Styles of Engagement: How Groups Work over Large Visual Surfaces

ABSTRACT
To build collaborative interfaces for large digital displays such as tabletops and whiteboards, we need to understand how small groups work together over large visual worksurfaces. In this paper, we present the results of two observational studies conducted to understand how individuals transition between personal and group work over large digital worksurfaces. These results suggest that individuals frequently and fluidly engage and disengage themselves from the group task, but that coordination suffers when individuals in a group are not fully engaged. We explore observable causes for this coordination problem, exposing a set of design consequences for digital tabletop surfaces.

Author Keywords
Computer supported cooperative work (CSCW), collaborative tabletop displays, single display groupware, mixed focus collaboration, engagement.

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

INTRODUCTION
Recent interest has shifted from the technical underpinnings for building collaborative tabletops to the design of interfaces for such tabletops (Russ, Katie, Stacey & German guy). This shift has been motivated by the assumption that to build good collaborative interfaces for such displays, we need to first understand the nature of collaborative work over tabletop displays, both in terms of collaborative dynamics (Ryall et al., Tang, Scott et al.) and in terms of the type of work groups will engage in over tabletops (Scott et al. – taxonomy of work tasks).

Many kinds of activities involve *mixed-focus collaboration*, where individuals frequently shift between independent tasks and shared tasks with the group (Carl and Saul). The problem is that we do not have a systematic understanding of the dynamics of mixed-focus collaboration beyond recognizing the end points: individual work and shared work. Yet individuals do not instantaneously shift between independent work and group work; instead, a group's *style of engagement*—the manner in which collaborators are involved and occupied with the work of one another—fluidly changes through the course of work. Tabletop researchers have focused on understanding the dynamics shared tasks (Stacey, Ryall et al, Stacey & Stephan), yet we have only a limited understanding of how individuals conduct individual work in mixed-focus collaboration contexts, and the transitions individuals undertake to engage and disengage with others to perform group tasks. In this paper, our primary focus is to understand these styles of engagement.

While mixed-focus collaboration has been the primary interest of most tabletop research (Stacey, Tang, Ryall), research has focused on contexts where tabletop objects can be moved about with minimal consequence (e.g. photo collections [Stacey & Stephan], word collections [Ryall]). Yet, many visual contexts are spatially embedded with information and therefore cannot be mobile (e.g. parts of a map cannot be moved around without moving them out of spatial context): what collaborative dynamics arise out of these visual contexts? Our second interest in this paper is examining how people work collaboratively over tables with fixed spatial information.

To understand the dynamics of the transitions between individual and group work in fixed visual contexts, we conducted two observational studies of pairs working on mixed-focus tasks on a digital tabletop. Our analysis revealed six styles of engagement that users transition through between highly independent individual work and highly engaged group work. Recognizing these engagement styles demonstrates a need for fluid interface mechanisms to support group engagement.

This paper makes three contributions: first, we identify a working set of engagement styles collaborators use in mixed-focus collaboration; second, from this set of engagement styles, we derive a set of design consequences for digital tabletops supporting mixed-focus collaboration, and finally, we provide a preliminary investigation into the dynamics of collaboration over a spatially fixed data set.

We begin by outlining the collaborative tabletop literature, which acknowledges that individual work occurs, but largely focuses on the *collaborative* aspect of mixed-focus collaboration. We then review more recent work which has recognized the need to provide mechanisms to coordinate independent and group work in collaborative tabletops. This work motivated our own interest in understanding the transitions between independent and group work. We then describe our studies, which were designed to provide insight in this aspect of mixed-focus collaboration. The results of our studies then provide a context for our discussion about engagement in mixed-focus collaboration, where we explore design consequences of our findings.

BACKGROUND: THE NATURE OF TABLETOP WORK

Because traditional tabletops are a flexible and powerful context for collaborative work (Stacey et al.), studying their collaborative use has been of interest to researchers in both the distributed groupware (e.g. Tang, Carl & Saul) and tabletop design communities (e.g. Stacey). Early research was carried out using observational studies of group work over traditional tabletops (e.g. Tang, Bly??). For example, Tang studied groups performing a design task over a traditional table, providing valuable insight into work practices occurring over tabletops, including the role of gestures, the use of space on the tabletop to mediate work, the fluidity of work activity, and the role of collaborators' tabletop orientation in structuring activity. This kind of early work has strongly influenced and framed recent efforts to design interfaces for collaborative tabletops, which has aimed to facilitate the fluid interactions found on traditional tabletops (Stacey et al.).

Spatially Understanding Group Work

Because the most striking feature of large displays is their size, much of the research literature has focused on spatial aspects of collaborative tabletop work practices. Based on the formative work described earlier (Tang), investigations have focused on how to support mixed-focus collaboration from this spatial perspective (Stacey, Ed, Russell). (Stacey) identify three fluid and dynamic areas of shared tabletop displays: personal territories, group territories, and storage territories. These territories provide a spatial context for different types of activities. Personal territories are areas within easy reach, and are used for fine manipulation of and reservation of resources. Group territories provide a context for the "group" task, and generally encompass areas that are not personal territories. Finally, storage territories provide collaborators with areas to temporarily place task resources. This spatial framework provides researchers a means to interpret spatial interactions on tabletop displays, and a broad set of implications for tabletop applications to exploit territorial behaviour (e.g. (Stacey & Stephan)).

Spatial Interference and Engagement

Beyond this spatial framework for group work, researchers have also investigated how low level interactions with a shared interactive surface can cause interference (Morris et al., Morris reference), and how collaborators implicitly avoid interference (Ed). Several tabletop researchers have documented problems with the shared displays, noting common occurrences of spatial interference (Morris references). The frequent incidence of interference indicates that the oft-cited "social protocols" are sometimes inadequate to quell the problem (Morris et al.). Morris et al. articulate a set of coordination strategies to resolve these interference problems: first by identifying three types of conflicts (global, whole element, sub-element), and then by identifying a three classes of strategies (owner-controlled, mixed-initiative, reactive).

The common reports of interference and the development of coordination strategies seem to paint a picture of socially inept collaborators (Morris et al.). If so, what of Tang's reports of "[participants] coordinated actions demonstrating an awareness of the other participants," [Tang] or Scott et al.'s "fluid and dynamic" [Stacey] interactions? A more likely interpretation of these seemingly contradictory findings is that the foundational work for tabletop displays occurred in traditional media environments [tang, scott, blah], yet early attempts to build tabletop displays and interfaces (while successful in pushing forward the state of the research) have failed to support some fundamental aspects of collaborative tabletop work. Not all previous work has shown a noticeable incidence of interference—for example, (Ed et al.) and (Stacey) show that people spatially separate their actions, thereby avoiding interference. The context for these works are different (multiple mice with a standard upright monitor and traditional media, respectively), but it seems unusual that this natural work practice is not evident on tabletop displays (Morris et al., etc.).

Intuitively, it seems plausible that if one is working in an extremely engaged manner with someone else, the goals and intentions of the pair would be well-known to the pair, thereby reducing the incidence of "surprise actions." This reduction of interference would be due in part to mutual workspace awareness (Carl), but furthermore, the pair would have a congruent shared mental model of the next set of required low-level interactions (Pinelle). That is, if a group is to work in a highly engaged fashion, it necessitates greater coordination, thereby reducing the incidence of interference.

Our primary focus then is on this collaborative engagement in mixed-focus collaboration. Prior research has recognized the existence of at least two forms of engagement (independent work and shared work), yet this dichotomy is likely a simplified model of mixed-focused collaboration: intuitively, we believe that groups vary their engagement level throughout a task. If this is true, then collaborative interfaces for tabletops need to provide fluid support for transitions between independent and group work; otherwise, interfaces can simply provide modal access to two different views of the workspace: one that supports independent work, and one that supports highly engaged group work. To understand the dynamics of these transitions, we designed a two complementary studies: in Study 1, participants were assigned a task without specific roles, and although the task itself could be accomplished in a divide and conquer fashion, we did not explicitly make it so; in Study 2, we explicitly assigned roles, designing the task such that part could or should be accomplished independently while another would require more coordination.

Fixed Spatial Data

Beyond studying engagement issues, we are interested in digital tabletops as a means to facilitate collaboration over spatially fixed data (such as maps). In these types of data, the semantics are spatial, meaning that they are implicitly a single unit. Our interest complements existing work that has developed a rich picture of tabletop behaviour for tasks involving fine manipulation of mobile objects, where the focus is on rearranging, reorienting, resizing and sharing of objects (Ringel et al.). Introducing spatially immutable data types is not merely an exercise: data sets that cannot be moved from their default location bring into question the effectiveness of certain tabletop work practices (e.g. Stacey, Russ et al.). How does tabletop territoriality manifest itself when areas of the table *must* be shared? How do individuals establish personal territories? Further, does the incidence of interference increase intolerably? Beyond territoriality, how do groups manage their orientation to the space and each other when there is no "optimal" orientation (e.g. text has an optimal orientation) and the space cannot be reoriented? Some fundamental aspects of our understanding of tabletop behaviour need to be re-examined when simply examining a different data type.

Exploring Fixed Spatial Data

We already have several conventional tools used to explore fixed spatial data sets. Since our studies made use of modified versions of these tools, we briefly describe them here. Since many of these tools can be found in popular mapping applications, such as Microsoft® Streets and Trips, we use a map to describe their functionality.

Supporting overlapping information with filters and layers

Sometimes, maps have overlapping information (such as which streets are freeways, and the path between two locations). With filters, users can selectively display data layers atop the city map using a set of buttons. In Figure 1, both layers are turned on, but layers can be turned on and off independently. Turning filters on and off change the entire display. At times, one may only want to view parts

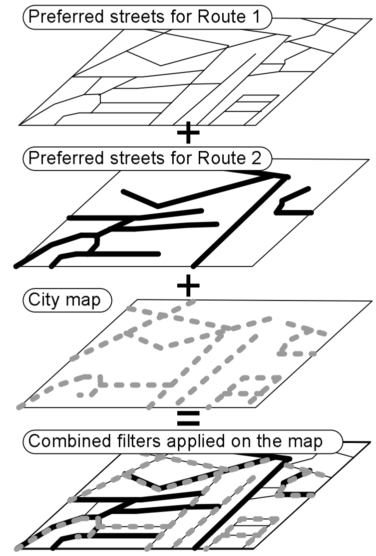

Figure 1. (a) A set of layers being applied to a map, (b) lenses being applied to the same map, and (c) an ShadowBox relocating a region of the map. NOTE: figures b and c not here yet

of a data layer. Magic Lenses (Stone) allow users to create moveable and resizable windows that provide a view of the data layer. When these lenses overlap, they can often show both layers of data (Figure 1). In our studies, filters could be turned on and off using a set of buttons; similarly, lenses could be created by tapping on a set of buttons. Lenses could be moved by dragging their borders, and resized by dragging their corners.

Supporting mobility with ShadowBoxes

DragMag [Ware] allows participants to interactively select an area of the map, and zoom into that area in a moveable window elsewhere on the map. For Study 1, we created the ShadowBox, which has similar characteristics to DragMag, namely that it allows participants to view a given region of the map elsewhere. Beyond this similarity, we "shadowed" participants' interactions in the ShadowBox so that interactions (e.g. drawing, erasing) in the ShadowBox would be immediately reflected in the original map, and vice versa.

OVERVIEW OF OBSERVATIONAL STUDIES

In the next section, we describe two observational studies designed to explore collaborative engagement and the effect of fixed spatial data. To foreshadow, Study 1 was designed to explore the effect of the spatial data type, but participants surprised us by not working independently at all—instead, pairs always worked closely together. In addition to working together, collaborators often moved around the table together, obviating personal territories. Surprised by the lack of individual independent activity, Study 2 gave each participant a specific role, requiring both independent and group activity. We observed a far wider range of behaviour in this second study, providing us with insight into collaborative engagement.

STUDY 1: EXPLORING GROUP WORK

To explore how collaborators work over a spatially fixed data set, we designed a map-based route finding task. Pairs would find two separate bus routes connecting four different end points in a map of a fictitious city. Beyond simply finding roads to connect the end points, participants were given a set of constraints, such as to routes traveled along "preferred" streets. Data to aid participants in meeting these constraints were presented in various data layers (Figure). We were interested in exploring collaborative opportunities afforded by two different ways of presenting this data (Figure): whole screen manipulations (filters) *vs*. small window manipulations (lens). Further, to mobility for our maps, we provided ShadowBoxes (Figure).

Even though this initial foray into fixed spatial data sets was designed as an observational study, given the literature, we expected participants to behave in a particular way.

- *Divide-and-conquer*. We expected that participants would use a divide-and-conquer approach to the task (e.g. divide-and-conquer – Carl, Tony), with each individual working on a separate route.
- *Preference for tools supporting individual work*. We expected that participants would prefer the use of lenses and ShadowBoxes over the filters since they better supported individual work (i.e. an individual could view regions of the data without disturbing the work of the other). In contrast, we expected the global filtering condition to induce interference (since one person's view of the space would affect the entire group's view), thereby inducing a more sequential approach to the overall group strategy.
- *Preference to avoid physical movement.* Since the ShadowBoxes supported mobility of the fixed spatial data set, we expected participants to prefer the use of the ShadowBoxes as opposed to physically moving around the table.

Design

Our study was a 2 (filters vs. lens) × 2 (with ShadowBoxes *vs*. without ShadowBoxes) within-subjects design. The presentation order of the conditions counter-balanced across groups using a balanced Latin square.

Apparatus

We used a large tabletop display (5 × 4 feet) with high resolution (1534 × 1024 pixels) supporting two simultaneous contact points via SMART Technologies' DViT. Participants could interact with the table using pens or fingers, but most only used the pens. The custom tabletop groupware application was built with C#, Direct3D and the Trans2D library, and ran on a dual-Xenon 2.8 GHz Windows XP PC. Dragging the pen on the map could draw routes on the custom-made map while dragging a digital eraser widget would erase the routes. Depending on the condition, the application also provided a set of widgets to control the display of additional data. Finally, we placed two chairs within easy reach of the tabletop display, but no participants chose to use them.

Figure 2 NOTE: this is a figure of a group working on the task. It will show the map and a partially completed route

Participants

Eight paid participants (four pairs: six males, two females) with normal or corrected-to-normal vision were recruited from the general university population. Individuals in each pair knew each other well. Six participants reported being right handed, none had previous experience with tabletop setups, and most had minimal experience with mapping software (two reported having used MapQuest in the past).

Method

Participants first filled out a set of questionnaires to collect demographic information and to assess their experience with mapping applications. Participants were then given a short tutorial on how to use the table display, and general instructions on the task. The route planning task involved finding two routes between a set of four end points while managing a set of constraints (ensuring the bus routes were reasonably direct, that the bus routes traveled along "preferred" streets, that they pass through commercial areas avoiding industrial areas, and to avoid overlapping routes). Prior to each of the four trials (one per condition), participants practiced using the widgets they would be using for that trial. Participants were introduced to the tabletop while standing, and were given the option to use chairs. During the trials, we asked that the participants use the talk aloud protocol, and videotaped their interactions with each other and the tabletop for later analysis. Once the trials were complete, participants participated in a semi-structured design feedback session. This session allowed participants to design better widgets for the purposes of completing the route finding task.

Data Collection and Analysis

Collaborators' interactions on the tabletop were logged and videotaped. Since our primary interest was in the collaborative dynamics of the interactions, we focused our efforts on sections of the videotapes excluding practice times. In total, we analyzed 53, 59, 72, and 57 minutes of video footage from the four groups, respectively. This video was analyzed using a thorough multi-pass, open coding approach (Tang, or some other ref). Field notes from the sessions informed initial coding categories (e.g. which area of the table are they working from), and subsequent coding passes were driven by iteratively refined coding schemes based on further study of the videos. Figure X shows a visualization of a small portion of the analysis log for Study 2. This methodology facilitates an intimate familiarity with the intricate, subtle mechanics occurring in the sessions, providing a very rich understanding of the underlying collaborative processes. Given the observational nature of this study and the low number of groups, statistical hypothesis testing would have been inappropriate.

Results

Our field notes, combined with the results of the video analysis and the feedback from the design sessions brought out two major findings.

Tendency to work together

Contrary to our expectations, pairs worked together across all conditions—visibly working independently less than X% of the time. In only X trials of a total of 16 did pairs even attempt to divide up the task to work alone. Groups generally worked together closely to find one route before finding the other route. We found that groups themselves were highly mobile, with individuals frequently moving around the table to gain a better perspective of the area of interest, or to gain a shared perspective on a particular street. Beyond simply moving together, groups also worked in tandem—often, one person would control the widget (either lenses or filter buttons) while the other would draw the actual route on the display. In some sense, this division of labour could be considered as divide-and-conquer, but our contention is that in these instances, the pairs were working as a single entity on the same problem as opposed to two autonomous entities working independently on different aspects of the same problem.

Group 3 was a notable exception to this result. In the Filter conditions, this pair worked in parallel on different routes independently. To facilitate this parallel operation, the pair used the filter layers in a "time sharing" mode: when one needed to see a given data layer, he would tap and view his layer for as long as he needed while the other worked from memory. Group 3's working style suggests that at least some groups desire to work independently. Group 3 found an awkward way to support their independent working style since our interaction widgets did not provide fluid support for them.

Maintaining context

In addition to working together most of the time, pairs did not use the data widgets in the way we had expected. Most strikingly, the viewing condition that we expected to be least successful (filters) turned out to be the most desired as well as the most efficient. In the other view conditions, users essentially mimicked the functionality of the filters (even when the filters had not yet been presented). For example, with global lenses, users would create table-sized lenses, and move them in and out of the workspace to cover the working area. With the ShadowBoxes, users would simply move them out of the way, preferring to use widgets that affected the global space.

Participants reported that the lens widgets in general suffered from a several usability problems. First, they were somewhat cumbersome—resizing and moving the lenses around was tedious, requiring a mental mode switch from the actual activity (drawing or viewing the workspace) to fine grained manipulation activity. Secondly, participants reported that the lenses did not support the way in which they were working (i.e. as a group rather than independently), and so found them extraneous. Finally, the participants reported that lenses could not be resized to meaningfully partition the space. That is, since the two routes necessarily crossed in the middle of the city, each lens needed to be larger than half of the table anyway. Since the task required optimizing a global route, the participants preferred tools providing global, contextual information rather than simply local data presentation.

Discussion

Pairs were mobile and non-territorial when working together over the spatially fixed data. In general, they worked and moved together as a single entity, and did not exhibit a strong desire to work independently: the *entire workspace* was therefore group territory, with no established personal territories. As a consequence, the tools we had envisioned being used to establish personal territories (lenses and ShadowBoxes) were simply not used.

Our results hinted that participants did not use the lens and ShadowBox tools because they were frequently working in a very engaged manner. For instance, they often worked on the same route together or were looking at the same area of the map together. Beyond this, they preferred the visual mechanisms allowing them to view the space together (filters), thereby providing common ground. Even more compelling, they often stood in close proximity with one another—surprising since the workspace had no orientation cues.

Yet suggesting that participants disliked lenses and the ShadowBoxes because they worked in an engaged manner would be premature. Study 1's task required optimization on what amounted to the entire workspace; as a consequence, an individual would likely prefer filters, which would allow the maintenance of global context anyway.

Study 1 provided the starting point for many questions we were interested in exploring. While we saw engaged, group task behaviour, codifying these forms of engagement would

Figure 3NOTE: This should be a slightly different visualization. Should show: timestamp, position, collab style, that both position and collab style are fluid. Viz should: remove "problem" (say working on same prob), use space for "arrangement". Will also provide richer description of this visualization

have been premature since we had not seen the full range of behaviour from individual work to group work. Was the non-territorial, mobile behaviour a simple consequence of the fact that the workspace could not be meaningfully partitioned? Were our interpretations of group behaviour and context seeking an artifact of the restrictions in the first study's workspace? Finally, in what ways would working behaviours change when groups moved from independent to group work?

STUDY 2: TRANSITIONS IN GROUP WORK

While the first study provided us with insight into the design space of interactive tabletop tools for group work, many questions remained. We designed Study 2 to address the confounds of Study 1. The following summarizes these changes and the rationale for the change.

- *Giving participants different roles*. Participants worked closely together in Study 1, but this may have been because the task description which did not outline independent roles.
- *Explicitly introducing independent and group tasks.* By imposing activity at the extremes of mixed-focus collaboration where engagement would be lowest (during independent activity) and highest (during group activity), we could observe all activity leading up to these low and high points of engagement.
- *Completely conflicting data layers.* The data layers in Study 1 overlapped only in certain regions, meaning that participants in many cases could still work with both filters turned on. Completely conflicting data layers preclude this strategy. This design simulates situations where each individual needs so much information that other information is occluded.
- *Multiple sub-problems.* Study 2 had three related sub-problems that could be meaningfully partitioned (i.e. one person could work on one sub-problem without requiring the entire work surface). By design, two of these sub-problems could be worked on independently, while the third sub-problem overlapped somewhat with the other two. Study 1 had a single problem covering the entire space, precluding independent work altogether.
- *Redesigned lens widget.* We also redesigned the lens widget based on design feedback from Study 1: the lens was a single lens with filter buttons that could apply filters on the lens view independent of the workspace. Furthermore, we removed the ShadowBox condition to focus our efforts.

Pairs worked over a fully connected graph with X nodes and Y edges representing an abstract route planning task (such as airline routes). Two independent data overlays provided edge weight information (one provided "travel time," while the other provided the "financial cost"), where the weights could be 1, 2 or 3. Participants were to generate routes to connect four specific nodes on the graph (1 through 4—spread from the "top" of the table to the "bottom"). Depending on the condition, each participant was responsible for generating one of two *independent routes* (one for travel time, one for financial cost), or the pair was to generate a single *compromise route* (taking into account both travel time and financial cost). We also varied the visual tool participants had to use: global filters, or the redesigned lenses. In following with the protocol of the first study, we allowed teams to stand and walk around the tabletop.

Design

Our study was a 2 (global filters *vs*. lens filters) x 2 (individual routes *vs*. compromise route) within-subjects design. The presentation order of the conditions counter-balanced across groups using a balanced Latin square.

Hypotheses

In this study, the working space had three separate sub-problems (the route between 1 and 2, the route between 2 and 3, the route between 3 and 4). The route between 2 and 3 shared nodes with the other sub problems. Given our experience with Study 1, we had several hypotheses.

- *Individuals will work independently with lenses.* Since lenses support independent work, and independent work can occur with independent sub-problems, we expected independent work to occur frequently. This also applies to the compromise route, where each individual can work on different areas of the route.

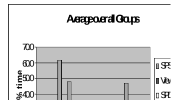

Table 3. NOTENeed some graphs here: engagement by task type 1&3, 2&4. engagement by tool (1&2, 3&4); a second row with arrangement by condition

- *Perspective sharing during group work.* When working together on the same sub-problem, we expected groups to share the same perspective on the problem by standing in close proximity to each another.

Apparatus

We used the same apparatus and setup as Study 1. A new custom-made, fully connected graph was generated. We also redesigned the lens filter widget with feedback from the design sessions of the first study. As in Study 1, the lens could be moved or resized by dragging on the border or the corner of the lens, respectively.

Participants

We recruited eight paid participants (four pairs: four males, four females), different from those in Study 1, with normal or corrected-to-normal vision from the general university population. Six were right handed, and although none had previous experience with tabletop setups, and most had minimal experience with mapping software (two reported having used MapQuest in the past).

Method

Study 2 used an identical protocol as Study 1—the exception being that participants practiced and used the global filters and the redesigned lenses in Study 2. We again videotaped the trials using two cameras, and had the participants use the talk aloud protocol. At the end of the trials, participants again participated in a semi-structured whiteboarding design session to gain insight into their impressions of the tools and the task.

Data Collection and Analysis

A separate video coding analysis was conducted for each study as the focus of each was different. In Study 2, we analyzed a total of 53, 59, 72, and 57 minutes of video footage from the four groups, respectively. In our initial coding passes for Study 2, we focused on individuals' positions around the table and which sub-problem each was working on. In the next section, we detail an open coding scheme specifying different styles of pair engagement that was generated and refined based on these coding passes and field notes. A small portion of the analysis log for Group 2 in Study 2 is visualized in Figure X. This visualization shows the dynamic nature of engagement and its relation with the physical position of collaborators, which we describe further in the next section. As in Study 1, the observational nature of this Study 2 combined the low number of groups rendered statistical hypothesis testing inappropriate.

Results

This study required mixed-focus collaboration, providing a rich range of independent and group work, and allowing us to explore the entire range of engagement behaviour. Our analysis revealed six different levels of group engagement throughout the task. These dynamic engagement styles were related to a range of other factors, including the experimental condition, collaborators' physical positioning around the table, and how interference was handled, providing strong support for our coding scheme. We begin by describing the six styles of group engagement, then describe other factors, and how these related to group engagement.

Styles of engagement

Originally based on field notes, we iteratively refined a coding scheme for the videos in Study 2 to abstractly capture the style of engagement of each pair. We recognized early in Study 1 that pairs fluidly engaged and disengaged with one another with varying styles: in some cases, pairs would be working extremely close together (Figure Xa), and at other times, one would simply be watching the other work (Figure Xb). In Study 2, we aimed to codify these behavioural styles, carefully noting the manner in which collaborators were involved with the work of one another. Figure Y shows a small snapshot of one group's transitions between these styles. Like all groups, this group transitioned in and out of engagement frequently and smoothly.

Our final coding key for collaborative engagement is as follows:

Figure 4. Note: two sub-figures – one with someone grabbing someone, another with someone waving at someone, a third where both are tracing at the same time

***[SPSA]**: (Same Problem, Same Area):* Collaborators are actively working together with the same mindset. They are together evaluating a route, tracing a route together or actively drawing the route together (one person points at landmarks while the other connects them with a pen). Often, this is accompanied with conversation. (Figure X – tracing together)

***[VE]**: (One working, another viewing engaged):* Pair is working together, but one has a more active role in the task, and the other is viewing intently. For instance, one may be showing another a route, or the other may simply be watching very carefully. In this latter case, the individual was watching closely enough to be able to suggest corrections. This state is often accompanied with conversation. (Figure X – one tracing, one watching intently)

***[SPDA]**: (Same Problem, Different Area):* Collaborators are working simultaneously on the same sub-problem, but are clearly working "independently" in the sense that one person's thoughts and actions were unrelated to the others. (Figure X – both working, but in different areas)

***[V]**: (One working, another viewing):* One collaborator is working on a task, and the other is watching the first, but not engaged sufficiently to help or offer suggestions. The latter individual is only able to perceive high level activities, such as that the first has stopped working, or that s/he needs a tool, etc. (Figure X – one person watching from a distance)

***[D]**: (One working, another disengaged):* One collaborator is completely disengaged from the other, not paying attention at all to the work of the partner. (Figure X – one person sleeping)

***[DP]**: (Different problems):* Collaborators are working completely independently, and on separate sub-problems. Largely, one person's interactions with the workspace do not affect the other in any way. (Figure X – opposite sides of the table)

This coding scheme allowed us to keep a running classification of the engagement and collaborative style of the groups throughout the session. After codifying each of the 16 sessions (four groups, four conditions), we ran a set of analyses to understand how engagement styles related to the experimental condition.

When groups worked in the individual routes condition, groups were considerably less engaged with one another compared to when they worked in the compromise condition (Figure X). Whether participants used lenses or global filters also affected the engagement style, with collaborators engaging one another more with global filters (Figure X). By design, the lenses facilitated individual work: an individual could partition the space and work independently. As a consequence, in the lens+individual route condition, participants almost always worked independently (XX% of the time, Figure X). The global filters are somewhat the opposite: with them, only one data layer can be viewed at a time, implying a shared view. When teams were to use global filters, they worked together, both when working to complete the individual routes, as well as when the compromise route was being constructed (XX% of the time, and YY% of the time, NOTE: this is defined as the sum of everything minus (DP + 1W2VD) Figure X). These results were in keeping with our expectations. We were most surprised by the remaining condition (lens+compromise route) where individuals could operate in parallel by working on separate sub-problems with the lenses. Instead, we found that three groups worked together for most of the task, and literally never worked independently of one another (XX% of the time, Figure X). Group 2 was an exception to this result, working in a sequential, independent manner before reengaging and working closely to determine their final solution. Akin to Group 3 in Study 1, this group shows that independent work is possible with global filters, but their behaviour indicates a preference to work simultaneously on this kind of problem.

Positioning

We performed a similar analysis of the arrangement of individuals around the tabletop. Figure X summarizes our orientation-independent classification scheme for the positional relationship of the two collaborators.

In keeping with the results of the first study, we found that collaborator proximity was closely related to how collaborators worked with one another: when collaborators worked together, they physically stood close together. We observed this relationship with respect to our task conditions (Figure X—broken down by condition), but more importantly, with respect to the engagement style (Figure X). With higher engagement states, groups stood physically closer to one another. Although this effect is confounded by the fact that participants were physically closer when working on the same sub-problem, the result corresponds with results from our first study, which had no sub-problems. A notable exception to this observation is that the SideBySide arrangement is closer than the StraightAcross, yet SideBySide was more frequently used for independent work. This result is likely related to the ergonomics of the work surface: the SideBySide arrangement did not facilitate working on the same sub-problem (because one would not be able to see the surface); however, because two sub-problems were spatially independent, SideBySide facilitated concurrent independent work.

Individuals' physical positioning also related to territorial behaviour: individuals tended to explicitly interact only with areas physically close to them (i.e. in their personal territories). Yet, these "territories" were not permanent: as individuals moved around the table, others were no longer restricted from operating in those areas. When a pair became engaged, we observed frequent occurrence of "perspective taking," where a second person might take on the first person's perspective (Figure – the one where the guy is standing right behind the girl). In these instances, the second person would never displace the first: that location and perspective belonged to the first person. Even if the second person wanted to gesture toward the table, he would move to a different location around the table before doing so.

Handling interference

We also viewed many instances of interference (Figure – show one where someone is grabbing someone, show another where someone is waving at someone). The nature of these instances of interference varied: in some cases, visual interference occurred when one collaborator's body interceded another collaborator's view of the workspace (Figure X), in others, physical interference occurred when one collaborator's body physically blocked another from interacting with the workspace.

What varied visibly was the way in which interference was handled with respect to the level of engagement. The more collaborators were engaged with one another, the more smoothly interference was handled. Conversely, the less engaged collaborators were with one another, the more abrupt and visibility awkward interference was handled. For instance, when collaborators were closely engaged, interference was handled extremely gracefully—like a choreographed dance (ref to Tang, Carl), with one collaborator moving out of the way just as another collaborator moved into the space, and back out again in unison. When collaborators were less engaged, we saw instances of gestures, pushing, and even a few instances of grabbing.

Transitions between engagement states

There appeared to be three mechanisms that collaborators used to transition between engagement states: explicit verbal or gestural cuing, consequential communication, and visual interference.

With explicit verbal or gestural cuing, one collaborator would signal his/her partner to evaluate a certain route, to ask for help, or to otherwise engage in closer engagement. For example, we observed the "come here" gesture, or verbal statements such as, "I found a '7' (referring to the cost of the route)," or "Look here." Subsequently, collaborators often transitioned into a higher engagement state. Alternatively, the second collaborator would finish what s/he was doing before moving to see what the first was referring to.

Consequential communication, the information one generates as a consequence of one's interactions with the workspace (Carl), also played a large role in marking the transitions between individual work and group work. Collaborators used this mechanism to maintain an awareness of others in the workspace without explicit cuing. We coded for many of the brief glances that occurred throughout group work, and this provided some insight into the timing of transitions between individual and group work. For instance, we coded many instances of "uncued help," where one collaborator seemed to spontaneously help the other (e.g. to reach or interact with a tool, to interact with the workspace on behalf the other person) without explicit cuing. On subsequent coding passes, we found that these were often preceded by a brief glance. Beyond cuing mundane house-keeping tasks, brief glances also provided insight into timing of higher-level engagement transitions. For instance, some participants were fairly conservative in their gestures over the workspace. When these participants reached over into the workspace, it was a signal that they had found a good working solution, and this action acted as an implicit cue to their partners to re-engage with each other.

Finally, visual interference often led to higher levels of engagement. Because the workspace was shared physically, collaborators who were accustomed to pointing out routes and counting out loud (equivalent to trying out possible chess moves) would often visually and/or physically interfere with his/her partner. In many of these cases, the latter partner would simply "give up" on working independently (for the time being) to see what the first was doing, leading to higher engagement. In other cases, these incidents would lead to the more amusing incidents of interference management (e.g. Figure X).

Discussion

Participants preferred tools that matched their overall engagement strategy and task for each trial. When working on the compromise routes, individuals preferred the filters since it provided a global visual context for work, which accorded with their desire to work together. For working on individual routes, individuals preferred the lens tool since it supported the independent work strategy that they preferred.

Participants also moved to gain similar visual contexts. Beyond simply moving closer together, participants often cocked their heads, visibly tilting their heads to gain similar perspectives on the data. Note that the data set itself had no default or preferred orientation. Instead, it seemed that when pairs engaged one another, one individual would already have had an established view; that individual would persuade his/her teammate toward that perspective so that discussions and deictic references could more easily progress.

The smooth interaction between pairs was particularly striking: when highly engaged, pairs worked in synchrony without interference as a single entity with four arms. The smoothness is likely the cause of collaborators being highly aware of others' actions, activities and goals in the workspace, and moving to accommodate one another (Carl). When pairs were less engaged with one another, and more focused on their individual work, interference was handled awkwardly, likely because pairs were less aware of each others' actions.

Likely, some interference can be interpreted as an individual getting "too much awareness" of a teammate's interactions with the workspace (Carl's usability of awareness and Garth). As described earlier, individuals would often interrupt his task and see what the other person was doing since attempting to continue working independently was futile. These incidents of interference may be mitigated by providing a higher dots-per-inch (dpi) resolution for tabletops, providing individuals with meaningful work areas.

Finally, we observed a wide variance in the overall strategies pairs used to complete the tasks. Some pairs relied on turn-taking strategies when working closely together; others simply worked simultaneously, relying on unspoken coordination to manage the space. This variance is a strong indicator that interface mechanisms to support group must be flexible and fluid to support the many different ways groups work.

We now bring together the results of our studies together, discussing key design implications for tabletop interface designers.

LESSONS FOR PRACTITIONERS

Taken together, our two studies have provided insight into the dynamics of mixed-focus collaboration. The findings suggest that individuals frequently and fluidly move from individual work to many stages of collaborative group work. Each stage is accompanied with slightly different behavioural mechanics: for example, when very tightly engaged, a group prefers to operate in very close proximity to one another—even when the workspace itself has no implicit orientation. Furthermore, we have seen that for the tasks in our study, different viewing tools support different kinds of engagement better: for individual work, independent views are preferred because they reduce interference, and allow each individual to work independently; for tightly engaged groups, groups preferred a tool that provided the same view, thereby providing everyone with a common ground.

1. Support a variety of engagement styles.

Mixed-focus collaboration clearly encompasses a wide variety of engagement styles; consequently, the nature of work and the associated requirements are fluid throughout the course of the collaboration. Most systems fail to provide support for engagement styles, falling back on social protocols to effect engagement styles (e.g. Russ, Ringel 3R's). Since digital tabletops have dynamic displays, we can present different views to support different engagement styles.

2. Provide fluid, flexible mechanisms to transition between engagement styles.

The problem is not trying to support individual work *or* group work in mixed-focus collaboration (Carl & Saul); instead, it is supporting the *transitions* between individual and group work. Providing only a single view of the workspace would limit individuals' abilities to work independently (Carl & Saul), yet providing separate copied workspaces would prevent many group collaborative dynamics (such as being able to see what others are doing) from occurring (Stacey et al.). *We do not recommend against mitigating interference altogether*: it seems clear that some forms of interference are desirable as they help to engage collaborators; instead, it is the unexpected awkward interference that arises when individuals' working desires are in conflict that should be mitigated via fluid interface components that provide all individuals of impending interference. Beyond providing warning of interference, the interface should provide fluid means for transitioning between engagement styles, supporting groups as they move in and out of engagement. For instance, rotating an object toward a fellow collaborator has been shown to temporarily increase engagement (Russ): the system may also use this cue to transition the presentation of the workspace to match the closer engagement.

3. Provide mobile high resolution personal territories.

The interference we observed in our studies was a direct result of individuals' desired working areas overlapping. This exacerbated three problems: the relatively low resolution of the display (relative to say paper), the relatively large body of the collaborator working in the workspace, and the consequence that most interaction occurred in what was effectively a group space (Stacey). Yet, with traditional media, we have seen that interference does not appear to ever be a problem (Stacey, Tang). It may be that we have natural social mechanisms in the real world that facilitate personal space and private work without disturbing others (e.g. working in an area close to oneself), while other mechanisms exist that we can work while attracting the attention of others (e.g. working in the group space). The fixed spatial data set we worked with only allows working in what amounts to the group space.

Most simply stated, we can better support both individual and group work for fixed spatial data with higher dots-per-inch (dpi) resolution displays. Individual work is better accomplished when it is physically close (Stacey), and as a consequence, the larger group is not interfered with. Yet these personal territories also need to be mobile because individuals fluidly move about the workspace as they work over fixed spatial data.

Beyond simply providing higher resolution displays, interface designers need to provide a fluid means to transform a part of the group space into a personal space. In tabletop scenarios where the primary objects are small and mobile, personal spaces are established via orientation and proximity, yet in a spatially fixed data scenario, where the entire space is defined by a single immobile view, establishing personal territory is more difficult. Participants established personal territory with their physical location, but in many situations, physical mobility may not be feasible. The ShadowBoxes from our first study were a step in addressing this problem, but they were inadequate—a likely consequence of the poor interaction dynamics. We have begun exploring more fluid interaction mechanism for the ShadowBoxes, as well as trying to understand real life metaphors that the ShadowBoxes are attempting to mimic.

CONCLUSIONS

Through a series of two studies using digital tabletops, we have explored and exposed a working set of engagement styles collaborators fluidly move through in mixed-focus collaboration. The results of this study also provide some insight into varying observations of interference, where some report that groups work smoothly without interference (Ed, Stacey, Tang), while other researchers have found interference to be such large problem that system mechanisms need to be designed to mitigate interference (Ringel). We believe that a mixture of several factors probably primed some studies for interference more than others (e.g. relative size of embodiment in the workspace, are collaborators able to work independently in a personal space, do individuals need to work over the same working area), but that level of engagement as a factor in interference has not been adequately explored. Namely, if collaborators are highly engaged with one another, interference occurs much less; the less collaborators are engaged, the more frequent and disruptive the occurrence of interference. By first identifying these styles of engagement and understanding their dynamics, we have derived a set of design consequences for digital tabletops supporting mixed-focus collaboration over spatially fixed data sets.

ACKNOWLEDGMENTS

We thank CHI for good times and free food, HCI@UBC for free food and cash, Imager Lab for cheap pop, and Andre Agassi for showing us that glory is forever.

REFERENCES

1. Adobe Acrobat Reader 7. http://www.adobe.com/products/acrobat/.
2. Anderson, R.E. Social impacts of computing: Codes of professional ethics. *Social Science Computing Review 10*, 2 (1992), 453-469.
3. How to Classify Works Using ACM's Computing Classification System. http://www.acm.org/class/how_to_use.html.
4. Klemmer, R.S., Thomsen, M., Phelps-Goodman, E., Lee, R. and Landay, J.A. Where do web sites come from? Capturing and interacting with design history. In *Proc. CHI 2002*, ACM Press (2002), 1-8.
5. Mather, B.D. Making up titles for conference papers. *Ext. Abstracts CHI 2000*, ACM Press (2000), 1-2.
6. Schwartz, M. *Guidelines for Bias-Free Writing*. Indiana University Press, Bloomington, IN, USA, 1995.
7. Zellweger, P.T., Bouvin, N.O., Jehøj, H., and Mackinlay, J.D. Fluid Annotations in an Open World. *Proc. Hypertext 2001*, ACM Press (2001), 9-18.
8. Stone, M. C., Fishkin, K., and Bier, E. A. 1994. The movable filter as a user interface tool. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Celebrating interdependence* (Boston, Massachusetts, United States, April 24 - 28, 1994). B. Adelson, S. Dumais, and J. Olson, Eds. CHI '94. ACM Press, New York, NY, 306-312.
9. Ware, C. and Lewis, M. 1995. The DragMag image magnifier. In *Conference Companion on Human Factors in Computing Systems* (Denver, Colorado, United States, May 07 - 11, 1995). I. Katz, R. Mack, and L. Marks, Eds. CHI '95. ACM Press, New York, NY, 407-408.

The columns on the last page should be of approximately equal length.