

VisTACO: Visualizing Tabletop Collaboration

Anthony Tang¹, Michel Pahud², Sheelagh Carpendale³, Bill Buxton²

¹GVU Center
Georgia Institute of Technology
Atlanta, GA
tonyt@cc.gatech.edu

²Microsoft Research
One Microsoft Way
Redmond, WA
{mpahud | bibuxton}@microsoft.com

³Dept. of Computer Science
University of Calgary
Calgary, AB, Canada
sheelagh@ucalgary.ca

ABSTRACT

As we design tabletop technologies, it is important to also understand how they are being used. Many prior researchers have developed visualizations of interaction data from their studies to illustrate ideas and concepts. In this work, we develop an interactional model of tabletop collaboration, which informs the design of VisTACO, an interactive visualization tool for tabletop collaboration. Using VisTACO, we can explore the interactions of collaborators with the tabletop to identify patterns or unusual spatial behaviours, supporting the analysis process. VisTACO helps bridge the gap between observing the use of a tabletop system, and understanding users' interactions with the system.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design, Human Factors

Keywords: Tabletop, collaboration, information visualization

INTRODUCTION

Effective design of new technologies needs to be paired with equally effective techniques for evaluating those technologies. Collaborative tabletop technologies are no different: the recent explosion in popularity has triggered the need for both user requirements for design, and metrics for evaluating the utility of different interventions. Designing for tabletop collaboration builds on many of the mechanics of shared systems, leading designers to revisit issues of simultaneous input [1], synchrony [16], and awareness [4] when designing tabletop systems, and to rethink about how they should be evaluated. The tabletop context, however, presents many new fundamental challenges that influence group work such as orientation [12,31] and spatiality [18,26]. In light of this, early work on collaboration in shared physical workspaces (e.g. [26]) has proved to be invaluable as a way of helping us to understand how tables are used in collaborative activity: in particular, how users exploit the spatial affordance provided by tabletops to manage and coordinate activity.

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

ITS'10, November 7-10 2010, Saarbrücken, Germany

Copyright © 2010 ACM 978-1-4503-0399-6/10/11...\$10.00.

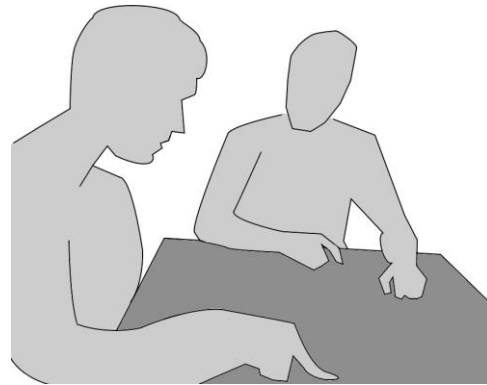


Figure 1. A still frame from a video capture of a study of tabletop collaboration. The challenge of performing video analysis is choosing what to focus on in order to understand what happened.

This paper develops a model of tabletop collaboration that focuses on users' individual interactions around tabletop groupware. Different designs of tabletop technology can have a strong effect (negative or positive) on users' ability to employ spatiality in their collaboration. The model we develop here can be used as a way of analyzing this spatial behaviour—in particular, drawing attention to patterns as well as unusual events—providing valuable insight into how a particular design affects the spatial affordance. Our model focuses on three axes of analysis: spatial (global vs. local), temporal (aggregate vs. particular), and subject (person vs. object vs. table).

Based on these axes, we designed VisTACO, an interactive visualization tool that allows researchers to study collaborative tabletop interaction data. VisTACO helps overcome the analytic gap between observing a study of tabletop collaboration and understanding what actually happened. A conventional analytic approach is to perform a time-consuming video analysis by watching and re-watching video (e.g. Figure 1) to understand users' interactions. This is extremely demanding—particularly for collecting aggregate data, or for exploratory analysis, where hypotheses are not yet fully formed. VisTACO facilitates visual exploration of users' interaction data, and rapid playback of a session. The visualizations emphasize spatial behaviour embedded in the interaction data, providing a deeper understanding of the dynamics of individuals' actions in a collaborative session.

This work makes two contributions to the research community. First, we develop a model of collaborative tabletop activity that focuses on spatial behaviours. Second, we present an interactive visualization tool that focuses on the principles highlighted on the model. Together, the model and the visualization provide an effective starting point for understanding how their tabletop designs affect collaborators' spatial interactions.

BACKGROUND

The research community has taken a many approaches to evaluate tabletop systems. This reflects the pragmatic needs of researchers in the nascent field: when evaluating tabletop interaction techniques, the focus tends to be on traditional performance metrics such as time and errors (e.g. [6]); other studies of highly novel tabletop technologies focus on experiential qualities (e.g. [10,19,20,21,23]); still other studies of tabletop technologies employ multiple methods, such as quantitative metrics (completion time), qualitative observation, and experiential metrics (e.g. [3,14,15,31]). Generally, these studies focus on the particular tabletop system design in question—how well has the system performed, how do users enjoy using the system, how can the system be improved, and so on.

In a few cases, researchers have worked from these observations to develop models of interaction at a higher level of abstraction (e.g. [9,12,18,25,26]). The utility of these models is that they provide a framework to understand collaborative activity in a way that extends beyond the specific systems being examined. Thus, they are useful for deriving design requirements for collaborative tabletop systems because they identify how affordances are used to facilitate or coordinate group work. Usually, these models are developed from iterative video analysis of collaborative activity, where categories of behaviours and function are identified and refined during the analysis. The purpose of performing this analysis is two-fold: first, to gain a deeper understanding of the activity beyond face-value; second, and perhaps more importantly, to understand the collaborative function of seemingly innocuous micro-actions.

The challenge with these models is that without undergoing a similarly painstaking analytic process, it is difficult for designers to assess whether their tabletop system is supporting these spatial processes appropriately. By merely observing users' behaviours, coarse-grained, provisional theories can be made of the tabletop collaboration; however, these observations are subject to bias. Our goal in developing VisTACO here was to address this concern: by visualizing concrete data, we would be able to explore and substantiate these provisional theories with meaningful visualizations.

Spatiality in Tabletop Collaboration

Many authors have argued that the spatial affordance provided by tabletops is both a fundamental collaborative and design resource (e.g. [18,19,26]). As articulated by Tang [26], the space on tabletops is used to mediate access to areas and therefore resources on the tabletop. Scott et al.

[18] further explore how the organic spatial partitioning that occurs during collaboration is a means for users to organize and coordinate activity, identifying three types of territories: personal (intended primarily from independent work), group (intended for shared resources), and storage. These territories emerge in traditional tabletop activity based on the locations of objects that are in the space itself. They reflect users' ability to reach into the space, but also what they are likely to interact with on the table—generally, the closer the object is to a collaborator, the more likely it is for him/her to interact with it. As an alternative interpretation of this formulation of territories [20], many tabletop designs have made territories explicit: UbiTable makes these zones visually explicit, and properties such as access control are bound to an object's location.

Kruger et al. [12] provides a rich descriptive account of the multiple roles (coordination, communication, and comprehension) that orientation of tabletop objects play in collaboration. For instance, objects that are oriented toward a particular collaborator are intended primarily for his/her access (comprehension). When one of those objects is reoriented toward another collaborator, this is a way to support communication, or to transfer/imply ownership/control of the object (coordination). Thus, we see that both the location and orientation of tabletop objects play important coordinating functions on the tabletop.

Broadly, this literature shows us that users meaningfully employ space on tabletops; yet, to what extent is this behaviour merely a function of users' bodies getting in the way? In a collocated context, it stands to reason that if one's arms are occupying the space at one side of the table, one's collaborators' access to the same space would be impeded—regardless of reach. Studies of distributed tabletops (e.g. [29,30,24]) provide an avenue to address this question: when collaborators' corporeal embodiments no longer impede one's access to parts of the space, do users take advantage of that space, or do users still employ tabletop spatiality to coordinate activities? The literature provides a somewhat mixed response to this question, suggesting that the use of tabletop spatiality (in a distributed context) functions either differently, or differentially based on the particular tasks at play. We return to this research question later when we visualize the data from a study of distributed tabletops to understand how users make use of space without the physical impediment of collaborators' bodies.

Visualizing Interaction

Many researchers have employed visualizations to understand how users make use of input devices (e.g. [11]). Figure 2(a) shows the 3D trace of the location of an input device. In accordance with Fitts' Law, it illustrates the fast, ballistic portion (top) and the slower, more deliberate corrective portion (bottom) of a targeting action. The example is instructive: targeting actions are fleeting—in *situ*, this behaviour would have been hard to see, whereas visualizations provide us with a mechanism to identify patterns of interaction. This problem is even more pronounced in stu-

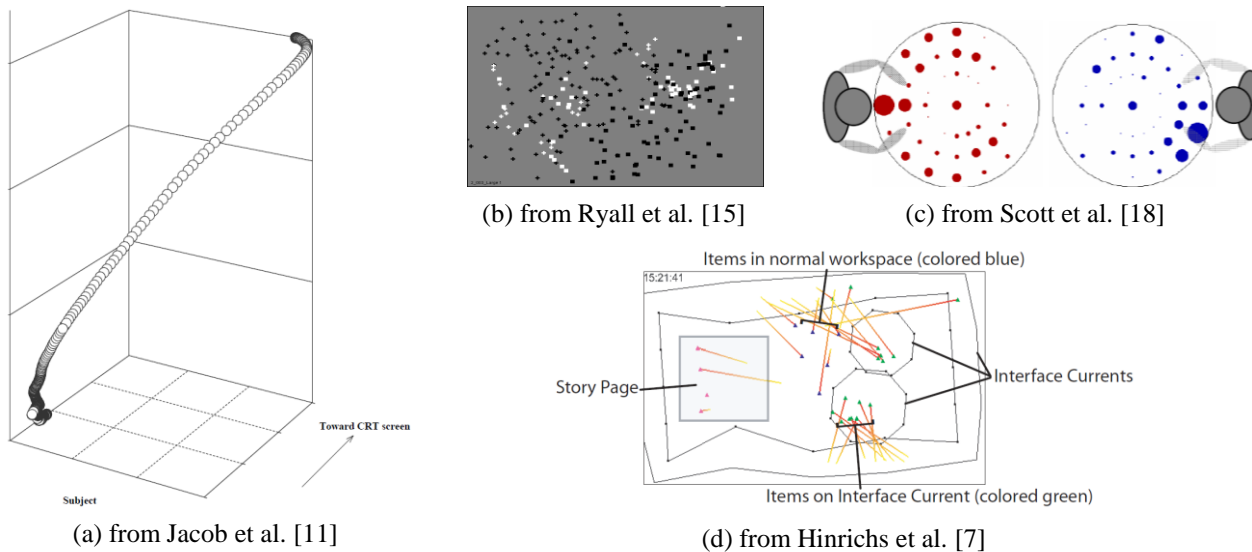


Figure 2. Four examples of interaction visualization: (a) shows two semantic phases (ballistic and correction) in targeting; (b) distinguishes between contact points of different collaborators on a tabletop; (c) allows us to see the relative locations (and amounts) of contact points of collaborators; (d) shows traces of collaborators’ contacts across different interface elements.

dies of collaboration, where multiple actors are interacting simultaneously (e.g. Figure 1): beyond choosing which actor to focus on, there are issues of choosing what to look for, interactions *between* the actors, and so on. The use of visualizations helps address this problem by factoring out the fleetingness of interaction, facilitating review and comparison—in essence, presenting dynamic data statically.

Several researchers have employed visualizations in the tabletop context (e.g. Figure 2(b-d)). Scott et al. [18] construct an activity plot of users’ traditional tabletop interactions based on a video analysis, illustrating the contact points of interaction in successive partial sectors. Ryall et al. [15] and Tuddenham et al. [30] generate activity plots based on users’ contact points on digital tabletops to similarly assess the spatiality of interaction. It is unclear whether these plots visualized the initial point of contact or the final “release” point of contact—both of which may be of interest depending on the type of analysis one is conducting.

Hinrichs et al. [7] provide one of the first visualizations of tabletop interaction where the direction of interaction is illustrated. For their purposes, they were interested in where tabletop artefacts were moved (i.e. to illustrate that they were moved into their interface currents). The visualization’s strength is in its ability to illustrate the *movement* of user’s interactions through the space—rather than illustrating just the beginning or end of the interaction, the streaks show how the contact point moved.

Although we have seen many examples of visualizations of tabletop collaboration, they have typically been for the purpose of illustrating some specific idea that is evident in the interaction data. In many cases, it has been to illustrate the spatiality of interaction [15,18,24,30], while in other cases it has been use to illustrate the flow of interaction [7]. If

we are to use these visualizations as part of the analysis process (as in [18]), we need a more generalized model of tabletop interaction that can be used as a basis for visualization construction and comparison. We begin the process of developing that model in the next section.

AN INTERACTIONAL MODEL OF TABLETOP COLLABORATION

As discussed earlier, tabletop collaboration unfolds extremely quickly: multiple users are interacting with one another and on the workspace, and there are multiple possible areas of interest on the work surface itself—all of which needs to be understood both piecemeal and as a whole. In this section, we distill aspects of tabletop collaboration that are generally captured for the purpose of analysis. We consider these as core components of a model of tabletop collaboration (summarized in Table 1).

Temporal: Aggregated vs. Specific

Researchers are interested here in both the *aggregate* set of interactions as well as *specific* interactions. The aggregate set is useful for tracking and understanding *patterns* of behaviour over the course of the entire interaction/session. For instance, in describing their theory of territoriality, Scott et al. [18] discuss the entire set of interactions by their users, describing a pattern of interactions that are consistent with a small set of territories.

In contrast, we have also seen focus on temporally local interactions—where the researcher focuses on a highly focused incident spanning the course of perhaps a few seconds—which are useful for illustrative purposes, as well as “interest cases” that sometimes deserve further analysis. Kruger et al. [12] provide specific vignettes that are used to illustrate specific concepts within their orientation framework. Similarly, Hinrichs et al. [7] provide clarity into how

Axis	Type	Example Research Question
Spatial	Global	How are users’ interactions distributed over the tabletop?
	Local	Who is responsible for interactions in this area of the tabletop?
Temporal	Aggregate	What is the overall distribution of interactions over the course of the entire task?
	Specific	How does the distribution of interactions change when we look at phases of the task?
Subject	Group	How does the group interact with the work surface?
	Individual	What does the individual with a sub-optimal orientation interact with?
	Object	Who makes contact with this specific tabletop media item?

Table 1. Summary of the interactional model of tabletop collaboration.

specific tabletop artefacts were moved between regions of space by focusing on specific temporal segments.

The requirement for an analysis tool from a temporal perspective is to facilitate the filtering of activities based on time. While this may seem fairly straightforward, the ability to focus on specific segments of time allows researchers to drill down and focus on particular incidents of interest—either for illustrative or analytic purposes.

Spatial: Global vs. Local

The spatial analogue of the temporal component has been equally important. Several researchers have asked questions of the *global* workspace (e.g. [15,24,30]): generally, where are users interacting with the workspace? The activity plots that have been generated to study this spatial behaviour typically focus on the entire workspace, thereby allowing a reader to examine the partitioning behaviour of users.

Similarly, there has been an interest in specific areas of the workspace. For Scott et al. [18], the focus was to examine how interactions were partitioned into semantic regions, where the interest was (for example) in front of each collaborator. In other examples in the literature, designers provide functionality based on spatial semantic “zones.” In UbiTable, the area in front of each collaborator is “locked” to other collaborators [20]—attempting to move an object in someone else’s personal zone results in a non-action. One might ask: how often does this actually occur, or is social protocol sufficient? Here, we have also seen that researchers are interested in users’ behaviours that take place inside, outside, or between spatial zones.

As with temporality, with spatiality, we are interested in being able to see both the global and local views of interaction, and to understand how the local view of interaction fits within the context of the global view.

Subject: Group vs. Person vs. Object

The third major axis of interest is in the subject of the analytic focus: are we interested in the activity of the *group*, the activity of an *individual*, or some *widget* in the workspace? In many cases, researchers are interested in the holistic activities of the collaborators in the workspace, regardless of the individual. For example, Ryall et al. [15] explore the question of whether a group’s behaviour changes with different sized tables, and show that *how* the

group interacts with pieces and distribute labour changes with group size. Similarly, we might be interested in whether an interface supports sharing. Here, we would be interested in the incidence of handover—or, whether objects have been truly shared (i.e. multiple collaborators make contact with it). We would want to focus on being able to see how many objects have had multiple users manipulate them, and to understand the sequence of actions that would have led to this behaviour.

Just as some researchers have been focused on group dynamics, other researchers are interested in understanding the actions of each collaborator independently. In these cases, the research question typically focuses how different configurations of users around the table affect their spatial use of the tabletop (e.g. [24,29]). In such cases, we would be interested in studying the actions of each user independently, thereby facilitating comparison across users and/or across conditions.

Just as some researchers are interested in specific individuals, others are interested in specific entities or widgets. For instance, some researchers have been interested in users’ interactions with specific widgets in the workspace (e.g. [7,17]): how often were they being used, how were they being used, and how were they manipulated? In these cases, the researchers would have been interested only in the interactions that involved the widget in question rather than the other interactions on the tabletop.

Other Considerations

To this point, we have discussed the major components of a generalized observational model for tabletop collaboration, and were sufficient for motivating the development of VisTACO. Based on our experience in designing VisTACO, we outline additional considerations that would be appropriate for inclusion in future iterations of this model.

Points of interaction vs. paths of interaction. Most activity maps for tabletop collaboration provide pinprick points that illustrate where tabletop interaction occurred (e.g. [15,18,30]). If tabletop interactions were similar to mouse clicks, this would be sufficient. However, most current tabletop interaction techniques rely on a three state model: contact-down, contact-move, and contact-up—more akin to mouse dragging. Thus research questions would more likely be about paths of interaction rather than points of inte-

raction; further, it would be inappropriate (in most cases) to reduce the data (for analysis) of a trace to a single point.

Additionally, the beginning and ends of such interactions should likely be interpreted differently. Depending on the research question, one may be more important than the other. For instance, if we were interested in a photo sorting task where the photos were initially strewn about the table, we would likely see that “begin” points of traces would likely be associated with users’ physical location, whereas the “end” points of traces might suggest a semantic organization. Regardless, both beginning and end should be represented and identified.

Semantic phases of interaction. In many tasks, there may exist several phases of interaction. For instance, in the task described by Scott et al. [18], participants were first given a set of still images from a popular television sitcom. Using these, participants were asked to create a storyboard using a subset of the images. One might imagine that a hypothetical group completing this task might go through several different phases to complete the task: first, simply sorting through the images to discover what images were present; second, generating a set of themes as a part of a creative brainstorming process; finally, sorting the images and selecting appropriate ones for the storyboard. Of particular interest here is that the tabletop behaviour of these participants would likely differ depending on the phase, and this differential behaviour may be important to the researcher (depending on the question of interest). It should therefore be possible to classify (and perhaps partition) these semantic phases of activity.

Semantic classification of interactions. At a finer granulation, it would be useful to support semantic classification and partitioning of interactions themselves. If we consider the conventional photo organization task, the main functional interaction is to drag a photo; however, as we have seen from prior literature, the meaning of this interaction depends on the context. If we simply classify the rotational aspect of this photo dragging interaction, it could be for any of comprehension, coordination or communication [12]. We might imagine some of the dragging actions to be none of these, too—perhaps it is simply to clear some space, or to organize parts of the tabletop workspace. Furthermore, many user interactions may be “canned gestures” (for example, the common “pinch” gesture is used for the purpose of zooming in/out of photos). Should these gestures be represented symbolically, or as raw contact point data?

With conventional desktop UIs, a functional classification is facilitated by the widget that captures the input (e.g. we can differentiate between a button press vs. a drag event on a scroll bar). Many current tabletop toolkits (e.g. DiamondSpin, Surface SDK, SMART SDK) support this widget-level functional classification; however, as illustrated above, we need to go beyond functional classifications, and consider the semantics underlying these interactions. For the purpose of analysis, certain types of interac-

tions may be of greater interest than others, and the ability to classify and filter these interactions would be useful.

Multitouch vs. single touch. A more recent development has been multi-touch interaction techniques, and the question is how to address these in an analytic tool. While multi-touch provides users with a theoretical 10 degrees of freedom (one for each finger), the reality is that most of the fingers will be unable to move independently of the others. In practice, we typically only observe users making use of one, two (and sometimes three or four) contact points in a meaningful way. Should additional contact points be discarded from the analysis (perhaps they were not deliberate?), or should we represent these as well?

Summary

We have outlined a general model of tabletop interaction that provides the basis for a generic tabletop analytic process (summarized in Table 1). In the next section, we describe a visualization tool we built to support analysis of tabletop collaboration based on these principles.

VISTACO: A VISUALIZATION TOOL FOR ANALYZING TABLETOP COLLABORATION

VisTACO (Visualization tool for TAbletop COLlaboration) was designed to help understand the use of tabletop systems. A researcher can make use of this system from a workstation with collected log data from any generic tabletop system to study and understand the interactions that took place on the table. Figure 3 shows the main screen of the system, and in section, this we outline its core features, describing how they address the analytic axes from the previous section: temporal, spatial, and subject partitioning.

Design

As shown in Figure 3, the interface is served by three major functional areas: (a) the main visualization pane, which provides an activity trace of users’ activities; (b) the selection areas for selecting the subject of inquiry (per user, per object, per interaction); (c) the time selection area.

Main Visualization Pane—This pane spatially represents the table surface (it is to scale, and has the correct aspect ratio), and shows an activity trace with each of the contact traces made by the users of the system: each contact, each drag of the finger across the tabletop, and each moved artefact. These traces are rendered as faded trails with a large endpoint. In practice, we found that for different research questions, it was sometimes more important to understand the start point of the interaction rather than the end point and vice versa. As suggested by Holten and van Wijk [8], the direction of the fade and the large endpoint of the trails can be flip-flopped (Figure 3(d)). Each user’s traces are shown in a different colour to facilitate quick identification.

In general we found that drawing all of the activity was overwhelming; we thus provide two mechanisms to view the activity traces:

- *plain*: all traces are displayed as single pixel-width lines
- *highlighting*: all traces are faded except for a selected subset, which are made bolder and highlighted (selection

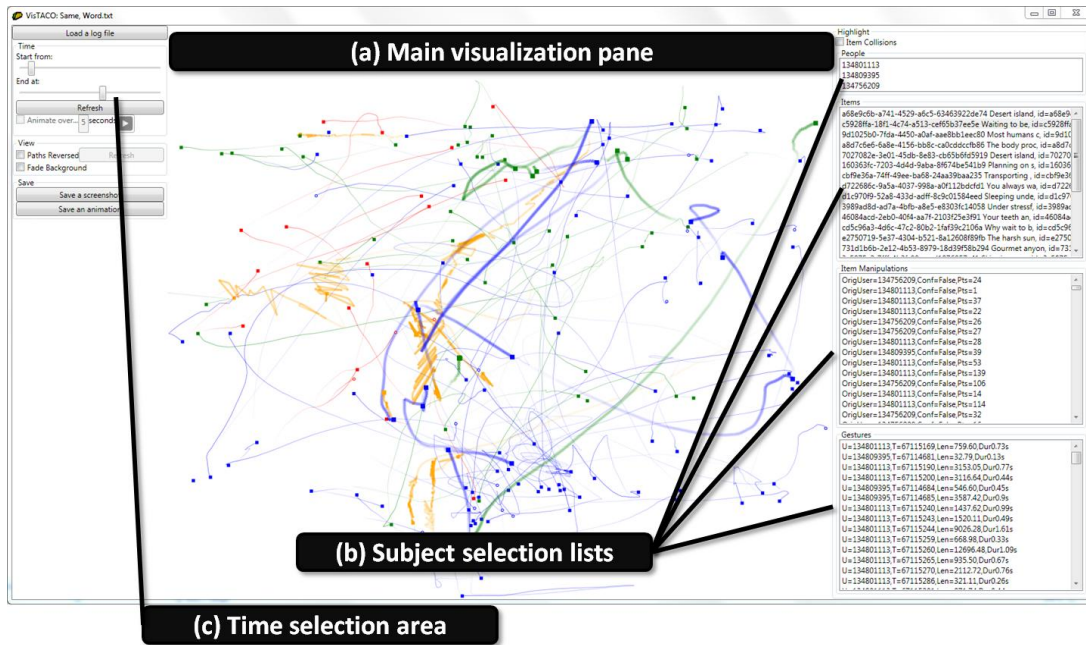


Figure 3. VisTACO user interface has three selection panes: (a) spatial, (b) subject, and (c) time selection.

can occur by hovering/selecting items from the subject selection lists, or by hovering over any of the traces in the main pane)

Subject Selection Lists—VisTACO automatically labels and groups trace data into five separate lists: (a) by user; (b) by artefact ID; (c) by trace ID (every contact); (d) traces with artefacts where there was a conflict (i.e. multiple users trying to move the artefact simultaneously), and (e) gestures, or traces with no associated artefact associated (these would be cases where the user had simply drawn/traced on the surface itself). Hovering over any of the items in the lists automatically highlights the associated trace in the main visualization pane. Explicit selections in these lists culls the remaining lists (e.g. if user 1 never touched item B, then item B is shown in faded form in the item list; deselecting user 1 or selecting additional users brings item B back into the item list). This functionality facilitates the subject-selection component of our model.

The main visualization pane can also be used to “reverse spatial query” the subject selection lists. The analyst can draw an outline on the visualization pane to select a set of traces, which then does two things: (1) it highlights the selected traces in the main visualization pane (based on a convex hull that examines either or both the start and end points of each trace); (2) it makes automatic list selections in the subject selection lists. As illustrated in Figure 6, this functionality facilitates local spatial querying: the analyst can quickly see who interacted with a given region; without this explicit drawn query, the system defaults to a spatially global visualization.

Time Selection Area—By default, the entire activity trace contained in the log file is visualized; however, the analyst can select to view temporal portions of the activity trace

that are relevant. This facilitates the aggregate vs. specific temporal spans that were outline above.

A simple extension window (not shown) provides a visualization of the amount of activity that is captured in the log file. This timeline view shows activity levels for each user during, or for each object. This facilitates drilling down into the data based on specific events that may have been noted, or semantic-level selection of temporal regions (for example, if there are multiple trials captured in the same log file, they can be visualized independently).

Additional Features—When an item has been manipulated multiple times (potentially by multiple users), the trace can become difficult to read/understand. VisTACO provides the ability to playback individual traces, or set of traces. This use of rapid motion gives the analyst a quick sense of the flow and direction of interaction.

CASE STUDY: 3-WAY DISTRIBUTED COLLABORATION

As an analytic tool, VisTACO is built to facilitate exploration of the data set, and then to generate visualizations/statistics that may be relevant for further analysis. It does not replace existing analysis techniques—instead, its role is to augment existing analysis processes, and to inform further analysis. To illustrate how VisTACO does so, we explore a case study to understand the spatial behaviour of users in a distributed tabletop system. The tabletop system in question (illustrated in Figure 4) was developed to support collaboration between users of three different tabletops distributed between multiple sites [24]. The tabletops were connected as a shared visual workspace, and a social proxy was provided for all remote participants for conversation, eye contact, and vice versa. Finally, the researchers employed video cameras mounted above the tables themselves to capture the arms of collaborators. These “video

arm shadows” were then transmitted to remote sites and overlaid atop remote tablespots.

As illustrated in Figure 5(b), a distributed system facilitates a configuration of users that is not possible with traditional tablespots: one where everyone’s orientation to the workspace is the same (e.g. [10,27]). Sharing an orientation to a workspace and document aids deictic reference, legibility [31], and comprehension [12]. It facilitates shared reading and several other types of activities which are more difficult with competing or disjoint orientations.

Yet, how does this configuration impact territorial behaviour that prior authors had seen in collocated scenarios (e.g. [18,26])? What are the factors that contribute to the territorial work practice?

The central question the authors were addressing was the extent to which the co-presence and corporeality of collaborators’ bodies affect territoriality: is territoriality a *consequence* of others’ bodies simply being physically in the way? If one’s collaborators’ bodies are not present (i.e. in the way of one’s work), will one still exhibit territorial behaviour? Further, how is territoriality manifest when all users share the same orientation with workspace? Tuddenham & Robinson [30] found that this work practice manifested differently in distributed tabletop scenarios; however, it left open the question of whether this differential behaviour was a consequence of collaborators’ physical bodies, or the task, and as a consequence, what would happen if users were configured around the workspace differently.

If this novel configuration afforded by distributed tablespots impairs users’ ability to employ the territorial behaviour (which provides a coordinating function [18,26]), then a designer needs to carefully weigh the benefits of each configuration based on the needs of the particular task at hand.

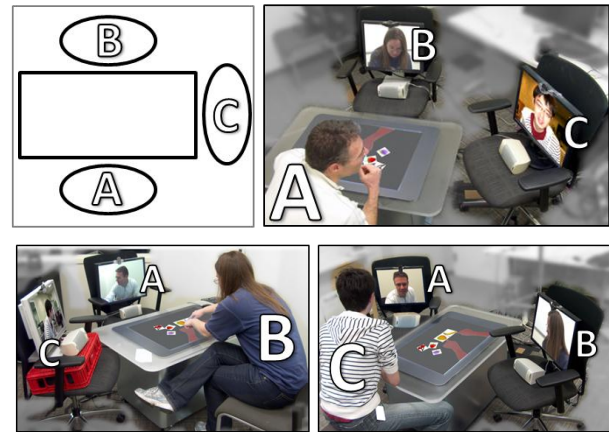


Figure 4. The case study investigated 3-way collaboration using distributed tablespots. Top-left illustrates the conceptual virtual table with three collaborators A, B, and C sitting around it.

The study from which we draw the data presented here is described in detail elsewhere [23]. Very briefly, participants were recruited from the surrounding community (most had no experience with tabletop systems). There were two tasks (photo sort and text sort), and both involved manipulating multiple tiles. Generally, the sessions lasted an hour. Hypotheses were generated from the field notes, and further examined using both video analysis and Vis-TACO to provide supporting or counter evidence.

Analytic Question 1: *What does overall territorial behaviour look like given a “photo sort” type of task?*

Figure 5 shows the two main user configurations: an around-the-table configuration (where users are positioned around the table—one per side), and a same-side configuration (where users are positioned virtually in the same location). Do users’ actions reflect a spatial partitioning? Due

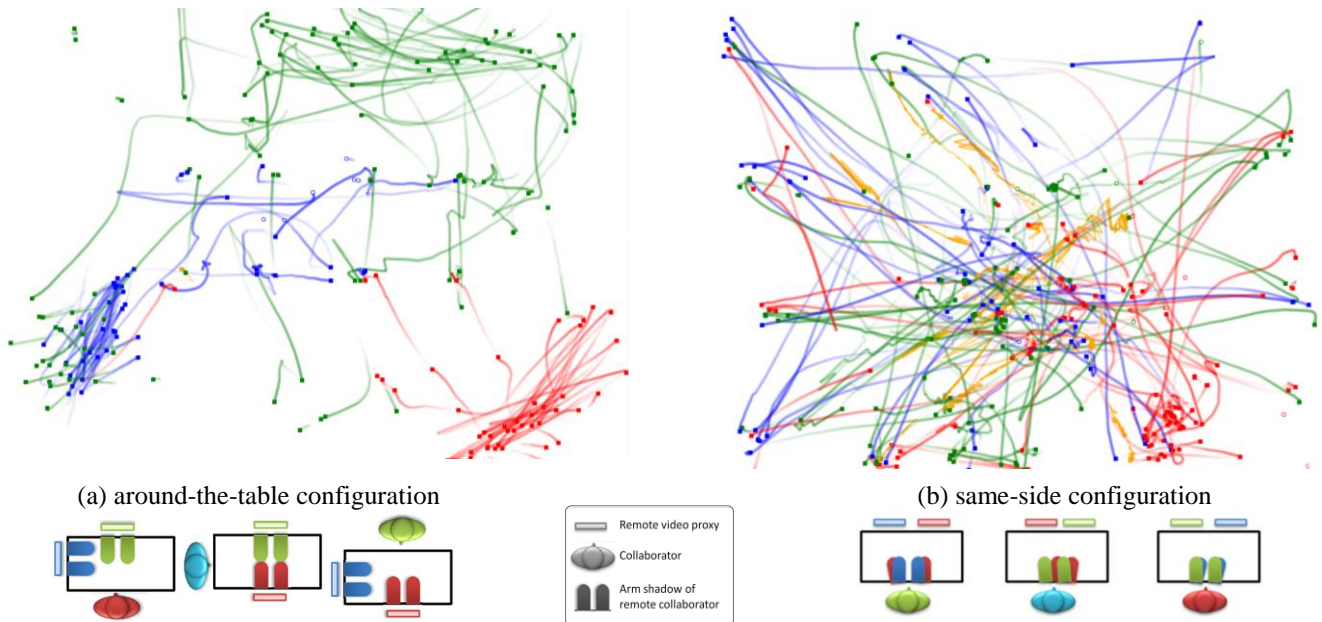


Figure 5. When performing the “photo sort” task, the configuration impacts how a group interacts with the space.

to space constraints, we only consider one group for illustrative purposes—recall that this tool would be used interactively, and so the analyst would iteratively refine his/her hypotheses through exploration.

Here, we are interested in temporally global, spatially aggregated interactions of the entire group. Figure 5(a) shows the distribution of the three users’ activities in an around-the-table configuration. That the users’ interactions are biased to different sides of the table is fairly evident. It also suggests they were right-handed because of the location of the traces (quickly revisiting the video supports this inference). Figure 5(b) shows the distribution of three users’ activities in a same-side configuration. What we see here is that their interactions are far less “side-biased.”

In this task, users were asked to construct figures that were typically quite large (1/2 - 2/3 the entire workspace)—thus, that Figure 5(b) shows users interacting over most of the workspace should not be surprising. Yet, who is responsible for the “pockets” of activity? Figure 6 focuses on the interactions in the corner where the analyst has performed a spatial selection in VisTACO. What we see here is that one user (red) is responsible for most of the interactions here. As it turns out, this area was being used by one participant for “storage” of some of the tiles when they were not being used. Figures 5 and 6 provide a researcher with the clue that “something may be going on the data,” which should prompt additional analysis (perhaps a qualitative analysis) in order to understand why this is occurring.

Analytic Question 2: *How does this behaviour change with a text-based organization task?*

The result from the above (which is suggestive of territoriality in the around-the-table configuration) seems somewhat in opposition to some results of prior work on distributed tabletops [30]. It thus serves to consider the extent to which *task* impacts territorial behaviour (in a spatial sense). The researchers in our case study investigated a second task—a text-based one where the tiles were all oriented in


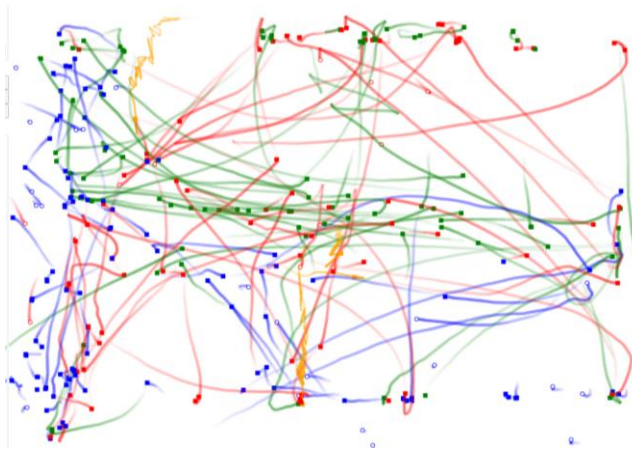



Figure 6. On the bottom-right, we have constructed a spatial query, which shows traces of all items that have been left here. This shows that red is responsible for the storage area.

one direction and could not be rotated. As we see in Figure 7, this task produced dramatically different results: in particular, the difference between configuration conditions is not readily evident.

If we consider the semantics of the text-task, this result is perhaps not overly surprising. Each tile in the text-task contained a string of text—between 17 and 37 words long—that needed to be read. Each of these tiles represented a resource (e.g. bottles of water, knife, etc.) that were collected from an imaginary shipwreck. Participants were asked to imagine themselves as survivors of this shipwreck, and that they could only carry a subset of these items. They were then asked to select a subset of the items, and to generate an agreed-upon ordering for priority of the subset. Figure 7(a) is one group’s trace, and is representative of other groups: participants often created spatial arrangements from left-to-right (this study was done in an English-centric culture) and top-to-bottom in a priority ordering. Items were moved from the columns/rows as deci-



(a) around-the-table configuration



(b) same-side configuration

Figure 7. The trace patterns are not readily differentiable given these two configurations in this task. The twirl traces in (b) represent incidents where one user was trying to get the attention of his collaborators (by moving the tile itself).

sions were made about the utility of each item.

In this type of task, the territoriality described by prior authors manifests differently (i.e. not spatially). The authors from the case study found that roles in the tasks were partitioned rather than the space being partitioned—a somewhat more general formulation of territoriality as described by Scott et al. [18]. Similarly, from the descriptions provided by Tuddenham & Robinson [30], it seems that the participants were employing a semantic territorial partitioning [28] rather than a spatial one.

Analytic Question 3: *How do different phases of task execution affect territorial behaviour?*

As suggested earlier, experimental collaborative tasks often consist of multiple phases of work. In the “photo sort” geometric task used by these researchers, the tiles were initially strewn around the center of the table. Participants were then to use these tiles to create geometric compositions using these atomic tiles—typically needing to reposition, reorient the tiles in a creative fashion. It should be noted that “too many” tiles were available.

Figure 8 shows a temporal partitioning of one group’s trial: the entire trial has been separated into three sections (each lasting ~4 minutes). Figure 8(a) shows the participants clearing the tiles from the center (or sorting). In Figure 8(b), they are actively constructing the composition: blue is doing most of the constructing while green and red look for tiles that are nearby. In Figure 8(c), we see that most of the interactions are small—this suggests they are performing fine-tuning on the compositions. Each of the sub-figures has a distinct signature, which can be gleaned through inspection. These patterns would be more time-consuming to discover through a video analysis.

DISCUSSION & CONCLUSIONS

VisTACO performs best when it is used as an interactive tool for exploring experimental study data. Whereas a researcher’s memory for a particular group/session fades quickly, VisTACO can be used to rapidly replay an entire session and/or to visualize the entire session as a still image. Furthermore, parts of the sessions (temporal, spatial, subject) can be segmented out and analyzed independently, allowing the researcher to explore and understand the data in a way that is not readily available through video analysis or other conventional techniques.

As we demonstrate in this paper, VisTACO is valuable to

support analysis of tabletop coordination—allowing the researcher to quickly ask and answer questions about some aspects of the collaboration (when it is manifest in the interaction data). While it is a useful tool, it should be considered as another tool in a researcher’s toolbox. Generally, we have found it useful as a starting point (to help explore the log data, and to generate questions and theories) rather than as a tool to provide conclusive answers to questions. For example, even when the tool reveals an interesting or unusual pattern of behaviour, it does not necessarily help to answer the question of *why* it happened—such a question may be better answered via detailed video analysis. Thus, a tool like VisTACO can help us to formulate and refine the areas of focus when we engage in deeper analysis of the data using other techniques.

VisTACO relies on interaction data that can distinguish between each collaborator. While it can be used with log data that does not do so, it is considerably less useful. This is a major concern, as FTIR [5] and related approaches (e.g. Microsoft Surface) become increasingly common. While these technologies are extremely efficient and precise in identifying contact points with the surface, they cannot distinguish between or correlate contact points to specific users. In contrast, the venerable DiamondTouch platform [2] and systems exploring distributed tabletops (e.g. [24,29,30]) are well suited for VisTACO. Some technologies can distinguish between touches if the contacts occur simultaneously; however, they cannot distinguish between users—these are less suitable for VisTACO.

In this paper, we presented a model of tabletop collaboration that focuses on observed the spatial activities of users. From this model, we designed a visualization system (VisTACO) that can be used to explore and understand tabletop interaction data. Through a case study, we demonstrated how an interactive system such as VisTACO can be used to help us more deeply understand users’ activities with interactive tabletops. Future work should move beyond VisTACO’s current focus on individuals’ activities, and also examine collaborative acts, such as order and flows of interactions. Such work is foreshadowed by VisTACO’s ability to view all interactions with an artefact, though a more sophisticated set of visualization primitives would likely be necessary to explore these questions. VisTACO, and tools like it will support better and deeper analysis of interactive tabletop designs, thereby facilitating stronger and informed




Figure 8. Splitting up a trial into three shows very different trace signatures. These traces correspond to different semantic phases in the activity: (a) sorting/inspecting the available tiles; (b) construction/searching; and (c) fine-tuning.

designs of this new, promising technology.

ACKNOWLEDGEMENTS

We thank Larry Guarez, whose Wednesday lunch special in the cafeteria was the inspiration for the name of our system. John C. Tang, Gina Venolia, Kori Inkpen, Sara Bly and Hrvoje Benko also provided feedback on the system itself.

REFERENCES

1. Bier, E. A. and Freeman, S. (1991). MMM: a user interface architecture for shared editors on a single screen. In Proc. UIST '91, 79-86.
2. Dietz, P. and Leigh, D. (2001). DiamondTouch: a multi-user touch technology. In Proc. UIST '01, 219-226.
3. Fleck, R., Rogers, Y., Yuill, N., Marshall, P., Carr, A., Rick, J., and Bonnett, V. (2009). Actions speak loudly with words: unpacking collaboration around the table. In Proc. ITS '09, 189-196.
4. Gutwin, C. and Greenberg, S. (2002). A Descriptive Framework of Workspace Awareness for Real-Time Groupware. CSCW 11(3), 411-446.
5. Han, J. Y. (2005). Low-cost multi-touch sensing through frustrated total internal reflection. In Proc. UIST '05, 115-118.
6. Hancock, M., Carpendale, S., and Cockburn, A. (2007). Shallow-depth 3d interaction: design and evaluation of one-, two- and three-touch techniques. In Proc. CHI '07, 1147-1156.
7. Hinrichs, U., Carpendale, S., & Scott, S.D. (2006). Evaluating the Effects of Fluid Interface Components on Tabletop Collaboration. In Proc. AVI '06, 27-34.
8. Holten, D. and van Wijk, J. J. (2009). A user study on visualizing directed edges in graphs. In Proc. CHI '09. 2299-2308.
9. Isenberg, P., Tang, A., and Carpendale, S. (2008). An exploratory study of visual information analysis. In Proc. CHI '08. 1217-1226.
10. Izadi, S., and Agarwal, A. (2007). C-Slate: Exploring Remote Collaboration on Horizontal Multi-touch Surfaces. In Proc. Tabletop '07, 3-10.
11. Jacob, R. J., Sibert, L. E., McFarlane, D. C., and Mullen, M. P. (1994). Integrality and separability of input devices. ACM ToCHI 1 (1), 3-26.
12. Kruger, R., Carpendale, M.S.T., Scott, S.D., Greenberg, S. (2004). Roles of Orientation in Tabletop Collaboration: Comprehension, Coordination and Communication. JCSCW, 13(5-6), 501-537.
13. Luff, P., Heath, C., Kuzuoka, H., Yamazaki, K., and Yamashita, J. (2006). Handling documents and discriminating objects in hybrid spaces. In Proc. CHI '06 561-570.
14. Morris, M. R., Ryall, K., Shen, C., Forlines, C., and Vernier, F. (2004). Beyond "social protocols": multi-user coordination policies for co-located groupware. In Proc. CSCW '04 262-265.
15. Ryall, K., Forlines, C., Shen, C., and Morris, M. R. (2004). Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. In Proc. CSCW '04, 284-293.
16. Salvador, T., Scholtz, J., and Larson, J. (1996). The Denver model for groupware design. SIGCHI Bull. 28, 1 (Jan. 1996), 52-58.
17. Scott, S.D., Carpendale, S., & Habelski, S. (2005). Storage Bins: Mobile Storage for Collaborative Tabletop Displays. IEEE CG&A, 25(4), 58-65.
18. Scott, S.D., Carpendale, M.S.T., & Inkpen, K.M. (2004). Territoriality in Collaborative Tabletop Workspaces. In Proc. CSCW '04, 294-303.
19. Shen, C., Lesh, N., and Vernier, F. (2003). Personal digital historian: story sharing around the table. interactions 10, 2 (Mar. 2003), 15-22.
20. Shen, C., Everitt, K., and Ryall, K. (2003) UbiTable: Impromptu Face-to-Face Collaboration on Horizontal Interactive Surfaces. In Proc. UbiComp '03, 218-288.
21. Ståhl, O., Wallberg, A., Söderberg, J., Humble, J., Fahlén, L. E., Bullock, A., and Lundberg, J. 2002. Information exploration using The Pond. In Proc. CVE '02, 72-79.
22. Streitz, N.A., Tandler, P., Müller-Tomfelde, C., and Konomi, S. Roomware. (2001). Towards the Next Generation of Human-Computer Interaction based on an Integrated Design of Real and Virtual Worlds. In: J. A. Carroll (Ed.): Human-Computer Interaction in the New Millennium, Addison Wesley (2001), 553-578.
23. Tandler, P., Prante, T., Müller-Tomfelde, C., Streitz, N., & Steinmetz, R. (2001). ConnecTables: Dynamic Coupling of Displays for the Flexible Creation of Shared Workspaces. In Proc. UIST '01, 11-20.
24. Tang, A., Pahud, M., Inkpen, K., Benko, H., Tang, J. C., and Buxton, B. (2010). Three's company: understanding communication channels in three-way distributed collaboration. In Proc. CSCW '10, 271-280.
25. Tang, A., Tory, M., Po, B., Neumann, P., and Carpendale, S. (2006). Collaborative coupling over tabletop displays. In Proc. CHI '06 1181-1190.
26. Tang, J. C. 1991. 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. (1991). VideoWhiteboard: video shadows to support remote collaboration. In Proc. CHI '91 315-322.
28. Tse, E., Histon, J., Scott, S.D., & Greenberg, S. (2004). Avoiding Interference: How People Use Spatial Separation and Partitioning in SDG Workspaces. In Proc. CSCW '04, 252-261.
29. Tuddenham, P., and Robinson, P. Distributed Tabletops: Supporting Remote and Mixed-Presence Tabletop Collaboration. (2007) In Proc. Tabletop 2007, 19-26.
30. Tuddenham, P. and Robinson, P. (2009). Territorial coordination and workspace awareness in remote tabletop collaboration. In Proc. CHI '09 2139-2148.
31. Wigdor, D. and Balakrishnan, R. (2005). Empirical investigation into the effect of orientation on text readability in tabletop displays. In Proc. ECSCW '05, 205-224.