of T3, we have added the functionality to extract arm outlines from video camera images, and to render them as translucent shadows on the workspace at each tabletop. A commodity webcam is positioned above each table and sends images to the local computer, which identifies the outlines of any arms present on the tabletop surface. These contours are then sent to the other clients, which render the outline on their own displays.

The arm segmentation process is particularly difficult for front-projected displays like ours because the projected light discolours the hands and creates extra shadows. The majority of arm-shadow work uses rear-projected displays from which arms can easily be segmented using background subtraction and/or skin-colour segmentation [e.g. 24]. Nevertheless, we have managed to use a simple algorithm to re-

liably segment arms and hands on a front-projected display. It is well known that skin is a poor reflector of blue light, and thus by controlling the colours displayed on the tabletops we can perform the segmentation using background subtraction in the blue colour plane. We then find the contours in the segmented image, after first filtering out small shapes that represent noise. To reduce network and processing bandwidth, we approximate the contour by a polygon before transmitting it to the other clients. More complex segmentation algorithms [e.g. 32] may well produce reliable results without the need to control the colour, but our goal here was to produce a simple prototype that would allow us to investigate our design guidelines.

Camera images at a resolution of 640x480px are received, processed and rendered at 15fps with no noticeable delay. This frame rate is limited by the camera, rather than the system. The artifacts themselves are rendered at 60fps, the highest rate supported by the projector, with no noticeable delay. We use the OpenCV computer vision library, and each client is controlled by an Intel Core 2 2.4GHz PC.

4.2. Early Observations

We have not yet conducted formal user studies of Distributed Tabletops, but early observations of the system in use are promising.

We tested the system using the poetry application in two sessions with 6 participants in total. In each session two participants sat at one tabletop and one participant at the other. Participants reported that the system felt responsive to use. None of them had any problems interacting with the interface, and all easily managed to string words together.

Even in this very limited trial, there was some evidence to suggest that participants were using the notions of territory and orientation. In both sessions, all participants were observed establishing personal territories by orienting words towards themselves in these regions of the table, and subsequently used these regions to construct fragments of poetry away from the other members of the group.

All participants in both sessions were observed taking words from each other’s personal territories. However, this was not unintentional; rather they were enjoying themselves and began deliberately and playfully “stealing” words from each other and moving them into their own personal territories, resulting in exclamations like “Give that back!”. Participants would also move words away when they thought they were about to be “stolen”. This suggests that participants were aware of the personal territories of their collaborators and also that they felt a sense of ownership over the words in their own personal territories. Furthermore it indicates that, through the embodiments, participants had a good awareness of each other’s actions and intentions.

Overall, all the participants enjoyed interacting with each

other in both sessions, and we did not observe a conversation disparity between the two tabletops.

5. Design Trade-offs in Groupware

Our design principles and the Distributed Tabletops system are motivated by the need for natural awareness mechanisms in tabletop collaboration. This aids mixed-focus tasks, which involve both individual and group work.

In monitor/mouse remote collaboration systems there are well-known trade-offs between supporting individual work and group work. Gutwin and Greenberg [6] show that allowing individuals to interact in a powerful way often prevents the group from maintaining an awareness of each other's actions in the workspace. They identify three areas where such trade-offs occur, and in this section we discuss each of these in the context of Distributed Tabletops.

Workspace navigation. For many tasks, conventional screens are too small to allow two collaborators to work side by side and to display a reasonable-sized workspace. Accordingly, collaborators must scroll independently in the workspace to accomplish individual work, leading to a lack of awareness. By contrast, Distributed Tabletops provides a large display with high awareness. There is sufficient space for collaborators to work side by side and, although the workspace cannot be scrolled, the size is presumably sufficient for the tasks currently carried out on physical tables. Furthermore, unlike monitor/mouse interaction, the task is split into multiple artifacts which can then be grouped into piles, allowing efficient use of the display space without resorting to scrolling. For example, a newspaper article layout task might use artifacts representing pages and articles which can be overlaid, so that the task can be completed effectively even though there is not sufficient space on the display to view all pages and articles at once. That said, the trade-off still arises when collaborators work on the same artifact at the same time: making copies of artifacts makes individual work easier but limits group awareness [25]. We believe, therefore, that this trade-off can be mitigated if the task is partitioned in a way that reduces the need for collaborators to work on the same artifact at the same time.

Artifact manipulation. In monitor/mouse remote collaboration a trade-off exists between allowing powerful artifact manipulation for individuals and providing awareness information to the group. The authors suggest powerful individual actions be made more perceivable to the group using techniques like animation and sounds. This design trade-off undoubtedly applies to Distributed Tabletops, such as in the location of controls [15]. However, we believe that the effect will be mitigated by support for natural awareness mechanisms like territoriality and orientation.

Workspace representation. The authors consider the problem of providing a shared workspace for collaborators

who each wish to view different representations of the same underlying data. The same problem applies both to co-located tabletop interfaces and to Distributed Tabletops, and has yet to be explored in either.

6. Discussion and Conclusions

We believe that our design guidelines are an appropriate foundation for mixed-presence and remote tabletop collaboration over digital artifacts. Even in our early observations of Distributed Tabletops, we have seen evidence to suggest that participants were using mechanisms of territory and orientation as they would in a co-located setting, had a good awareness of each other's actions and intentions, and collaborated in a mixed-presence setting without a conversation disparity between the two tabletops.

Further work will involve formal user studies to investigate the differences between the roles of territory and orientation in co-located and mixed-presence collaboration, and the roles that they can play in different tasks.

Further development of the computer vision algorithm is also required if the system is to support tabletop applications that use unrestricted colours, and other more complex algorithms may be more appropriate in these cases [e.g. 32]. However, many applications are nevertheless possible using our algorithm, and we are currently using the Distributed Tabletops system to investigate remote and mixed-presence document collaboration and command and control tasks. It may also be possible to combine this kind of interface with systems that investigate other aspects of group-to-group collaboration such as faithfully conveying eye-gaze and body posture [16].

This paper makes three contributions. Firstly, we observe that previous work investigating mixed-presence and remote tabletop collaboration over digital artifacts seems to be uncertain as to what constitutes tabletop collaboration and why these techniques might be important in such a different setting. We begin by discussing the various aspects of co-located tabletop collaboration and show that supporting natural awareness mechanisms like territoriality, orientation, and a careful choice of remote embodiment are all crucial to remote and mixed-presence tabletop collaboration over digital artifacts.

Secondly, based on this discussion, we pose design guidelines for systems to support such collaboration. Finally, we present the Distributed Tabletops as a method. Unlike prior work, Distributed Tabletops preserves a consistent "virtual seating arrangement" between tables, uses direct input mechanisms, allows orientation of artifacts, and uses arm shadows as remote embodiments. Early observations indicate that our design principles are valid, and we discuss trade-offs in groupware design.

In summary, Distributed Tabletops represents a first step

towards design principles and a system for remote and mixed-presence tabletop collaboration over digital artifacts.

Acknowledgements

We gratefully acknowledge Mark Ashdown, members of the Rainbow group, and the anonymous reviewers, for useful comments; and Thales Research and Technology (UK) and the EPSRC, who jointly funded this work.

References

- [1] Ashdown, M. and Robinson, P. Escritoire: A Personal Projected Display. *IEEE MultiMedia* 12, 1 (Jan. 2005), 34-42.
- [2] Brave, S., Ishii, H., and Dahley, A. Tangible Interfaces for Remote Collaboration and Communication. In *Proc. CSCW 1998*, 169-178.
- [3] Coldefy, F. and Louis-dit-Picard, S. DigiTable: an interactive multiuser table for collocated and remote collaboration enabling remote gesture visualization. In *Proc. PROCAMS 2007*.
- [4] Dourish, P. and Bellotti, V. Awareness and coordination in shared workspaces. In *Proc. CSCW 1992*, 107-114.
- [5] Esenther, A., and Ryall, K., RemoteDT: Support for Multi-Site Table Collaboration. In *Proc. Int. Conf. Collaboration Technologies (CollabTech)*, 2006.
- [6] Gutwin, C. and Greenberg, S. Design for individuals, design for groups: tradeoffs between power and workspace awareness. In *Proc. CSCW 1998*, 207-216.
- [7] Gutwin, C. and Greenberg, S. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. *Comput. Supported Coop. Work* 11, 3 (Nov. 2002), 411-446.
- [8] Hauber, J., Regenbrecht, H., Billingham, M., and Cockburn, A. Spatiality in videoconferencing: trade-offs between efficiency and social presence In *Proc. CSCW 2004*, 413-422.
- [9] Hutterer, P., Close, B. S., and Thomas, B. H. Supporting Mixed Presence Groupware in Tabletop Applications In *Proc. TABLETOP'06*, 63-70.
- [10] Ishii, H. and Kobayashi, M. ClearBoard: a seamless medium for shared drawing and conversation with eye contact. In *Proc. CHI 1992*, 525-532.
- [11] Kirk, D. and Stanton Fraser, D. Comparing remote gesture technologies for supporting collaborative physical tasks In *Proc. CHI 2006*, 1191-1200.
- [12] Kruger, R., Carpendale, S., Scott, S. D., and Greenberg, S. Roles of Orientation in Tabletop Collaboration: Comprehension, Coordination and Communication. *Comput. Supported Coop. Work* 13, 5-6 (Dec. 2004), 501-537.
- [13] Kruger, R., Carpendale, S., Scott, S.D., and Tang, A. Fluid integration of rotation and translation. In *Proc. CHI 2005*, 601-610.
- [14] Luff, P., Heath, C., Kuzuoka, H., Yamazaki, K., and Yamashita, J. Handling documents and discriminating objects in hybrid spaces. In *Proc. CHI 2006*, 561-570.
- [15] Morris, M. R., Paepcke, A., Winograd, T., and Stamberger, J. TeamTag: exploring centralized versus replicated controls for co-located tabletop groupware. In *Proc. CHI 2006*, 1273-1282.
- [16] Nguyen, D. T. and Canny, J. Multiview: improving trust in group video conferencing through spatial faithfulness. In *Proc. CHI 2007*, 1465-1474.
- [17] Ryall, K., Forlines, C., Shen, C., and Morris, M. R. Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. In *Proc. CSCW 2004*, 284-293.
- [18] Scott, S., Grant, K., and Mandryk, R. System guidelines for co-located collaborative work on a tabletop display. In *Proc. ECSCW 2003*, 2003.
- [19] Scott, S.D., Carpendale, M.S.T., and Inkpen, K.M. Territoriality in collaborative tabletop workspaces. In *Proc. CSCW 2004*, 294-303.
- [20] Shen, C., Vernier, F.D., Forlines, C., and Ringel, M. Diamondspin: an extensible toolkit for around-the-table interaction. In *Proc. CHI 2004*, 167-174.
- [21] Stewart, J., Bederson, B. B., and Druin, A. Single display groupware: a model for co-present collaboration. In *Proc. CHI 1999*, 286-293.
- [22] Streitz, N. A., Geissler, J., Holmer, T., Konomi, S., Miller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R. i-LAND: an interactive landscape for creativity and innovation. In *Proc. CHI 1999*, 120-127.
- [23] Tang, A., Boyle, M., and Greenberg, S. Understanding and mitigating display and presence disparity in mixed presence groupware. *J. Res. and Practice in Inf. Tech.* 37, 2.
- [24] Tang, A., Neustaedter, C., and Greenberg, S. VideoArms: Embodiments for Mixed Presence Groupware. In *Proc. HCI 2006*, 85-102.
- [25] Tang, A., Tory, M., Po, B., Neumann, P., and Carpendale, S. Collaborative coupling over tabletop displays. In *Proc. CHI 2006*, 1181-1190.
- [26] Tang, J. C. Findings from observational studies of collaborative work. *Int. J. Man-Mach. Stud.* 34, 2 (Feb. 1991), 143-160.
- [27] Tang, J. C. and Minneman, S. VideoWhiteboard: video shadows to support remote collaboration. In *Proc. CHI 1991*, 315-322.
- [28] Tse, E., Histon, J., Scott, S. D., and Greenberg, S. Avoiding interference: how people use spatial separation and partitioning in SDG workspaces. In *Proc. CSCW 2004*, 252-261.
- [29] Tuddenham, P., and Robinson, P. T3: Rapid Prototyping of High-Resolution and Mixed-Presence Tabletop Applications. In *Proc. IEEE TABLETOP 2007*.
- [30] Wellner, P. Interacting with paper on the DigitalDesk. *Commun. ACM* 36, 7 (Jul. 1993).
- [31] Wesugi, S. and Miwa, Y. "Lazy Susan" Communication System for Remote, Spatial and Physical Collaborative Works. In *Proc. IEEE TABLETOP 2006*, 35-42.
- [32] Wilson, A. D. PlayAnywhere: a compact interactive tabletop projection-vision system. In *Proc. UIST 2005*, 83-92.
- [33] Wilson, A. D., and Robbins, D. C. PlayTogether: Playing Games across Multiple Interactive Tabletops *IUI'07 workshop on Tangible Play*.