

Coordination and Awareness in Remote 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 Remote collaboration technologies frequently provide a shared visual workspace of the task at hand, but often lack support for the visual cues and work practices of co-located collaboration. This is particularly acute in design tasks, in which the shared workspace is the focus of collaboration. Lately there has been growing interest in *remote tabletop interfaces*: large horizontal interactive surfaces that provide shared workspaces for remote collaboration. These interfaces may afford some of the cues and beneficial work practices associated with collaboration at conventional tables. If so, they may offer benefits over other remote collaboration technologies. However, while a number of remote tabletop interfaces have been constructed, there are few empirical findings around these interfaces in practice. This chapter reviews current work in remote tabletop interfaces and then presents an exploratory study investigating two work practices in remote tabletop collaboration.

Introduction

Effective support for synchronous remote collaboration has long proved a desirable yet elusive goal for computer technology. Consumer videophone technology was unveiled in 1964, amid forecasts of replacing standard telephony by the early 1970s [1]. Forty years later, video conference technology is more widely available yet remains largely unused, and so perhaps little has changed in practice. Researchers have discussed a variety of problems, most notably poor reproduction of the visual cues, such as eye gaze, that mediate face-to-face conversation [2], and the inability to initiate and conduct the informal collaborative interactions that oc-

cur outside of formal scheduled collaboration [3]. Both problems may soon be solved. The falling cost and commoditisation of large displays, camera equipment, and network bandwidth are making always-on remote “video windows” between spaces more feasible in practice, while novel techniques faithfully reproduce visual cues that were absent or distorted in previous technologies [4, 5]. These advances are currently making their way into commercial meeting-room systems.

Technologies that provide a shared visual workspace of the task at hand, rather than of the other remote collaborators, remain rather less advanced. Researchers have demonstrated large-format remote whiteboards to support fixed sketching [6]. However, remote collaboration for other workspace-based activities is largely confined to collaborative versions of conventional desktop-computer applications. These systems lack support for the visual cues and work practices that underpin visual workspaces in co-located collaboration, leading to well-documented problems [7]. The problems are particularly acute in tasks, such as collaborative design, in which a shared visual workspace is a natural focus [8]. A shared visual workspace is also important in a number of other prevalent collaborative tasks such as document review, data analysis and information-gathering.

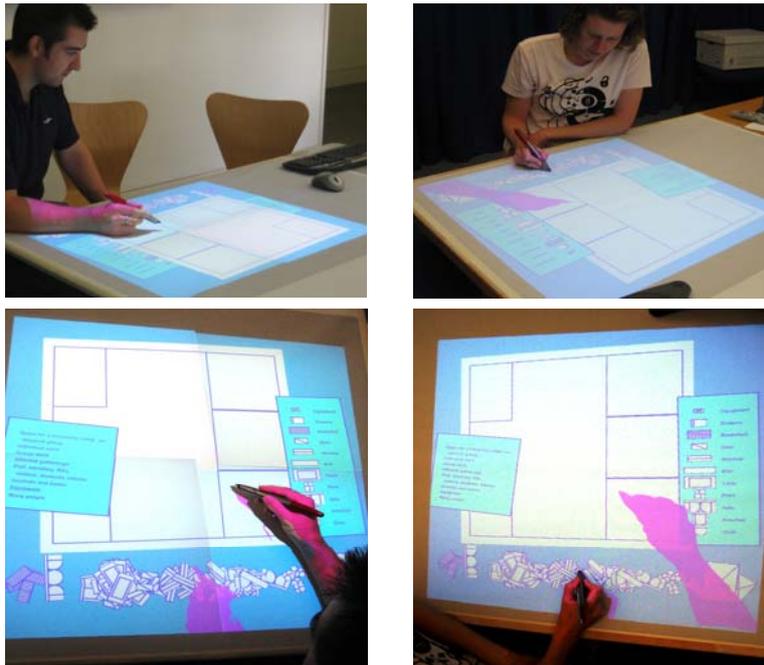


Fig. 1. Representative example of a remote tabletop interface. Two tabletop interfaces are linked to provide a large shared workspace with purple remote arm “shadows”. From [47]; © 2009 ACM, Inc. Included here by permission.

Recent developments in technology to support co-located collaboration may offer opportunities to address these problems of remote work. The increasing size and resolution of affordable display technologies have led to predictions of ubiquitous large displays [9]. Researchers investigating co-located collaboration have consequently constructed large horizontal interactive displays around which co-located collaborators can sit and interact using their hands. These tabletop interfaces present interactive shared task artefacts that appear on the display and mimic real-world task artefacts such as photos or puzzle pieces.

Scott [10] and Morris [11] argue that conventional tables are prevalent in work environments because they are well-suited to many collaborative two-dimensional information tasks, such as planning, scheduling, brainstorming, design, and layout. Their size and orientation enables collaborators to sit around, and to spread out and spatially-organise the task artefacts. Tabletop interfaces aim to enable access to interactive content during such tasks, and in a way that affords established collaborative work practices. They have been demonstrated for tasks including planning [12] and design [13], and can afford beneficial work practices observed at conventional tables, such as fluid switching between individual and group work [14], and territorial partitioning as a coordination mechanism [13].

A tabletop approach might similarly offer benefits to shared workspaces for remote collaboration. A number of research projects have used large horizontal interactive surfaces to present a shared visual workspace to support remote collaboration [15, 16, 17, 18, 19, 20, 21, 22]. Figure 1 shows a representative example of these *remote tabletop interfaces*.

However, there has been little consensus among researchers on how to design these remote tabletop interfaces. Moreover, many remote tabletop projects have concerned themselves with constructing remote tabletop interfaces, rather than characterising the extent to which they support remote collaboration. This has been necessary because the interfaces present technical challenges. Nevertheless, motivations of supporting remote collaboration have often not been developed into well-grounded claims, and there has been little evaluation to determine whether such claims would hold in practice.

In particular, the design of remote tabletops is based in part on an untested assumption: *that a tabletop design will afford for remote collaborators the beneficial work practices as co-located tabletop collaboration, such as fluid transitioning between individual and group work, and coordination based on spatial partitioning*. Conventional groupware suffers from limited support for the work practices and visual cues of co-located collaboration [23]. Accordingly, if this assumption holds, then remote tabletop interfaces may provide a more effective remote collaboration medium. However, the assumption has not been tested in practice, and in contrast to co-located tabletop collaboration, we know little about the work practices afforded by remote tabletops.

The aims of this chapter are threefold: (i) to review work in this area, (ii) to identify key areas of exploration, and (iii) to present a study that begins to explore these areas.

We begin by presenting a review of remote tabletop interfaces, which details the prior systems and discusses their origins and the empirical findings. Based on the review, we identify work practices as a key area of exploration for remote tabletop interfaces. We then present a study that compares work practices in co-located and remote tabletop collaboration. We conclude with a discussion of future trends and a summary.

Background: From Remote Sketching to Remote Tables

Remote Tabletop Interfaces

This section adopts a fairly narrow definition of remote tabletop interfaces: remote horizontal interactive surfaces linked together to create a shared workspace in which the task artefacts are composed of computer-generated imagery. Large shared workspaces for whiteboard-style fixed remote sketching are discussed later (see “Large Format Remote Sketching”). Similarly, remote augmented reality interfaces, in which some task artefacts are physical rather than computer-generated, again have different design issues and are discussed later (see “Remote Augmented Reality Surfaces”).

Escritoire [24, 16] was one of the earliest remote tabletop interfaces by this definition. Instead of real paper, it used projected light to create “virtual sheets of paper” in a large horizontal shared workspace for remote collaboration. Two projectors were arranged to create a large (A0 size) low-resolution peripheral area, for storing virtual sheets, surrounding a smaller high-resolution foveal area, into which virtual sheets could be dragged for manipulation or reading. This arrangement was partly a product of the available display technology, and was inspired by the use of space on desks in offices rather than collaborative work practices. Remote collaborators used styluses and bimanual techniques to move and annotate “virtual sheets of paper”, and to gesture to each other with telepointer traces.

RemoteDT [17] presented a large horizontal shared workspace containing a Windows XP desktop. Remote collaborators used direct-touch input to interact with legacy applications, to sketch, and to gesture to each other with telepointers. Multiple co-located collaborators could interact concurrently. TIDL [18] provided similar functionality for legacy Java applications using multiple mice. Regenbrecht et al. [19] demonstrated a remote tabletop system in which remote collaborators could move and rotate photos using touch input. The system did not provide a remote gesture representation such as a telepointer or arm shadow, and did not allow multiple co-located collaborators to interact concurrently.

VideoArms [20] enabled remote collaborators to use touch input to move task artefacts, such as photos, and to sketch. Each collaborator’s arms were captured using a camera and presented to remote collaborators as an image overlaid on the

workspace. The system supported group-to-group collaboration since each surface was large enough for multiple co-located collaborators to stand around, though the touch input system did not allow co-located collaborators to interact concurrently. Distributed Tabletops [25] and Digtible [21] provided similar functionality and remedied this problem using multiple styluses and multi-touch surfaces respectively.

Lastly, C-Slate [22] provided a shared workspace in which remote collaborators could use touch input to reposition task artefacts, such as virtual sheets of paper, and could annotate using a stylus. Like VideoArms, Distributed Tabletops and Digtible, each collaborator's arms were captured using a camera and presented on the surface to remote collaborators. The image of the arm became translucent as the arm was lifted from the surface. The surface was not sufficiently large for co-located collaborators to work side-by-side.

In summary, although the systems all provide large horizontal interactive surfaces, they differ in design: some use indirect mouse input whereas others use direct stylus or touch input devices; some provide remote representations of collaborators' arms whereas others use telepointers or traces; some use interaction techniques such as rotation that favour around-the-table interaction, whereas others focus on support for legacy applications; and some use surfaces sufficiently large for collaborators to work side-by-side while others do not.

Large Format Remote Sketching

Remote tabletop interfaces have their origins in large-format remote sketching interfaces, which in turn were motivated by studies of co-located collaborative design work.

Tang [26] studied the use of shared paper drawings in design meetings on conventional tables. He observed that the process of creating and discussing drawings is often more important than the resulting drawings themselves. He notes that collaborators fluidly interleave drawing and writing with gestures that not only express ideas but also mediate interaction by negotiating turn-taking and focusing attention. Tang's study was preceded by a similar study by Bly [8], and prompted several further studies of co-located collaborative design [27, 28].

The work prompted a shift in remote collaboration technologies away from conventional monitor and mouse interaction and towards larger direct-input remote collaboration interfaces that could afford interleaved sketching and gesturing, unmediated remote representation of hand gestures, and space to mediate interaction.

Commune [29] and VideoDraw [30] linked horizontally-mounted monitors to provide a shared workspace for remote design work. The systems supported simultaneous sketching using styluses, gesturing, and instant visual updates to all sites as users sketched. Commune used a digital stylus system for drawing and

gesturing using a cursor, whereas VideoDraw used video links to enable whole-hand gesturing. In an evaluation of Commune using a two-person design task, the authors observed three novel aspects of the system: both collaborators shared the same orientation of the horizontal workspace; both collaborators were able to mark and point to exactly the same place in the workspace without interference from each other's hands; and collaborators were able to switch seamlessly between writing, drawing and gesturing. They also observed that cursors convey only pointing gestures rather than the rich variety of gestures observed in co-located collaboration, and that the drawing space provided by a monitor was too small.

VideoWhiteboard [6] linked large vertical projected displays to support remote sketching. As with VideoDraw, collaborators could sketch simultaneously and see instant visual updates as others sketched. Instead of whole-hand gestures, collaborators now saw the shadows of their remote collaborators' entire bodies. The authors observed that the larger display now allowed collaborators to work side by side at each site. Although the shadows conveyed many of the gestures used in whiteboard interaction, they observed problems when pointing to precise locations, when conveying subtleties such as head-nods from a distance, and when collaborators' shadows overlapped. This may have been exacerbated by the lack of local feedback of the shadows presented to remote collaborators.

Ishii [31] later demonstrated ClearBoard, which provided the remote sketching task space of VideoDraw, combined with a head-and-shoulders video view of the remote collaborator presented on the same surface. The design imitated collaboration through a clear glass board on which collaborators could draw. The camera at each site was positioned behind the surface in order to achieve greater spatial fidelity.

Remote Augmented Reality Surfaces

Following the early remote sketching systems, researchers created various large-format remote collaboration systems that augmented tangible task artefacts with visual information. Double DigitalDesk [32, 15] was one of the earliest such systems. Each collaborator sat at his or her own desk and interacted with their own paper copy of the same page of information. The system augmented this paper with a video image showing annotations made by the remote collaborator, and an image of the remote collaborator's hands.

Double DigitalDesk followed TeamWorkStation [33], which provided similar functionality but displayed the shared workspace on a conventional vertical monitor, rather than augmenting the paper on the desk itself. Double DigitalDesk inspired a number of later projects that pursued the approach of augmenting physical paper with remote annotations, such as Tele-Graffiti [34], LivePaper [35], Agora [36], and PlayTogether [37]. Kirk et al. [38] demonstrated a similar system for remote physical assembly tasks on horizontal surfaces.

Augmenting tangible artefacts to create a shared workspace tends to lead to problems maintaining consistency across the two remote sites. It is not clear how Double DigitalDesk, for instance, addresses the problem of one collaborator moving their copy of a shared document in the workspace, or turning the page of a multi-page document; the remote annotations would then project onto the incorrect page. Some interfaces, such as PlayTogether, have addressed this by opting not to replicate tangible task artefacts at each site. Instead, they have presented each collaborator with a projected image of tangible task artefacts that are not physically present at their own site. In such systems, only the collaborator who is co-located with a given task artefact can manipulate it; the other collaborators can only see it (typically at a low display resolution) and gesture to it. This asymmetry is inherent in tasks such as remote surgery, and advantageous in tasks such as remote bomb disposal. It may, however, be problematic for the kinds of information tasks traditionally performed on tables, in which collaborators share task artefacts by passing them among themselves. Remote tabletop interfaces avoid the issues of asymmetry and consistency by using computer-generated imagery instead of tangible objects.

Empirical Findings

There have been relatively few empirical studies of remote tabletop interfaces. In contrast to the wealth of knowledge about collaboration around conventional tables and tabletop interfaces, we know relatively little about remote tabletop collaboration.

The majority of findings relate to the design of the remote arm representation. Firstly, Kirk et al. [39] conducted a number of experiments using an instructor-follower remote physical assembly task and a remote augmented reality surface. The follower participant sat at a table and assembled physical components on the table by following instructions from the instructor participant. The instructor sat at a different table and saw a video image of the follower's workspace. A video image of the instructor's arms was projected onto the follower's table, allowing the instructor to gesture to the follower. The authors found that this unmediated video representation of hands could convey a variety of complex gestures [38]. Furthermore, Kirk et al. [39] showed that this representation yielded performance benefits over a video-only condition with no gesturing. Their analysis of language showed that gesturing enabled deictic utterances that may have replaced lengthier descriptive utterances, and that gesturing enabled turn-taking. A further study using a similar task [40] compared this unmediated video representation of the hands with an alternative in which the instructor could sketch on the follower's workspace, but could not gesture with their hands. The hand gesture representation yielded faster task completion than the sketching approach or a combination hands and sketch approach, with no loss of accuracy.

Secondly, Tang et al. [20] hypothesised that high-fidelity remote arm representations in the shared workspace enable collaborators not only to gesture to each other but also, when combined with direct input devices, to infer the actions being generated by the movement of the arm. Impoverished, indirectly-controlled representations such as telepointers do not, for instance, allow collaborators to see each other reaching with their hands for a task artefact or tool, or retracting their arms from the display to think. In an exploratory study of design and puzzle tasks using remote arm representations and direct input devices, they observed that remote collaborators spent considerable periods of time watching each other work, which supports this hypothesis.

A number of further studies relate to remote tabletop interfaces more generally, but do not address the issue of work practices. Escritoire [16] was trialled using a laboratory study involving three pairs. In questionnaire responses, they all agreed strongly that the shared workspace was useful, and were less sure that a video view showing the remote collaborator's face was also useful. Hauber et al. [41] investigated a decision-making task and compared three conditions: co-located tabletop collaboration; a conventional GUI remote collaboration interface with a shared workspace and a head-and-shoulders view of the remote collaborator; and remote tabletop collaboration, sat at opposite ends of the table, with a vertical screen showing the head-and-shoulders view of the remote participant. Participants in the remote tabletop condition talked more about the technology, took longer to complete the task, reported being more aware of their remote partner, and reported feeling more like they were in the same room.

Pauchet et al. [42] used Digtale to compare remote and co-located tabletop collaboration using a puzzle task. Remote collaborators again were provided with head-and-shoulders video of each other using vertical screens. They reported faster task completion times in the remote interface, though it is not clear whether this generalises to other tasks.

Work Practices in Remote Tabletop Collaboration

The works reviewed in the previous section establish that remote tabletop interfaces are technically feasible, and suggest that high-fidelity remote arm representations may help remote collaborators to gesture and to remain aware of each other's actions.

However, it remains unclear whether they support collaborative activities to the same extent as the conventional tables and tabletop interfaces that they mimic. In particular, tabletop interfaces have been found to be successful at supporting collaboration in part because they can afford some of the beneficial work practices observed in collaboration at conventional tables. It is not clear from the prior work whether it is possible to design remote tabletop interfaces to afford these work practices.

We suggest this is a key topic for further exploration in this area. With this in mind, we review two collaborative work practices that have been observed in collaboration at conventional tables and at tabletop interfaces, and discuss whether they may also be afforded by remote tabletop interfaces.

Workspace Awareness

The ability to maintain an awareness of other collaborators' actions in the workspace is central to many collaborative tasks, and underpins a number of beneficial work practices. Co-located collaborators use a variety of visual and auditory cues to maintain awareness, whereas remote collaborators are reliant on the cues conveyed by the system. Gutwin and Greenberg [7] describe how “[poorly-designed remote collaborative systems] often seem inefficient and clumsy compared with face-to-face work”. Consequently, the generation, presentation and use of awareness cues are important factors in the design of a remote collaboration system.

More formally, Gutwin and Greenberg define workspace awareness as “the up-to-the-moment understanding of another person’s interaction with the shared workspace” [7]. They argue that workspace awareness underpins a number of beneficial collaborative work practices. One such work practice is *collaborative coupling*, the way in which collaborators move regularly between working individually and working closely together as part of a shared task. These changes tend to be opportunistic and unpredictable [43]. For instance, collaborators may switch from working individually to working closely together in order to discuss a decision or because one collaborator needs the other’s involvement in order to make progress with their work. Gutwin and Greenberg argue that supporting this work practice is important, but difficult to achieve in practice. The system must provide awareness cues so that each collaborator remains aware of the state of their peers and of the task in order to be able to instigate discussion or individual work at appropriate points, and to identify when their peers are trying to do so. Workspace awareness therefore underpins collaborative coupling. A number of other work practices also rely on workspace awareness, such as coordinating intricate actions when working closely together; anticipating the actions of collaborators; and identifying appropriate times to offer assistance [7].

Gutwin and Greenberg [7] suggest that collaborators not only maintain awareness by conversing and gesturing, but also by directly and peripherally watching each other work. In particular, they gain awareness through *consequential communication* (watching a collaborator’s arms) and from *feedthrough* (watching changes in task artefacts as they are manipulated).

In an observational study of a group design task at a conventional table, Tang [26] observed how peripheral awareness enables collaborators to coordinate intricate actions and switch appropriately between working individually and working

closely with others. Similar behaviour has been observed in studies of tabletop interfaces, for a variety of tasks including information-gathering and planning [12], photo-layout [13], and route-planning [14]. This suggests that, like conventional tables, tabletop interfaces can afford high levels of workspace awareness and so the practice of collaborative coupling.

Workspace awareness principles have been applied to the design of remote arm representations in remote tabletop interfaces. As described earlier (see “Empirical Findings”), Tang et al. [20] suggest that remote arm representations can help collaborators watch each other work, and hence aid consequential communication. However, in order to design for awareness we must consider not only the arm representations but interaction more generally. For example, in conventional groupware, consideration of awareness cues when collaborators manipulate task artefacts and navigate the workspace is also important [44]. It is unclear how remote tabletop interaction techniques impact workspace awareness, and so we do not understand how to design remote tabletops to afford the desirable work practices that depend on it.

Territorial Coordination and Seating Arrangement

A number of studies have shown that collaborators at conventional tables partition the workspace into regions. When sat around a conventional table, each collaborator has a distinct area of table in front of them in which they carry out individual work as part of the task. Tang [26] observed a collaborative design task using paper on a conventional table, and noted that designs sketched in this area were “within a personal boundary and not intended for others to perceive”. Scott et al. [45] observed a similar task and observed that participants moved task artefacts, such as paper or scissors, into this personal territory to reserve them for themselves, and would later return them to the table centre to indicate their availability to the group. Collaborators also implicitly partitioned space in the centre of the table so that each adopted responsibility for the nearest region. Similar territorial behaviour was observed in a tabletop interface using a collaborative photo-layout task [13]. Collaborators would, for instance, move a virtual container of digital photos to the table centre when working together, and when working individually would move containers into and out of their personal work area as necessary, without disrupting their partner. Territoriality is therefore a prominent coordination mechanism in tabletop collaboration.

However, it is unclear whether remote tabletop interfaces will afford this practice. Territoriality among remote tabletop collaborators may depend upon seating arrangement. They can arrange themselves either around the table, as if co-located, or alternatively can both sit at the same place relative to the workspace. This second, *overlaid*, arrangement is unlikely to afford territoriality since collaborators’ personal areas will overlap, and so may lead to coordination difficulty.

Seating preferences in remote tabletop collaboration are unclear. Kirk and Stanton-Fraser [40] compared overlaid and non-overlaid arrangements in a remote instructor-follower physical-assembly task. They found a non-significant trend that the overlaid arrangement was easier and led to more task progress. However, the instructor-follower character of the task meant that territorial coordination was not possible. By contrast, co-located collaborators tend to interact at a distance that preserves their personal space [46].

Exploratory Study

Having identified work practices as a key area of further exploration for remote tabletop interfaces, we now present an exploratory study that investigates this area. The full results are described in our previous publication [47].

This study seeks to understand:

- Whether remote tabletops afford the coordination mechanism of territoriality, and whether this is affected by seating arrangement.
- How the design of remote tabletops impacts workspace awareness and the work practices that depend on it.

Method

Technology

The study used the Distributed Tabletops system [25] (Figure 1). The system provides a large shared workspace in which co-located or remote collaborators can move, reorient, and otherwise manipulate task artefacts using digital styluses. All collaborators may interact concurrently. Collaborators' arms are displayed using a translucent remote "shadow" representation. Each collaborator's shadows are also shown locally in order to provide local feedback of the remote representation. The system runs at 60 fps with a latency of around 100 ms. Arm shadows are captured at 15 fps. Sites were also linked using a speakerphone.

Study Design

The study compared three conditions (Figure 2):

- *Co-located-adjacent (CA)*: Collaborators sat at adjacent corners of the same table. Each had their own stylus, and both could interact concurrently.
- *Remote-adjacent (RA)*: Collaborators sat in separate rooms at tables which were linked using the system and using a speakerphone. Collaborators were again positioned at adjacent corners. Again, both collaborators could interact concurrently.
- *Remote-overlaid (RO)*. As remote-adjacent, except that collaborators were positioned in the overlaid arrangement.

The study followed a within-subjects design. Each pair of participants tested each of the three conditions, and the presentation order was counterbalanced across pairs using a Latin square. Pairs attempted each of three design briefs in turn using the system; one in each condition. The presentation order of design briefs was counterbalanced across the conditions and pairs. Data was captured using system log files and video recordings. Post-condition and post-study questionnaires asked about preferences and difficulties; these were further explored in semi-structured interviews.

Task and Interaction

The study used a furniture layout task in which participants were asked to work together to fulfil a design brief by arranging diagrammatic furniture on a floor plan. For example, one such brief asked participants to design a communal space for a new library and to provide, among other things, as much seating as possible, and areas for photocopying. The other briefs were similar and asked for designs for a common room and for a research lab. Participants were asked to colour furniture items to illustrate the different parts of the brief. After each task, participants were asked to give a short joint presentation and to answer questions about their solution.

This task is representative of various tabletop design tasks. It requires discussion and manipulation of 2-D task artefacts in order to explore different approaches and tradeoffs. Participants were also instructed to apply their own prior knowledge of communal spaces. This led to implicit constraints which, along with the tradeoffs and joint presentation, ensured that collaborators had to maintain an

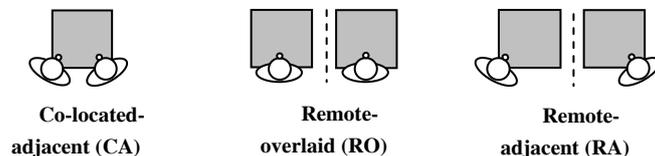


Fig. 2. Study conditions. From [47]; © 2009 ACM, Inc. Included here by permission.

awareness of each other's actions and to coordinate if they were to produce a mutually-satisfactory design.

The shared workspace presented a 75cm x 75cm blue square containing an empty white floor plan, "piles" of diagrammatic furniture nearest the participants, a task brief, and a key explaining the furniture representations (Figure 1). At the task outset, the floor plan was empty except for lines marking room boundaries.

Each participant had a stylus with which they could move any of the task artefacts (including the plan). The system used the popular Rotate 'N' Translate interaction technique [48], which uses a pseudo-physics model to enable task artefacts to be rotated and translated simultaneously in a single stylus stroke. Once moved onto the floor plan, furniture would snap orientation to multiples of 45° and would "stick" to the floor plan if the plan were moved. The colour of an item of furniture could be changed by tapping twice with the stylus to open a colour menu, and then tapping on the appropriate menu item.

Participants and Procedure

18 paid participants aged 20-39 were recruited from a Computer Science department to form 9 pairs (16 males; 2 mixed-sex pairs). Participants in each pair reported having met previously. Two participants were left-handed. One had limited experience with tabletop interfaces; the others had none.

Each pair was first given a short tutorial about interacting with the system, after which each participant practiced individually until comfortable. Once the experiment structure had been explained, the pair completed each condition in turn. In each condition, the participants were arranged appropriately and asked to stay in their seats and not to move the seats. Participants practiced together until comfortable using a practice brief, to minimise learning effects during the recorded sessions. All pairs took about 10 minutes before their first condition, and about 2 minutes thereafter. The session workspace was then loaded. Participants were instructed to work together to arrange the task artefacts to fulfil the brief to the best of their abilities. They were advised that it might take between 15 minutes and half an hour, and that at the end they should make a short presentation and answer questions on their finished design.

After the presentation and questions, each participant completed a post-condition questionnaire. Once all three conditions were complete, each participant individually completed a post-study questionnaire, and the pair took part in a semi-structured interview.

Results: Initial Observations

Pairs in each condition worked for an average of 21 minutes. Participants often used their prior experiences (for example, that photocopiers should be situated away from working areas because of the noise). This led them to explore different layouts and adjust their design as they proceeded, to produce an outcome that was most appropriate in the context of their own prior experience. We observed both individual and group work, and both turn-taking and concurrent work.

Pairs in all three conditions appeared to have few difficulties interacting with the system. The post-condition questionnaire used 7-point Likert scales to ask about ease of task completion, ease of communication, and the extent to which the pair worked together (Table 1). Friedman rank tests for repeated-measures ordinal data did not yield significant differences among conditions.

At the end of the study, participants were asked individually which of the remote conditions they preferred. 11 preferred overlaid, 6 preferred adjacent, and 1 had no preference. A chi-squared test yielded no significant overall preference.

Results: Territorial Coordination

System log files were used to generate *activity maps*. Each map shows the locations of one pair’s interactions with furniture in the workspace in a single session [45]. Figure 3 shows activity maps for a single representative pair. The black and white marker distributions in the co-located-adjacent conditions show a trend for partitioning the floor plan and surrounding space according to proximity, so that each participant worked broadly in the half of the table nearest themselves. By contrast, the marker distributions in the remote cases suggests that participants broadly partitioned the workspace into regions in which they worked, but the par-

Question	CA	RO	RA
“We worked together throughout the task.” (1=strongly agree, 7=strongly disagree)	2.0 (0.8)	2.1 (1.2)	2.1 (0.9)
“How easy or hard was the task to complete using this technology?” (1=very easy, 7=very hard)	2.5 (0.9)	2.2 (0.9)	2.4 (0.8)
“How did you find communicating this way?” (1=very easy, 7=very hard)	1.7 (0.8)	1.9 (0.6)	2.3 (0.8)

Table 1. Mean (standard deviation) Likert scale responses. From [47]; © 2009 ACM, Inc. Included here by permission.

tioning forms a patchwork rather than a strict left-right arrangement.

Such a quasi-naturalistic open-ended task inevitably leads to variations in partitioning, depending on how activity unfolded at the time, and so it is difficult to make inferences from the visual data alone. Accordingly, a quantitative statistical analysis of the extent of left-right partitioning was performed on the underlying location data. We began by quantifying the extent to which participants' interactions were partitioned to the left and right.

Following the approach of Scott et al. [45], we calculated the proportion of interactions carried out by each participant on each side of the table to find, for instance, that the left side was 30% participant A and 70% participant B. However, we cannot use this figure to infer partitioning because B may have been more ac-

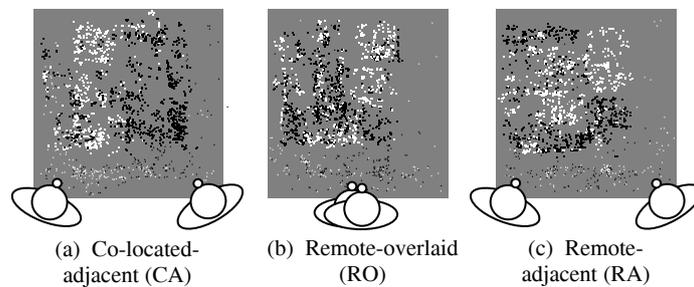


Fig. 3. Activity map showing interactions of one pair in the workspace. Each marker corresponds to a task artefact being picked up or dropped. Colour indicates the person interacting. From [47]; © 2009 ACM, Inc. Included here by permission.

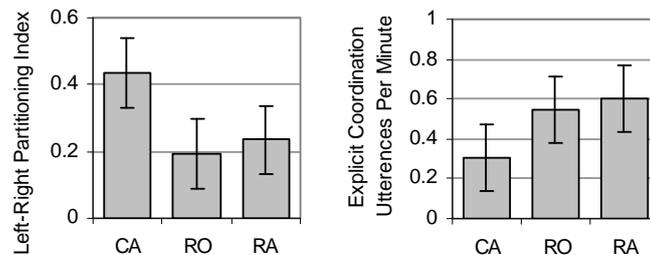


Fig. 4. Left-right partitioning index and rate of coordination utterances for each condition. Error bars are 95% confidence intervals and indicate variation within pairs (i.e. the error considered by a repeated-measures ANOVA). From [47]; © 2009 ACM, Inc. Included here by permission.

tive than A on both sides. Similarly, we may calculate that, for instance, 80% of B's interactions were on the right-hand side of the table, but again we cannot infer partitioning since 80% of A's interactions may also have been on that side.

Instead, we desire an alternative metric to measure the extent of partitioning between collaborators. We calculate the *left-right partitioning index* by first taking the proportion of A's interactions that lie within the left half of the table, and then the same for B's, and then taking the absolute difference between the two figures. For example, if 80% of A's interactions fell on the left side, and 30% of B's, the resulting index would be $|0.8-0.3| = 0.5$. If A's interactions were entirely on the left, and B's entirely on the right, the index would be $|1.0-0.0| = 1.0$. If both participants interacted to the same extent on the left, the resulting index would be 0. Accordingly, this metric is an aggregate measure that we can use quantitatively to highlight differences in partitioning.

The partitioning indices across different conditions are shown in Figure 4 (left). The trends are consistent with the patterns observed in the activity maps. The mean index for the co-located-adjacent condition is 0.4, corresponding to a 70%:30% split between sides of the table. The mean index in both remote conditions is 0.2, corresponding to a 60%:40% split. This is perhaps surprising: the extent of partitioning in the remote-adjacent case is quantitatively indistinguishable from the remote-overlaid case, in which participants were overlaid and hence could not partition by proximity.

The difference among the conditions was significant using a one-way repeated-measures ANOVA ($F(2,16)=7.02, p<0.01$). Pairwise t-tests found significant differences between CA and RA ($t(8)=3.64, p<0.01$), and between CA and RO ($t(8)=2.83, p=0.022$).

The activity maps shown in Figure 3 show the location of interactions relative to the shared workspace. Figure 5 shows the marker locations transformed to appear relative to the participants' final furniture layout. This provides some insights as to how the remote participants spatially partitioned the task into the patchwork arrangement: the markers are often partitioned according to either the "walls" on the floor plan at the task outset, or to "islands" of furniture or new "walls" created by the participants during the task.

That participants in the remote conditions did not partition based on proximity, as is the social norm, suggests that they may have worked to explicitly coordinate their activities. Although much of the coordination seemed implicit in the activities of the participants, they sometimes coordinated with explicit utterances to make clear to each other the activities they had completed or were about to start (e.g. "I'll do the common room now", "I've finished doing the windows", "You can start on the secretary's room"). We coded the dialogue and found that such explicit coordination utterances were twice as frequent on average in the remote conditions (Figure 4, right). The difference between conditions was significant using a one-way repeated-measures ANOVA with the conservative Greenhouse-Geisser sphericity correction ($F(2,16)=6.142, p=0.03$). Pairwise t-tests identified signifi-

cant differences between CA and RO ($t(8)=4.25$, $p<0.01$), and between CA and RA ($t(8)=2.40$, $p=0.04$).

Results: Workspace Awareness

We now consider workspace awareness, and the beneficial work practices that rely on it. As described in the review of empirical findings, prior research has investigated intentional gesture in detail, and so we focus on the remaining sources of workspace awareness, namely consequential communication and feedthrough. We begin by examining the work practice of collaborative coupling, which relies on workspace awareness (described in the review of work practices).

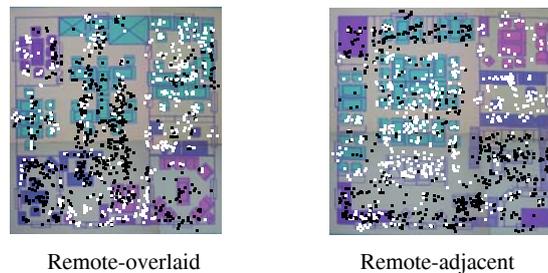


Fig. 5. Interactions of one pair overlaid on their final furniture layouts. Furniture layouts and “walls” are shown as lines and coloured areas under the markers; participants themselves coloured the furniture during the task. Each marker corresponds to a task artefact being picked up or dropped. Marker colour (black or white) indicates the person interacting. From [47]; © 2009 ACM, Inc. Included here by permission.

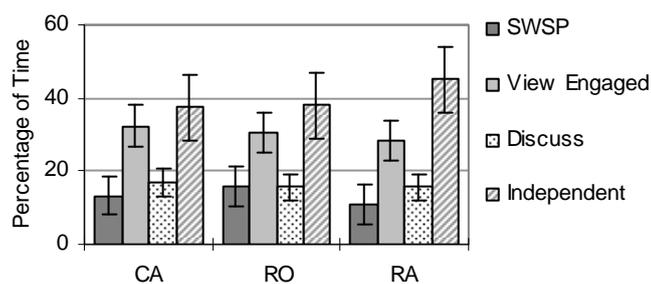


Fig. 6. The proportion of time spent in each coupling style in each condition. Error bars indicate variation within pairs and show 95% confidence intervals. From [47]; © 2009 ACM, Inc. Included here by permission.

Coupling Styles

We iteratively refined a coding scheme for the coupling styles observed in the present study. Randomly-selected segments of video data from each condition were repeatedly analysed using a similar approach to prior exploratory studies of co-located collaboration [45, 14]. The initial coding categories were informed by field notes and by the co-located tabletop coupling styles identified by Tang et al. [14]. This initial analysis yielded four distinguishable coupling styles that exhaustively and mutually-exclusively classify the coupling arrangement at any time:

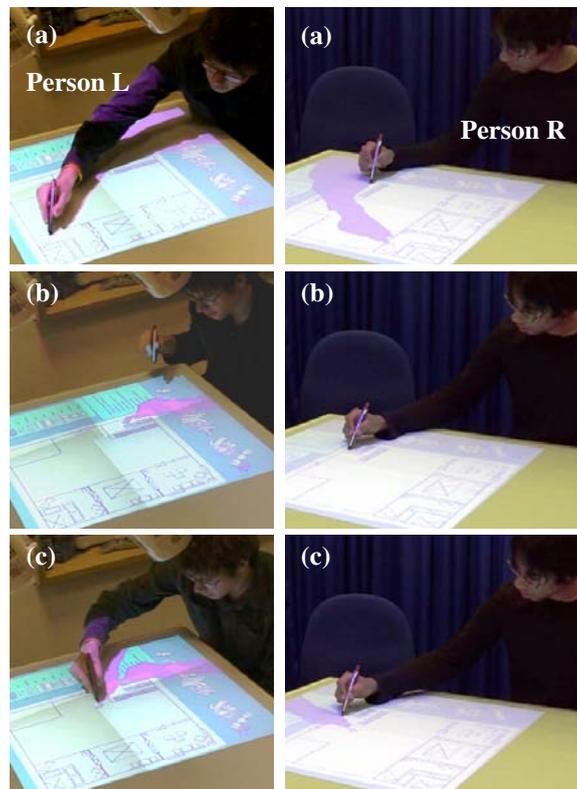
- *Simultaneous work on the same problem, (SWSP)*: the partners are actively working together to solve the same problem, such as both creating a wall.
- *View engaged*: As Tang et al. [14]. Though the partners are working together, one is watching closely while the other manipulates the display, for instance while taking turns or demonstrating ideas to each other.
- *Discuss*: the partners are discussing ideas together and neither is manipulating the display
- *Independent work*: the partners are working independently while either manipulating or looking at the workspace. Partners may be conversing or working in silence.

We explored whether the different conditions affected the ability to work in the different coupling styles. We used the above scheme to code the entire video of every session to calculate the proportion of time spent in each coupling style (Figure 6). There was large variation among different pairs: some tended to be predominantly closely-coupled throughout; others tended to work independently. This is to be expected in a quasi-naturalistic task. Nevertheless, large differences among conditions in the proportion of time spent in each style were not apparent, and analysis using repeated-measures ANOVAs found no significant differences.

This result is perhaps surprising: working remotely did not seem to have a practical impact on collaborators' behaviour at this aggregate level, in a fairly representative design task. Participants seemed able to work independently and closely coupled to the same extent regardless of whether co-located or remote.

Coupling Transitions

Our observations of the video record suggest that participants in all conditions were able to change fluidly between different coupling styles. The changes seemed swift and opportunistic and were often preceded not by explicit gestures or conversation but instead by one collaborator watching another's arm movements (consequential communication) or another's manipulation of task artefacts (feedthrough).

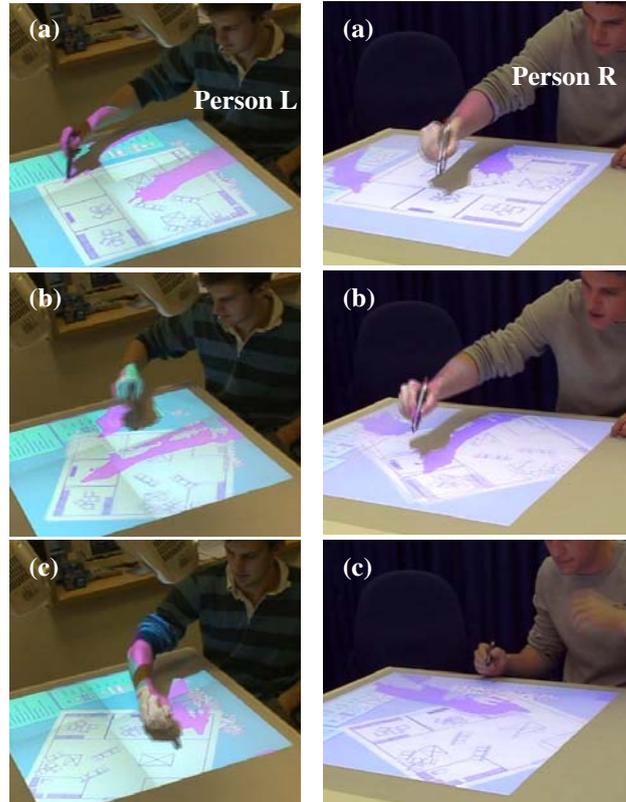


(L and R work independently. L glances at R's work.) (a)
 (L stops work and watches R.) (view engaged) (b)
 L: eh it's good need a door
 (L assists R by adding a door to the room that R is arranging.) (simultaneous work on the same problem) (c)

Fig. 7. A series of coupling transitions. From [47]; © 2009 ACM, Inc. Included here by permission.

Figure 7 shows a representative example. The collaborators L and R are working independently, and in different parts of the workspace (Figure 7a). L notices R working in the corner and starts to watch (view engaged, Figure 7b). L then begins to help R by moving a piece of furniture into position, while ensuring that his arm does not block R's activity (simultaneous work on the same problem, Figure 7c). At this point, the collaborators begin to talk. The video also revealed similar instances of anticipation and assistance, such as by watching closely and then instigating discussion, or taking turns to demonstrate different ideas to each other.

Support for different coupling styles and fluid transitions between is desirable and suggests high levels of awareness. Yet it is not clear how this arose as a result of the design of remote tabletop interfaces. However, the video also revealed that two occasionally-used interaction techniques tended to cause awareness problems. The use of the techniques by one partner tended to be unanticipated by the other,

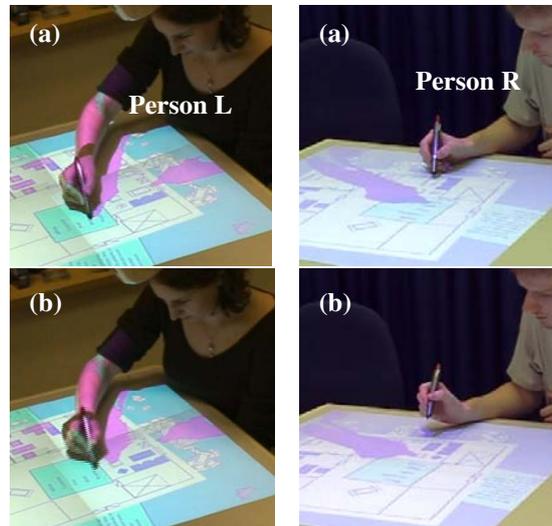


L: do you want to spin the plan so you can get at it more easily
 (R reaches towards the top of the plan start working) (a)
 (L starts to rotate the plan. R has not anticipated this and has
 to abort the reaching action.) (b)
 (L still rotates the plan. R backs off looking puzzled.) (c)
 R: oh I see yeah

Fig. 8. Deliberate movement of the floor plan by one participant
 was unanticipated by their collaborator. From [47]; © 2009
 ACM, Inc. Included here by permission.

and in contrast to the otherwise-fluid transitioning between coupling styles, this resulted in confusion that had to be resolved verbally. Though undesirable, these techniques offer an opportunity to gain insights by exploring how they impaired workspace awareness while the remainder of the system afforded the otherwise-high awareness level.

The first occasionally-used problematic technique was the movement of the large floor plan. Like the furniture items, the plan could be simultaneously translated and rotated by touching any part of it with the stylus and then dragging, using the popular Rotate 'N' Translate technique [48]. Observations of the video suggest that movement of the large floor plan by a collaborator was often unanticipated by their partner. Figure 8 shows a representative example. L states his



(L has opened a menu and is about to press on it.) (a)
 (R opens a menu as L presses, causing L to press on this new menu instead of her own menu. Consequently, the wrong item of furniture changes colour.) (b)
 L: oh?

Fig. 9. Opening of a menu by one participant was unanticipated by their partner. From [47]; © 2009 ACM, Inc. Included here by permission.

intention to spin the floor plan while R begins to reach towards the top of the plan to interact (Figure 8a). However, as R's hand nears the top of the plan, L begins to spin it (Figure 8b). R stops the reaching action, retracts his hand (Figure 8c) and says "oh I see yeah".

Though the participants quickly recovered, this sequence contrasts with the smooth coupling transitions observed throughout the majority of the sessions as participants moved furniture (for example, Figure 7). This is curious because both the movement of the furniture and the movement of the floor plan used the same Rotate'N'Translate interaction technique. The difference may result from the large size of the floor plan. Because furniture items are small, a collaborator has to reach towards a furniture item in order to interact with it. This reaching motion can be peripherally observed by their partner, either in the motion of the arm (if co-located) or the shadow (if remote). This consequential communication enables them to anticipate the action. By contrast, the movement of the floor plan can be instigated from any point within its large area. Because it does not require reaching to a particular point, the action is likely harder to anticipate. In Figure 8a, for example, R cannot infer from his view that L is about to rotate the plan. The large size of the floor plan may also impact negatively in a second way. Furniture items

are small and so any movement of them is localised to the vicinity of the interacting arm or arm shadow. A collaborator can safely assume that if their partner's arm or arm shadow is far from their own then their actions will not interfere. This is not the case for the large floor plan, since it is so large that its movement is not localised near to the hand causing the movement.

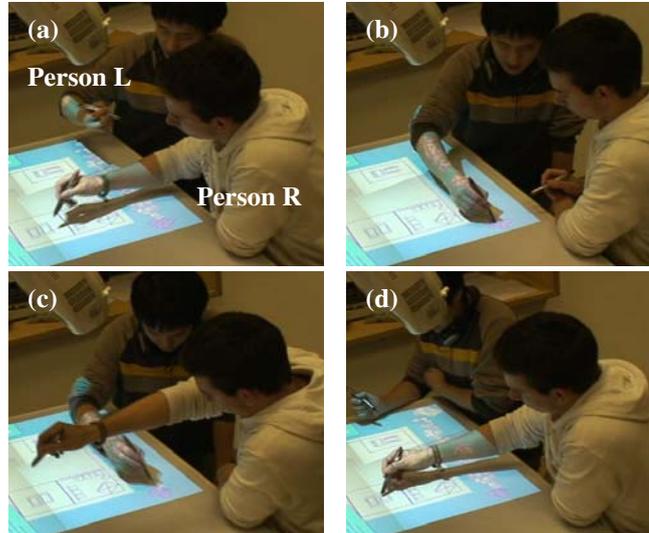
The second problematic interaction technique was the colour menu, which opened when a participant tapped twice on a furniture item. Like the movement of the floor plan, this action often seemed unanticipated by the instigator's partner. Figure 9 shows a representative example. L has opened a colour menu (Figure 9a) but as she presses on it to choose a colour, R opens another colour menu that overlaps the original. L unintentionally presses on the overlapping menu instead. This selects the wrong colour and the wrong furniture item. The collaborators are confused and discuss what happened (Figure 9b). Again, though quickly resolved, this contrasts with the fluid coupling switches observed when moving furniture (Figure 7). Like the movement of the floor plan, this problem may again be explained using the awareness framework. Figure 9a illustrates that the tapping action by R to open the menu is not obvious from the arm shadow seen by L, and so the action provides no consequential communication. Furthermore, there is no feedthrough as the action unfolds because the menu appears instantly once the tapping action is complete. By comparison, the movement of the furniture is continuous and so is perceptible as the action unfolds, providing feedthrough. Furthermore, the furniture moves with the shadow of the dragging arm, providing consequential communication.

Discussion

Territorial Coordination

The study condition significantly affected the collaborators' partitioning of the workspace. Participants in the co-located-adjacent condition partitioned the floor plan and surrounding space broadly according to who was nearest. This proximity partitioning is consistent with the territoriality findings of Scott et al. [45]. Their study observed two pairs at conventional tables. The left-right partitioning indices from their data can be calculated as 0.51 and 0.42, which is consistent with our results.

Participants in the remote-overlaid condition were not able to partition and delegate by proximity since their personal regions overlapped. Nevertheless, each pair partitioned the floor plan into a patchwork of areas in which one or the other participant worked. The partitioning followed the boundaries visible at the start of the task, and new boundaries created using the task artefacts.



(R and L are working independently.) (a)
 L: yeah I'll start doing pigeon holes
 R: okay
 (R leans back and L leans across and takes a bookshelf.) (b)
 L: um bookshelves are pigeon holes can I just steal like loads?
 (L leans across to take more and R tries to work.) (c)
 (L finishes taking and R leans forward to work again.) (d)
 L: cheers

Fig. 10. Co-located-adjacent collaborators have difficulty working across each other. From [47]; © 2009 ACM, Inc. Included here by permission.

Participants in the remote-adjacent condition behaved much as those in the remote-overlaid condition; the adjacent seating arrangement did not lead to the proximity-based partitioning observed in the co-located-adjacent condition. This indicates that territorial coordination may arise from more than seating arrangement alone.

Video observations of participants reaching may help explain this difference between the co-located-adjacent and remote-adjacent conditions. Participants in the co-located-adjacent condition seemed to avoid reaching across their partners, perhaps because this would block them from working. Before reaching they would often wait for an opportune moment, and sometimes asked permission (Figure 10). In the remote conditions, by contrast, partners would unhesitatingly work directly across each other, work within the arm shadows themselves, and reach to take furniture from immediately in front of each other (Figure 7). This blocking hypothesis helps explain the differences in partitioning and is supported by recent findings in the literature. Nacenta et al. [49] found that action-at-a-distance techniques (like

“tractor beams”), that avoid blocking problems, lead to less proximity-based partitioning in co-located tabletop collaboration than when using conventional techniques like Rotate+Translate. Similarly, Ha et al. [50] found that when using a mouse (which avoids blocking problems), a majority of participants reported being more likely to interact with objects on their partner's side of the table than when using a stylus.

These results suggest that, unlike co-located collaboration at conventional tables and tabletop interfaces, remote tabletop collaboration does not afford the work practice of territorial coordination. As described earlier, social norms dictate that space immediately in front of a collaborator is reserved as an area in which they can try ideas away from the group, and can move task artefacts to reserve them for their use [45]. This cannot be the case in the remote-overlaid condition, because these regions are overlaid. That remote-adjacent collaborators unhesitatingly took task artefacts from immediately in front of others, suggests that this condition also does not afford personal territories.

We have described two further results. Firstly, like Kirk [40], there was a non-significant trend for remote collaborators to prefer the overlaid arrangement. 10 of the 11 who expressed this preference reported having difficulty reaching parts of the table when remote-adjacent. However, this difficulty was not reported when co-located-adjacent, despite the participants being sat in the same positions relative to the workspace. Video observations suggest this effect may be caused by a limitation of the remote gesture system. Co-located participants would point to the far corner of the table, using the 3-D trajectory of their finger or hand to avoid having to reach. The remote arm representation did not convey these 3-D depth cues and so remote collaborators would instead gesture by reaching out and hovering their hands over the far corner. This problem was most acute in the remote-adjacent condition because participants were seated to one side, and may have led to the preference for remote-overlaid.

Secondly, we showed that the patchwork partitioning in the remote cases coincided with more frequent explicit coordination utterances. This suggests that the lack of support for social norms may require greater formal coordination effort.

Workspace Awareness

Our investigation of workspace awareness began by examining collaborative coupling. As discussed in the review of work practices, this is a beneficial work practice that relies on a high level of workspace awareness and is afforded by co-located tabletop collaboration. We identified four coupling styles and found no significant differences among conditions in the proportion of time spent in each style. Furthermore, the confidence intervals on these proportions indicate that moving among experiment conditions impacted little on collaborators' abilities to work in the different styles. This contrasts with Tang et al. [14], who observed a

fifty percentage-point difference in the proportion of time spent in similar coupling styles as a result of manipulating the design of a tabletop interface. We would not claim that the condition has no impact on ability to work in different styles in general; an alternative task may result in a more pronounced difference. Nevertheless, for this fairly representative design task, the conditions seemed to have relatively little impact at an aggregate level.

Collaborators' abilities to switch fluidly between coupling styles seemed dependent on the interaction techniques. In particular, two techniques used occasionally tended to impair workspace awareness and to lead to confusion that was resolved verbally. Though the collaborators recovered quickly, these instances provide insights into how the design afforded the otherwise-high level of awareness. Firstly, the movement of the floor plan lacked consequential communication because the action could be instigated from any point without having to reach first, and so was hard to anticipate. Secondly, the opening of the menu lacked consequential communication because the instigating double-tap action was not conveyed by the arm shadow; and lacked feedthrough because the menu appeared instantly and at the end of the action.

By comparison, most coupling changes seemed smooth and opportunistic, mediated by consequential communication and feedthrough as collaborators moved furniture items. Partners also offered each other assistance at appropriate times and closely-coordinated their actions when working in the same part of the floor plan, providing further indication of a high level of workspace awareness. The movement of a furniture item provided consequential communication because the collaborator had to reach towards an item in order to move it, and because the movement of the item was then coupled to that of the interacting hand. The action further provided feedthrough via the continuous movement of the item, which could be observed by collaborators as the action unfolded.

Implication #1: Visual boundaries to aid coordination

Whereas co-located tabletop collaborators coordinate based on proximity, remote tabletop collaborators partition the task in a patchwork arrangement based on visual boundaries. Furthermore, our observations suggest that neither remote arrangement affords a personal territory in which to try ideas away from the group and to reserve task artefacts. Applications and techniques designed for co-located tabletop collaborators may therefore encounter coordination difficulties if applied naively in a remote tabletop setting. One way to address this problem may be to provide flexible visual boundaries for remote tabletop collaborators. For instance, each collaborator could be provided with their own moveable coloured palette region to act as a personal territory onto which task artefacts could be placed.

Implication #2: Localised, continuous visual changes

Our observations of interaction techniques suggest that a high level of awareness and the resulting beneficial work practices can be afforded when interaction techniques:

- localise visual changes to the vicinity of the interacting hand;
- and localise the area in which the action can be instigated;
- and produce continuous visual changes.

We observed how two interaction techniques proved difficult to anticipate, and led to confusion, because they did not satisfy these properties.

Many interaction techniques demonstrated in prior tabletop interfaces satisfy these properties, and the resulting high awareness level may contribute to the effectiveness of these interfaces in affording collaborative activities. For instance, techniques like Rotate'N'Translate [48], for moving small photos, are localised and produce continuous visual changes, similar to the movement of the furniture items in this study. Similarly, dragging from a pile to create new task artefacts, dragging to a "recycle bin" to remove task artefacts, and using individual movable lenses [14, 51] to view different visualisations all satisfy these properties.

The properties can also be applied to suggest ways of improving the two problematic interaction techniques in the current study. Instead of double tapping to open the menu, for instance, the menu could open gradually while the collaborator drags from a point, so that the movement of the hand and menu are observable as the action unfolds. The menu could alternatively be kept permanently open as a moveable box or a toolglass [52]. Like the furniture items, this permanent menu would lead to visible reaching actions and, when moved, would provide continuous localised visual changes. When dragging a large task artefact like the floor plan, the visual effects can no longer be localised to the dragging hand because of its size. Nevertheless, a small "drag handle" can be attached to the task artefact to localise the area in which the action can be instigated, so that collaborators can see each other reaching towards the handle and so anticipate the action.

Our suggestion that continuous visual changes may improve awareness in tabletop interaction techniques follows a similar proposal of Gutwin and Greenberg [44] for conventional GUI groupware. However, we believe that localising interaction has not previously been explored as a tool to improve workspace awareness. Moreover, just as Gutwin and Greenberg [44] argue that enforcing continuous visual changes in conventional groupware can increase awareness but negatively impact individual power (e.g. by prohibiting fast keyboard shortcuts), so enforcing locality may lead to a similar tradeoff: localising the point at which an action can be instigated may boost awareness through increased reaching, but requires extra reaching effort by collaborators. This may impact productivity and ultimately lead to fatigue. Prior experience suggests that a reasonably balance may exist, because collaborators at conventional tables frequently reach to acquire task artefacts. This is supported by Morris et al. [53], who explored reach in a tabletop

interface by comparing a single central shared menu against replicated menus near to each collaborator. When using the central shared menu, participants were concerned not by the ergonomics of reaching, but by the socially-awkward proximity of their hand to others' hands. This physical proximity arose perhaps because the task required frequent menu use by different collaborators, but would not occur in remote tabletops.

Future Trends

There are considerable opportunities for further research in this area. Firstly, the limitations of the exploratory study offer a number of opportunities to broaden the findings. Further work could establish the effects of collaborating remotely for tasks where a lack of spatial partitioning is particularly disadvantageous, and investigate the effectiveness of the proposed visual boundaries and coloured palettes. Further work could also explore whether the findings extend to larger groups of mutually-remote collaborators and to group-to-group collaboration.

Secondly, further work could also explore further the proposed awareness interaction techniques (such as dragging from a point to create a menu). The properties of localised and continuous visual changes may also be applied to boost workspace awareness in other kinds of remote collaboration interface, such as conventional GUI groupware. Novel 3-D display technologies may enable 3-D depth cues in remote arm representations and hence address the problem of leaning to gesture, in remote tabletops and other large-format interfaces.

Thirdly, future work may begin to consider how to combine the shared workspace of remote tabletops with a head-and-shoulders video view. Such an addition may offer benefits in tasks, such as negotiation tasks, that rely on social cues. However, the person-space technology must faithfully convey spatial cues such as eye gaze, body orientation, and arm location, to avoid impairing the accurate spatial cues conveyed by the remote tabletop task-space technology. Some systems, for instance, use an overlaid remote tabletop arrangement yet present a video view of the remote collaborator on a vertical display on the far side of the table [22]. Such systems may present conflicting spatial cues because a user sees the remote arm reaching away from them on the table but towards them in the video view.

Finally, perhaps most importantly, as the cost of large displays decreases, and remote tabletop interfaces become easier to construct over time, there may be an opportunity to move away from laboratory-based work and towards field studies of long-term installations and real-world tasks.

Conclusion

This chapter began by reviewing work in remote tabletop interfaces, an emerging class of interfaces for remote collaboration in a shared workspace. We showed that there has been little empirical work in the area, and in particular that the work practices afforded by remote tabletop interfaces are not well understood. As a starting point to explore this area, we presented a study of two work practices in remote tabletop collaboration. The study compared co-located and remote tabletop collaboration using two remote seating arrangements. Neither of the remote seating arrangements led to territorial partitioning of space in the way observed in the co-located arrangement. We propose that this difference was caused by differences in reaching. All three conditions afforded individual and group work as part of a shared task. However, two interaction techniques impaired workspace awareness in the remote cases, and so impaired the ability to transition fluidly between working styles. We propose that these problems were caused by a lack of localised interaction, and a lack of continuous visual changes. The results yield implications for the design and further study of these interfaces.

Acknowledgments We gratefully acknowledge Mark Ashdown, Alan Blackwell, Darren Edge, Cecily Morrison, and the anonymous referees, for feedback on drafts; and the EPSRC and Thales Research and Technology (UK), who funded this work.

References

- [1] Egido C (1988) Video conferencing as a technology to support group work: a review of its failures. In: CSCW'08: Proceedings of the 1988 ACM conference on Computer-supported cooperative work, ACM Press, pp 13–24
- [2] Okada KI, Maeda F, Ichikawaa Y, Matsushita Y (1994) Multiparty videoconferencing at virtual social distance: Majic design. In: CSCW '94: Proceedings of the 1994 ACM conference on Computer supported cooperative work, ACM Press, pp 385–393
- [3] Bly SA, Harrison SR, Irwin S (1993) Media spaces: bringing people together in a video, audio, and computing environment. *Commun ACM* 36(1):28–46
- [4] Nguyen D, Canny J (2005) Multiview: spatially faithful group video conferencing. In: CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press, pp 799–808
- [5] Nguyen DT, Canny J (2007) Multiview: improving trust in group video conferencing through spatial faithfulness. In: CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press, pp 1465–1474
- [6] Tang JC, Minneman S (1991) Videowhiteboard: video shadows to support remote collaboration. In: CHI '91: Proceedings of the SIGCHI conference on Human factors in computing systems, ACM Press, pp 315–322
- [7] Gutwin C, Greenberg S (2002) A descriptive framework of workspace awareness for real-time groupware. *Comput Supported Coop Work* 11(3):411–446

- [8] Bly SA (1988) A use of drawing surfaces in different collaborative settings. In: CSCW '88: Proceedings of the 1988 ACM conference on Computer-supported cooperative work, ACM Press, pp 250–256
- [9] Weiser M (1999) The computer for the 21st century. SIGMOBILE Mob Comput Commun Rev 3(3):3–11, doi: <http://doi.acm.org/10.1145/329124.329126>
- [10] Scott SD (2005) Territoriality in collaborative tabletop workspaces. PhD thesis, University of Calgary, Calgary, Alberta, Canada
- [11] Morris MR (2006) Supporting effective interaction with tabletop groupware. PhD thesis, Stanford University
- [12] Rogers Y, Lindley S (2004) Collaborating around vertical and horizontal large interactive displays: which way is best? Interacting with Computers 16(6):1133–1152
- [13] Scott SD, Carpendale MST, Habelski S (2005) Storage bins: Mobile storage for collaborative tabletop displays. IEEE Computer Graphics and Applications 25(4):58–65
- [14] Tang A, Tory M, Po B, Neumann P, Carpendale S (2006) Collaborative coupling over tabletop displays. In: CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems, ACM Press, pp 1181–1190
- [15] Wellner P (1993) Interacting with paper on the DigitalDesk. Commun ACM 36(7):87–96
- [16] Ashdown M, Robinson P (2005) Escritoire: A personal projected display. IEEE MultiMedia 12(1):34–42
- [17] Esenther A, Ryall K (2006) RemoteDT: Support for multi-site table collaboration. In: Proc. CollabTech'06: International Conference on Collaboration Technologies (Collab-Tech)
- [18] Hutterer P, Close BS, Thomas BH (2006) Supporting mixed presence groupware in tabletop applications. In: TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems, IEEE Computer Society, pp 63–70
- [19] Regenbrecht H, Haller M, Hauber J, Billingham M (2006) Carpeno: interfacing remote collaborative virtual environments with table-top interaction. Virtual Reality 10(2):95–107
- [20] Tang A, Neustaedter C, Greenberg S (2006) VideoArms: Embodiments for mixed presence groupware. In: Proc. HCI 2006: Proceedings of the 20th British HCI Group Annual Conference, pp 85–102
- [21] Coldefy F, dit Picard SL (2007) Digitable: an interactive multiuser table for collocated and remote collaboration enabling remote gesture visualization. In: Proc. PROCAMS'07: IEEE Workshop on Projector-Camera Systems, pp 1–8
- [22] Izadi S, Agarwal A, Criminisi A, Winn J, Blake A, Fitzgibbon A (2007) C-Slate: A multi-touch and object recognition system for remote collaboration using horizontal surfaces. In: Proc. TABLETOP'07: Second Annual IEEE International Workshop on Horizontal Interactive Human-Computer Systems, pp 3–10
- [23] Gutwin C, Greenberg S (1999) The effects of workspace awareness support on the usability of real-time distributed groupware. ACM Trans Comput-Hum Interact 6(3):243–281
- [24] Ashdown M (2004) Personal projected displays. PhD thesis, University of Cambridge Computer Laboratory
- [25] Tuddenham P, Robinson P (2007) Distributed tabletops: Supporting remote and mixed-presence tabletop collaboration. In: Proc. TABLETOP'07: Second Annual IEEE International Workshop on Horizontal Interactive Human-Computer Systems, pp 19–26
- [26] Tang JC (1991) Findings from observational studies of collaborative work. Int J Man-Mach Stud 34(2):143–160
- [27] Olson GM, Olson JS, Carter MR, Storrøsten M (1992) Small group design meetings: an analysis of collaboration. Human-Computer Interaction 7
- [28] Bekker MM, Olson JS, Olson GM (1995) Analysis of gestures in face-to-face design teams provides guidance for how to use groupware in design. In: DIS '95: Proceedings of the 1st conference on Designing interactive systems, ACM Press, pp 157–166

- [29] Bly SA, Minneman SL (1990) Commune: a shared drawing surface. *SIGOIS Bull* 11(2-3):184–192
- [30] Tang JC, Minneman SL (1991) Videodraw: a video interface for collaborative drawing. *ACM Trans Inf Syst* 9(2):170–184
- [31] Ishii H, Kobayashi M (1992) Clearboard: a seamless medium for shared drawing and conversation with eye contact. In: *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM Press, pp 525–532
- [32] Wellner P, Freeman S (1993) *The DoubleDigitalDesk: Shared editing of paper documents*. Tech. Rep. EPC-93-108, Xerox Research Centre, Cambridge
- [33] Ishii H (1990) Teamworkstation: towards a seamless shared workspace. In: *CSCW '90: Proceedings of the 1990 ACM conference on Computer-supported cooperative work*, ACM Press, pp 13–26
- [34] Takao N, Shi J, Baker S (2003) Tele-graffiti: A camera-projector based remote sketching system with hand-based user interface and automatic session summarization. *Int J Comput Vision* 53(2):115–133
- [35] Robinson JA, Robertson C (2001) The LivePaper system: Augmenting paper on an enhanced tabletop. *Computers & Graphics* 25(5):731–743
- [36] Kuzuoka H, Yamashita J, Yamazaki K, Yamazaki A (1999) Agora: a remote collaboration system that enables mutual monitoring. In: *CHI '99: CHI '99 extended abstracts on Human factors in computing systems*, ACM Press, pp 190–191
- [37] Wilson A, Robbins DC (2006) PlayTogether: Playing games across multiple interactive tabletops. In: *IUI'06 Workshop on Tangible Play: Research and Design for Tangible and Tabletop Games*
- [38] Kirk D, Crabtree A, Rodden T (2005) Ways of the hands. In: *ECSCW'05: Proceedings of the ninth conference on European Conference on Computer Supported Cooperative Work*, Springer-Verlag New York, Inc., pp 1–21
- [39] Kirk D, Rodden T, Fraser DS (2007) Turn it this way: grounding collaborative action with remote gestures. In: *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM Press, pp 1039–1048
- [40] Kirk D, Fraser DS (2006) Comparing remote gesture technologies for supporting collaborative physical tasks. In: *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, ACM Press, pp 1191–1200
- [41] Hauber J, Regenbrecht H, Billingham M, Cockburn A (2006) Spatiality in videoconferencing: trade-offs between efficiency and social presence. In: *CSCW '06: Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, ACM Press, pp 413–422
- [42] Pauchet A, Coldefy F, Lefebvre L, Picard SLD, Perron L, Bouguet A, Collobert M, Guerin J, Corvaisier D (2007) Tabletops: Worthwhile experiences of collocated and remote collaboration. In: *Proc. TABLETOP'07: Second Annual IEEE International Workshop on Horizontal Interactive Human-Computer Systems, 2007.*, pp 27–34
- [43] Dourish P, Bellotti V (1992) Awareness and coordination in shared workspaces. In: *CSCW '92: Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, ACM Press, pp 107–114
- [44] Gutwin C, Greenberg S (1998) Design for individuals, design for groups: tradeoffs between power and workspace awareness. In: *CSCW '98: Proceedings of the 1998 ACM conference on Computer supported cooperative work*, ACM Press, pp 207–216
- [45] Scott SD, Carpendale MST, Inkpen KM (2004) Territoriality in collaborative tabletop workspaces. In: *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work*, ACM Press, pp 294–303
- [46] Hall ET (1966) *Distances in Man: The Hidden Dimension*. DoubleDay, Garden City, NY
- [47] Tuddenham P, Robinson P (2009) Territorial coordination and workspace awareness in remote tabletop collaboration. In: *CHI'09: Proc. 27th ACM SIGCHI Conference on Human Factors in Computing Systems*, pp 2139–2148

- [48] Kruger R, Carpendale S, Scott SD, Tang A (2005) Fluid integration of rotation and translation. In: CHI'05: Proceedings of the 2005 SIGCHI conference on Human Factors in computing systems, pp 601–610
- [49] Nacenta MA, Pinelle D, Stuckel D, Gutwin C (2007) The effects of interaction technique on coordination in tabletop groupware. In: Proc. Graphics Interface 2007, pp 191–198
- [50] Ha V, Inkpen KM, Mandryk RL, Whalen T (2006) Direct intentions: The effects of input devices on collaboration around a tabletop display. In: Proc. TABLETOP 2006, pp 177–184
- [51] Forlines C, Shen C (2005) DTLens: multi-user tabletop spatial data exploration. In: UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology, ACM Press, pp 119–122
- [52] Bier EA, Stone MC, Pier K, Buxton W, DeRose TD (1993) Toolglass and magic lenses: the see-through interface. In: SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques, ACM Press, pp 73–80
- [53] Morris MR, Paepcke A, Winograd T, Stamberger J (2006) Teamtag: exploring centralized versus replicated controls for co-located tabletop groupware. In: CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems, ACM Press, pp 1273–1282