

Table-Centric Interactive Spaces for Real-Time Collaboration

Daniel Wigdor^{1,2}, Chia Shen¹, Clifton Forlines¹, Ravin Balakrishnan²
¹Mitsubishi Electric Research Labs
Cambridge, MA, USA
www.merl.com
shen | forlines @merl.com

²Department of Computer Science
University of Toronto
www.dgp.toronto.edu
dwigdor | ravin @dgp.toronto.edu

ABSTRACT

Tables have historically played a key role in many real-time collaborative environments, often referred to as “war rooms”. Today, these environments have been transformed by computational technology into spaces with large vertical displays surrounded by numerous desktop computers. However, despite significant research activity in the area of tabletop computing, very little is known about how to best integrate a digital tabletop into these multi-surface environments. In this paper, we identify various design requirements for the implementation of a system intended to support such an environment. We then present a set of designs that demonstrate how an interactive tabletop can be used in a real-time operations center to facilitate collaborative situation-assessment and decision-making.

Categories and Subject Descriptors

H.5.2 [User Interfaces]: Interaction styles

General Terms

Design, Experimentation, Human Factors.

Keywords

Tabletop interaction, interactive spaces, groupware, real-time collaboration.

1. INTRODUCTION

Although centrally located tables have historically been the foci of collaborative activity, much of the research into computationally enabled interactive spaces [11, 13, 17, 21] has primarily concentrated on how to facilitate the transfer of data, redirection of application windows, and redirection of mouse and keyboard input amongst interactive whiteboards and personal computing devices. This body of research informs the design of interactive spaces for the common office environment, spaces in which collaborative brainstorming and spontaneous discussions dominate. The tasks carried out in these spaces are usually “open-ended” without stringent time constraints or the risk of catastrophic consequences associated with the collaboration.

In Heath and Luff’s study of the London Underground [11], they identify several meta-level organisational and communicational activities that are engaged-in by those monitoring and attending to the real-time operation of the London subway system:

1. Surreptitious monitoring and interrelating tasks.
2. Rendering activities visible.
3. Overseeing the local environment of events and activities.
4. Shaping tasks and coordinating activities.

Copyright 2006 Association for Computing Machinery. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the U.S. Government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

AVI '06, May 23-26, 2006, Venezia, Italy.

Copyright 2006 ACM 1-59593-353-0/06/0005...\$5.00.

From the descriptions offered of some of these tasks, it is clear that a shared interactive table might benefit users by supporting and reducing the activity required to carry-out these meta-level activities. Heath and Luff point out, for example, the value of the coordinating influence of shared status-information displays and of overhearing telephone conversations in avoiding the need to manually coordinate tasks. This incidental coordination is enhanced by an interactive touch-table, where its direct-touch interface make both the operation and its target visible to all participants. It is operation centres, such as the one described by Heath and Luff, which are the primary targets of the interaction paradigms we have developed.

In addition to supporting real-time system control centres, our research is also intended to inform the design of dedicated meeting rooms, sometimes referred to as “war rooms” or “extreme collaboration” spaces, being used in sectors including large companies (e.g., www.pbworld.com) and operations centres (e.g., www.nyc.gov/html/nypd/). We envision that our techniques will be useful as part of a larger system which mixes table-based and direct ancillary display manipulation.

In this paper, we present solutions that enable complete control over multiple display surfaces *directly* from an interactive table for domains in which real-time collaboration is essential. We believe that recent developments in digital tabletops [8, 17] can be exploited to enable a return to table-centric spaces which can be valuable in supporting face-to-face real-time collaborative decision-making while simultaneously controlling and exploiting the additional information capacity of auxiliary displays.

We begin by reviewing related work, then present six design requirements for the building of an effective table-centred control space which includes multiple ancillary vertical displays. Finally, we present the design solutions which made-up a prototype system we developed.



Figure 1. The PB Cave (www.pb.com). A space containing many elements of our envisioned environment.

2. RELATED WORK

Computer-augmented collaborative environments have been the subject of a great deal of research. In the early projects, users were each equipped with a personal computer, generally at a nearly 1:1 computer-to-person ratio. Additionally, one or more large-scale screens may be used to display information of interest to the group as a whole, or as ancillary displays controllable from individual participants' workstations. In Engelbart and English's system [9], Begeman et al.'s *Project Nick* [3], and Xerox PARC's *CoLab* project [21] participants are seated at workstations arranged around a table, leveraging some of the affordances of table-centred interaction albeit without an interactive tabletop surface. In Koike et al's *EnhancedTable* [13], and Rekimoto and Saitoh's *Augmented Surfaces* [18], personal workstations are enhanced with either tables or other vertical surfaces.

Tabletop computers have also appeared in more recent multi-display collaborative environments. The iRoom project [11] extended CoLab in several ways, including the addition of a computer embedded into a table. The i-Land project [22] included personal workstations built into individual chairs, as well as wall and table systems. In the MultiSpace system [10], users interact directly with a shared table, a laptop, or a wall-sized display, passing objects between screens by dragging them to "conduits" or "portals" represented graphically on each device.

However, none of these prior research systems investigated solutions for the direct control of input interaction of ancillary displays from a shared tabletop.

3. DESIGN REQUIREMENTS

The prior work focused on allowing users to move among and use the different displays in a distributed manner. Although direct-interaction with the ancillary displays might provide more flexibility, it is desirable that the full-range of actions that can be performed on these displays be supported from the table in *real-time* for several reasons:

- all virtual elements are within reach of all participants
- participants can remain comfortably seated
- a consistent input and interaction paradigm is maintained
- advantages of table-centred interaction may be leveraged

Given these advantages and opportunities in the design space, and given Heath and Luff's observational evidence, as well as recent developments in interactive tabletop input technology [8,17], our goal is to explore the scenario where *all* interaction can take place on the tabletop, allowing multiple users to simultaneously interact – directly from the tabletop using multi point direct touch input – with the full content of multiple surrounding displays.

In order to attain this goal, we identified six design requirements that would need to be met in order to build a such a system.

a. Support Non-Linear Alignment of Displays

Many pointer-based systems that facilitate multiple-display interaction require that the edges of the various displays be linearly, or nearly linearly aligned with one another [1]. Although aligning the vertical displays with the edges of the table might simplify the problem, several factors make this unrealistic:

- the number of displays may outnumber the table edges.
- space is required for users and observers to sit.
- vertical displays can be positioned to optimise viewing angles.
- physical constraints in the room may prevent precise alignment.

Systems designed for our problem space must enable fluid and fast interaction despite the lack of rectilinear alignment of the edges of the displays. In our exploratory environment, two wall displays are located off the corners of the table. While this may seem like a minor requirement, its inclusion in the design space makes the adoption of many existing interaction techniques [1] infeasible, thus necessitating new designs.

b. Awareness of Input-Output Mapping

In a table-centred space, tabletop input must be shared amongst the various displays. Thus, appropriate visualizations must be provided so that users maintain an awareness of which display is being controlled at all times.

c. Contextual Association

Previous efforts [2,4,16,19] have shown the strength of visual contextual associations in large display applications. Objects that are related to one another are drawn together by a visible "connection" object, which visually articulates the link between them. Visual association of related items is particularly important as related items are separated across displays.

d. Maintain Direct and Absolute Input Paradigm

Typically, tabletop systems combine a touch-surface with a projected image, calibrated such that the input and visual space are directly overlaid. This affords an interface where on-screen artefacts are manipulated by direct touch in an absolute one-to-one mapping. To maintain a single consistent input paradigm, it is likely desirable to maintain the direct touch input style while using the interactive table to send input to the ancillary displays.

e. Multiple Granularities of Interaction

A table used as an input device must provide not only quick-and-coarse interaction, but also precise control of ancillary displays. Thus, multiple granularities of input are necessary, and switching between these granularities must be simple.

f. Support for Non-Interfering Concurrency

In collaborative spaces, it is also important that users be able to work on a particular sub-task without interfering with others. Two types of interference must be considered: physical interference, where one user's actions physically "collides" with another; and visual interference, where a task performed by one user interferes with the visual display space of another.

4. DESIGN SOLUTIONS

To facilitate interaction in our table-centred multi-display collaborative space, we designed interaction and visualization techniques to address the issues identified in the previous section.

4.1 Visual Connectivity between Displays

To provide a sense of visual and spatial continuity and connectivity among the various spatially non-aligned displays in our interaction space, we leveraged Gestalt principles, including closure and continuity. On the ancillary display, we placed a repeating pattern on the bottom edge of the screen, symmetrical to the pattern of a proxy to each ancillary display shown on the tabletop (Figure 2). Although subtle, this visual connection between the ancillary displays and the tabletop helps to promote a sense of spatial cohesiveness, and helps to establish the virtual topology of the system.




Figure 2. Schematic diagram of our system: the matching colours and shapes of the repeating pattern on the walls and proxies on the table allows precognitive connections.

To further reinforce the connection between the screens, when a user comes into contact with the proxy, the intensity of its colour subtly increases and decreases several times. This colour throbbing is mimicked by the proxy on the ancillary display, drawing attention and thus strengthening the perceived connection.

The use of both color-pairing and spatial proximity of the proxies is important: each provides a precognitive connectivity cue, as well as a clearly distinguishable characteristic for conscious searching. The power of each may be reduced in some environments: as more screens are added, the distinctions between colors will become more subtle, reducing the pop-out effect of matched coloring. Relying on proximity is also limiting, since some environments may be ambiguous (e.g., if two ancillary displays are vertically stacked, it is infeasible to distinguish them by proximity to proxies on the table's edge).

4.2 Enhanced Contextual Associations

As associated objects become separated across displays, there is an increased need to assist users in understanding this association. In previous work, contextual associations, dubbed “CoR²DS” by Shen et al. [19], were shown graphically on the same screen. We have built upon their ideas, which were originally intended to support contextual menu functionality, and extended them to assist in general contextual association. We have maintained the directionality of the relationships, by drawing the associations as growing “out” from some point on the underlying object, to nearly encompass the entirety of the sub-object (Figure 3). We also expanded the role of the proxy to allow it to act as a virtual conduit through which cords could be mapped to other displays. In this design, all cords between objects on different displays would begin or end at the proxy for the display on which the associated object was located. Although requiring that the cords pass through the proxy might limit visual connectivity across displays, this routing enables an understanding of the connection when the display and table are arranged at odd physical angles. As shown in Figure 3, cords passing through the same proxy are distinguished by colour. To address the issue of visual clutter, cords are displayed for a few seconds when an associated object is touched and disappear shortly thereafter. Touching a proxy object displays all cords passing through it. Although no attempt is made to maintain the visual continuity of the cord between displays, having the cords in the same colour, as well as appearing and disappearing in unison, provides some precognitive cues.




Figure 3. Cords are routed between displays via the proxy objects described previously. Left: Diagrammatic view. Right: Different view with real system screenshots overlaid.

4.3 World in Miniature (WIM)

Although there exist several techniques that facilitate control over remote surfaces [1, 2, 4], the demonstrated utility [14] of radar views led us to explore its use in our interaction space. A radar view is a world in miniature (WIM) [16], where a remote environment is displayed in a scaled format in the work area, and manipulations within the scaled miniature are transferred to the original space. The effectiveness of this technique has been demonstrated in a similar environment as a ‘coordinated synchronized view’ [22]. In our environment, interactions performed on the WIM on the tabletop would directly impact the corresponding region of the ancillary display. This has several desirable properties:

- a direct-touch, table-centric paradigm is maintained
- fast movement across multiple displays is possible by simply moving between multiple WIMs.
- every point on the remote display is selectable, which is not true systems [2, 4] that collapse white space around objects.
- users can comfortably view and manipulate screens that might otherwise require body contortions while seated at the table.

WIMs were integrated into the proxy objects, such that a WIM of the ancillary display was shown below the matching semi-circle. We included a control to display a second WIM window, visually tethered, which could be freely moved, resized and rotated about the table. Further, we surrounded the WIM with a graphical bevelled edge, shaded to match the colour of the proxy. This gives the appearance of depth, providing a window-like feel to the WIM (Figure 4).

In some systems, a WIM approach is already being used to control large ancillary displays from a control terminal using remote desktop software, such as VNC (www.vnc.com). In our work, we have augmented the WIM concept with dynamic orientation and zoom control, as well as multi-user telepointers and a meta-level zoom to be discussed shortly. Combined, these innovations make the WIM more appropriate for a table-centred control system.




Figure 4. Left: screenshot of an ancillary display. Right: screenshot of the tabletop, including the WIM view of the ancillary display (top-right) and additional proxies and their WIM (top-left and bottom-right).

4.4 Multiple Control Granularities

When manipulating objects on the ancillary display, users typically use the input space of the WIM, but keep their visual focus on the larger screen. Because movement in the WIM is exaggerated by the scale difference between the displays, fine operations become difficult. To enable fine-grained operations, we must reduce the control-display (CD) gain between the WIM and associated display, without sacrificing favourable aspects of the WIM design. We provide both a manual size control and a button to toggle between two preset sizes to reduce the CD gain.

We also need to provide users with the ability to adjust the gain without increasing the WIM's size, since a large ancillary display could mean that even a WIM occupying the entire table might not provide sufficiently fine-grained control. One solution is to allow the contents of the WIM, and the associated display, to be zoomed and thus increase/decrease the CD gain. Unfortunately, this affects the content shown on the ancillary display and can interfere with concurrent actions by other users. Our solution was to allow the user to zoom the WIM view, but leave the view on the ancillary display unchanged. By maintaining direct-touch control within a zoomed WIM view, the result is a reduction in the CD gain (Figure 5).



Figure 5. Partial screenshot of tabletop display: the three stages of a zoom of the WIM.

4.5 Telepointers

As described previously, we wished to facilitate the use of the WIM as an input-only space, allowing users to use the larger, ancillary displays as their visual focus. Further, we felt it necessary to provide peripheral observers not seated at the table with cues to facilitate their understanding of operations being performed on the ancillary displays. To support these goals, we added pointers to the ancillary display: whenever a user was touching the WIM, the corresponding point on the ancillary display would show a pointer, as depicted in Figure 6. We augmented these pointers with a colour coding to uniquely identify the actions of each user. These pointers can also be used by participants as a visual reference point for discussion (e.g., “look at this”), and also to reduce interference with concurrent tasks performed on the ancillary displays, by providing other users with awareness of the current focus of one's work.



Figure 6. Each user's point of contact on one or more WIMs is shown on ancillary displays, uniquely identified by colour.

4.6 Meta World in Miniature

Although the colourisation and positioning of the proxy are helpful in establishing the mapping of WIM to display, we found that it was sometimes difficult to achieve a quick understanding of which WIM matched which display. To improve this, we extended the WIM to represent not only a single ancillary display in miniature, but also the system as a whole. To achieve this, we added a control to the WIMs to allow the user to “zoom out” to a photograph of the work area. The WIM animates to the approximate location of the physical display it represents (Figure 7). The user can then touch any of the displays in the photo to cause the WIM to zoom into that screen. This approach is similar conceptually to ARIS [5, 6], where application windows are presented iconically in an abstract schematic of the system's screens. Our approach differs in that the actual screen content is shown, rather than iconic representations, which is more practical for environments where applications occupy the entire display.




Figure 7. Top: normal appearance of WIM. Bottom: WIM has animated to the ‘meta’ level, showing the space as a whole.

4.7 Moving Objects Between Displays

Moving objects between displays using the WIM is accomplished by dragging them from one WIM to another, or from a WIM to the general tabletop area and vice-versa. Figure 8 shows an object being moved from the tabletop to an ancillary display. Note that the orientation of the object changes once it has been moved off of the table. Although various orientations might be desirable for objects on the tabletop so that they can, for example, face a particular user, once they are moved to an ancillary display all users share a common “up” vector, and as such the object orientation is corrected for easier viewing.



Figure 8. An object is moved from the tabletop to an ancillary display by dragging it onto a WIM view. The orientation of the object is corrected once it is placed on the vertical display.

5. CONCLUSIONS AND FUTURE WORK

We have presented an exploration of table-centric interactive spaces focused on real-time collaboration, where interaction with both tabletop and multiple vertically mounted large displays are controlled solely from the interactive tabletop. Our contributions are twofold: identification of interaction and visualization issues that arise in the given problem space of single tabletop augmented with multiple ancillary displays, and the development of a suite of interaction and visualization techniques designed to address those issues. Coupled with a real-world usage scenario and user study, the end result is a better understanding of how such table-centric spaces can be best utilized for collaborative applications and a prototype interface that facilitates such use.

Concurrently with the preparation of the present paper, we have conducted two forms of evaluation of our interaction techniques: the first is the development of an application scenario, the second is a user study exploring the early learnability of the techniques. The application scenario is a mock-up of a real-time monitoring centre for the New York Police Department, dedicated to situational assessment and deployment of police forces. Details of the application scenario, including feedback from the New York Police Department, as well as the results of the user study, are documented in [23].

In the near future, we intend to further refine our interface and potentially deploy them in these real usage settings. We also intend to integrate our designs with existing interfaces already in use in these settings. Finally, the next steps in this research include looking at supporting multiple tables, much larger variety of displays, and participation by those working at displays away from the tabletop.

6. ACKNOWLEDGMENTS

We thank the AVI reviewers for their comments, Mark Hancock for assistance with application development, Ed Tse, Alan Essenther, and Kathy Ryall for valuable feedback on our system.

This study was partially supported by the Advanced Research and Development Activity (ARDA) and the National Geospatial-intelligence Agency (NGA) under Contract Number HM1582-05-C-0028. The views, opinions, and findings contained in this report are those of the author(s) and should not be construed as an official Department of Defense position, policy, or decision, unless so designated by other official documentation.

7. REFERENCES

1. Baudisch, P., Cutrell, E., Hinckley, K., & Gruen, R. (2004). Mouse ether: accelerating the acquisition of targets across multi-monitor displays. *CHI Ext. Abstracts*. p. 1379-1382.
2. Baudisch, P., Cutrell, E., Robbins, D., Czerwinski, M., Tandler, P., Bederson, B., & Zierlinger, A. (2003). Drag-and-pop and drag-and-pick: Techniques for accessing remote screen content on touch- and pen-operated systems. *INTERACT*. p. 57-64.
3. Begeman, M., Cook, P., Ellis, C., Graf, M., Rein, G., & Smith, T. (1986). Project Nick: meetings augmentation and analysis. *CSCW*. p. 1-6.
4. Bezerianos, A. & Balakrishnan, R. (2004). The Vacuum: Facilitating the manipulation of distant objects. *CHI*. p. 361-370.
5. Biehl, J. & Bailey, B. (2004). ARIS: an interface for application relocation in an interactive space. *GI*. p. 107-116.
6. Biehl, J. & Bailey, B. (2005). A scalable, portal-based interface for managing applications in interactive workspaces. *Technical Report UIUCDCS-R-2005-2542, 2005. University of Illinois at Urbana-Champaign*.
7. Dempsey, C., Laurence, D., & Tuovinen, J. (2002). Gestalt theory in visual screen design: a new look at an old subject. *Seventh World Conference on Computers in Education: Australian topics - Volume 8*.
8. Dietz, P. & Leigh, D. (2001). DiamondTouch: a multi-user touch technology. *UIST*. p. 219-226.
9. Engelbart, D. & English, W. (1968). A research center for augmenting human intellect. *AFIPS Fall Joint Computer Conference*. p. 295-410.
10. Everitt, K., Shen, C., Forlines, C., & Ryall, K. (2006 - in press). MultiSpace: Enabling electronic document micro-mobility in table-centric, multi-device environments. To appear in *IEEE TableTop 2006*.
11. Heath, C. & Luff, P. (1992). Collaboration and Control: Crisis Management and Multimedia Technology in London Underground Line Control Rooms. *Journal of CSCW*, 1.1, p. 24-48.
12. Johanson, B., Hutchins, G., Winograd, T., & Stone, M. (2002). PointRight: experience with flexible input redirection in interactive workspaces. *UIST*. p. 227-234.
13. Koike, H., Nagashima, S., Nakanishi, Y., & Yoichi Sato. (2004). EnhancedTable: Supporting small meetings in ubiquitous and augmented environment. *IEEE Pacific-Rim Conf. on Multimedia (PCM2004)*. p. 97-104.
14. Nacenta, M., Aliakseyeu, D., Subramanian, S., & Gutwin, C. (2005). A comparison of techniques for multi-display reaching. *CHI*. p. 371-380.
15. Pierce, J., Stearns, B., & Pausch, R. (1999). Two handed manipulation of voodoo dolls in virtual environments. *ACM Symposium on Interactive 3D Graphics*. p. 141-145.
16. Rekimoto, J. (1997). Pick and drop: A direct manipulation technique for multiple computer environments. *UIST*. p. 31-39.
17. Rekimoto, J. (2002). SmartSkin: an infrastructure for freehand manipulation on interactive surfaces. *CHI*. p. 113-120.
18. Rekimoto, J. & Masanori Saitoh. (1999). Augmented Surfaces: a spatially continuous work space for hybrid computing environments. *CHI*. p. 378-385.
19. Shen, C., Vernier, F., Forlines, C., & Ringel, M. (2004). DiamondSpin: An extensible toolkit for around the table interaction. *CHI*. p. 167-174.
20. Shen, C., Hancock, M., Forlines, C., & Vernier, F. (2005). CoR2Ds. *CHI Ext. Abstracts*. p. 1781-1784.
21. Stefik, M., Bobrow, D., Lanning, S., & Tartar, D. (1987). WYSIWIS revised: early experiences with multiuser interfaces. *ACM Trans on Info. Systems*, 5(2). p. 147-167.
22. Streitz, N., Geißler, J., Holmer, T., Konomi, S.i., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., & Steinmetz, R. (1999). i-LAND: an interactive landscape for creativity and innovation. *CHI*. p. 120-127.
23. Wigdor, D., Shen, C., Forlines, C., Balakrishnan, R. (2006). Table-Centric Interactive Spaces for Real-Time Collaboration: Solutions, Evaluation, and Application Scenarios. *MERL Tech Report # TR2006-017*. (<http://www.merl.com/publications/TR2006-017/>).