

Experiences with and Observations of Direct-Touch Tabletops

Kathy Ryall¹, Meredith Ringel Morris², Katherine Everitt³, Clifton Forlines¹, Chia Shen¹

¹*Mitsubishi Electric Research Laboratories, Cambridge MA, USA*

²*Stanford University, Stanford, CA, USA*

³*University of Washington, Seattle, WA, USA*

{ryall, forlines, shen}@merl.com, merrie@cs.stanford.edu, everitt@cs.washington.edu

Abstract

The design of multi-touch multi-user tabletop user interfaces is still in its infancy and is not yet well understood. To date, published experimental results have primarily focused on controlled user studies. In this paper, we present observations of user experience “in the wild” on interactive tables in four different real-world contexts – all non-controlled settings. We reflect upon our collective experience, report our observations, and summarize lessons learned by identifying design considerations relating to several aspects of interactive tables, such as simultaneous touching, ambiguous input, one-fingered touch, finger resolution, alternate touch input, crowding and clutter, text input, orientation, multi-user coordination, occlusion, ergonomic issues, and mental models.

1. Introduction

Our collective experiences from observing users of interactive tabletop systems in four distinct contexts over the past two years have revealed several interesting, recurring themes and issues in interactive tabletop computing. In this paper, we present the practical insights gleaned from our hands-on experiences. We have organized our observations and insights according to three key aspects of tabletop systems: (1) direct-touch interaction, (2) the content and layout of information, and (3) the physical setup of the interactive furniture. This collection of observations is intended to serve as a complement to the growing body of controlled experimental studies of the use of horizontal computing systems.

The observations in this paper are derived from experiences with a direct multi-touch, multi-user tabletop called DiamondTouch [2]. Direct multi-touch tables are surfaces on which input sensing and output displays are superimposed, and on which multiple touches can be detected simultaneously. DiamondTouch provides these capabilities and also

offers the utility of identifying which user is touching which particular location on the surface. While our experiences are with DiamondTouch technology, we believe our findings hold true across multi-user tabletop technologies. Throughout this paper when we use the term “table” we are referring to a direct multi-touch, multi-user tabletop.

2. Observational Contexts

Our observations stem from four contexts that vary by target user community, application type, and setting. In contrast to controlled user studies, these four observational situations were non-controlled settings, ranging from casual to focused usage scenarios. People chose to interact with the tables and applications of their own volition; their interactions were unstructured and free-form in nature. In some cases simple instructions were provided, while in others people simply walked up and started using the table. In all four contexts both individual and group interactions were observed by a subset of the authors, who took written notes during usage sessions for post-hoc analysis.

When discussing our experiences, we realized we had observed several common patterns of behavior across these varied use scenarios and contexts. We summarize these observations in this paper, as a knowledge base for other designers of interactive table systems.

2.1 Lobby Table

In June 2003 a DiamondTouch table was installed in the lobby area of Mitsubishi Electric Research Laboratories (MERL), an industrial research lab. The table is at coffee-table height, and centered amongst four comfortable leather chairs, representative of a typical waiting area or lounge (Figure 1). It is situated in a high-traffic area across from the receptionist’s desk. The applications running on the table include a

set of multi-user games along with a set of new research demos. The games run unsupervised, and most visitors play them with little or no instructions. The games (Figure 1, bottom) range in style from collaborative to competitive; they include E-Magnetic Poetry [13], in which users manipulate virtual magnets to create poems either independently or collaboratively, Pop-A-Bubble [2], an electronic version of Whack-A-Mole in which players compete against each other, and CircleMaze [1], a collaborative interactive maze. Hundreds of lab visitors, everyone from technologists to university students to researchers' children, have passed through MERL's lobby and casually interacted with this table over the past two years.

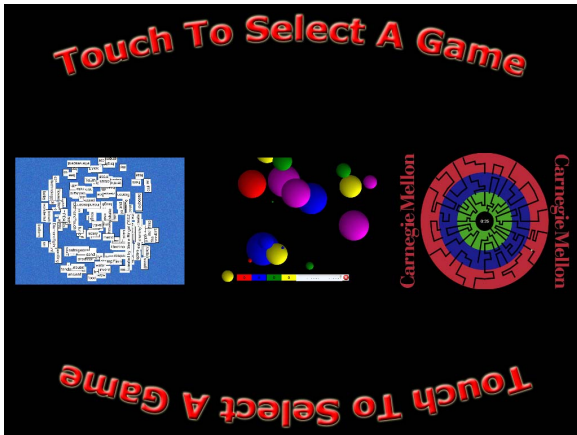


Figure 1. The table in the industrial research lab featured several game and research applications. Top: table in action. Bottom: screenshot of the game kiosk (E-Magnetic Poetry, Pop-A-Bubble, CircleMaze).

2.2 Biologists' Field Data Annotation Table

At Stanford University, we have developed TeamTag (Figure 2), a multi-user tabletop interface that supports collaboration among bio-diversity researchers. One to four researchers sit around the table to browse, label, and search through digital photographs of flora and fauna taken during field expeditions. Photos can be moved about the table, reoriented to face different members of the group, resized to view different levels of detail, and organized into piles. Labeling these images collaboratively allows the researchers to bring their collective expertise to bear on identifying the subjects of the photographs. About ten biologists have been observed using the system, in sessions ranging from thirty minutes to three hours in length. The table is at "desk height," with people seated at the table with their legs underneath.




Figure 2. TeamTag is a collaborative tabletop application for bio-diversity researchers. Top: table in action. Bottom: screenshot of the photo-labeling application.

2.3 NextFest 2004 Tables

Organized by WIRED magazine, the three-day NextFest 2004 [20] was designed to give the general public a close-up, hands-on view of innovative technology. We brought two tables to this conference, which were part of the Future of Design Pavilion. During the course of this event, the tables were used by almost 2,000 people. Visitors included children, educators, executives, designers, and engineers. On one table, the software included a general-purpose application where people view images and create text (Figure 3), an educational game called Habitat (in which multiple people match pictures of animals to pictures of their home environments: land, forest, sea, or ice floes), and a finger-paint program in which multiple people draw together on a large digital canvas. On a second table, we ran a set of multi-user games, a



Figure 3. At NextFest people tried out several tabletop applications, including photo browsing and annotation software. Top: table in action. Bottom: screenshot of photo browsing application.

subset of the games included on the lobby table described in Section 2.1. Both tables were at “desk height,” with people seated with their legs underneath.

2.4 GeoINT 2004 Map Visualization Table

GeoINT is a two-day symposium organized by the National GeoSpatial Agency [19] to foster collaboration and interoperability amongst technology providers in the sector of geospatial exploration, analysis, and presentation applications. We demonstrated a multi-user map application, DTLens [3], in which we used the IDELIX’s PDT lens API (www.idelix.com) for data exploration (Figure 4). Up to four users can simultaneously open and use personalized zooming lenses on geospatial map or diagrammatic data. Several hundreds visitors, all from the geospatial information analysis community, interacted with our tabletop application.



Figure 4. Multi-user zooming lenses were featured in the tabletop map-browsing demonstration at GeoINT. Top: Table in action. Bottom: screenshot showing two people’s lenses.

3. Observations

Observing the usage of interactive tabletops in these four contexts has been extremely valuable. We noticed many interesting usage patterns that we would never have predicted without watching people actually using and experimenting with the tables. More surprising, many of the same issues occurred in most or all of the four observation scenarios, which was interesting given the varied user populations, applications, and settings across the four contexts. This section highlights many of the recurring issues that we observed across three key dimensions: touch interactions, organization of content and physical setup.

3.1 Touch Interactions

3.1.1 Simultaneous Touching

At first, some people are hesitant to touch the table at the same time. Even though the DiamondTouch and its applications can support simultaneous touch input from up to four users, some people are hesitant to touch the table simultaneously, especially when they are first introduced to the technology. This has been observed to be more true with adult users than children. While some of this hesitance is likely due to conditioning from single-user technology, cultural issues also play a role. For instance, Japanese users are particularly hesitant to simultaneously interact on the surface due to cultural standards of polite behavior.

Even when users simultaneously interact with the tabletop, they are concerned about accidentally bumping arms with another user, or brushing against another user's hand. These concerns are more prevalent among groups that do not know each other well, and can impact the acceptability of tabletop UI designs. For instance, although a group of biologists that knew each other didn't mind using a tabletop on which they shared a common set of widgets in the center of the table, groups that were less familiar with each other complained that this design made them self-conscious about the possibility of physical contact with other group members. Work from the field of proxemics [5] (a field concerned with the distances people maintain amongst each other) could inform the design of more acceptable tabletop interfaces.

3.1.2 Ambiguous Touch

Accidental input is common, especially when pointing at something on the table. Every touch counts on a direct-touch input device. As a result, there is a general problem of distinguishing between deictic

pointing (or touching) and touching with the intent to interact. Accidental touches (such as when a person leans his wrists or elbows on the touch surface, as people are accustomed to do with traditional tables) are also misinterpreted as touch-to-interact.

3.1.3 One-fingered Touch

At first, users gravitate toward single-finger interaction. Since gesture-based interfaces are frequently described as "intuitive," we were surprised to notice that users rarely attempted to perform gestures on the tabletop. Although multi-handed and multi-fingered interfaces seem natural, the majority of users only attempted to use the tabletop in a manner that mimics standard stylus or touch-screen interfaces – by tapping and dragging with a single finger. It is interesting that users carry over this bias from other interface form-factors onto the tabletop display. However, once users are informed that the table is capable of recognizing multi-hand, multi-finger input, they then try to experiment with this style of interaction. Techniques for making acceptable gestures more apparent to users, such as Vogel and Balakrishnan's self-revealing help [16], which showed video of available gesture techniques to users who lingered in front of a display wall, might encourage more immediate multi-hand interaction with interactive tables.

3.1.4 Finger Resolution

GUI elements designed for a mouse need modification for finger-based input. Direct-touch manipulation creates occlusion difficulties that make standard GUI interfaces less usable. Fingers occlude the text of menu items or buttons they are pressing, making it desirable to display labels slightly offset from where the finger touches, rather than directly underneath as in mouse-oriented interfaces.

A second challenge due to the coarser resolution of fingers (compared with a mouse pointer) is present when interacting with some widgets designed for mouse interaction. People's fingers are different sizes, causing aiming and targeting precision problems for large-fingered individuals. The default dimensions of title bars, buttons, and other widgets in standard GUI toolkits are sized to be targetable by mouse input, but are sometimes too small to be accurately selected by fingers.

3.1.5 Alternate Touch Input

Some people preferred to use a stylus (or other input device) to interact with the table rather than their

hands. We believed that one of the benefits of a direct-touch manipulation interface was the ability to interact with the table directly with one's hands, removing the need for other physical input devices (e.g., mouse or keyboard). We have had a number of direct requests, however, to support stylus-based, or other tangible (non-hand) input to our tables. Some of the desire may be for finer-grained input resolution (as described in Section 3.1.4), but many people noted that they were hesitant or uncomfortable to touch the table with their hands, especially in situations when it was clear that many other people were interacting with the tables with their hands. Tabletop interaction designers should be aware of the possible hygiene-related concern.

3.2 Organization of Content

3.2.1 Crowding and Clutter

Users appreciate their elbowroom. When designing interactive tabletops, taking the size of the display into account can be important. Some prior work has looked at the issue of table size in controlled, experimental settings [11]. From our observations of unstructured use, we have observed that a table measuring 107 cm diagonal is a good minimum size. Observations of smaller table sizes (80 cm diagonal) reveal that users frequently bump elbows and arms with each other while interacting. Also, this smaller table size makes it very easy to reach into the "personal territory" [12] of another user, thus increasing the likelihood of "bad behavior" [9] such as "stealing" digital media from another group member. A square-shaped table may be desirable to facilitate equal-area personal zones for each user, but current tabletop technologies are available only in the 4:3 aspect ratio of standard monitors and projectors.

Crowding of the display area is also a challenge with smaller tables. Interactive tabletops lend themselves to applications such as layout and organization tasks, which involve having multiple (10's to 100's) of items on the tabletop at a time, or visualization tasks in which multiple users zoom in and enlarge the documents to whatever is maximally possible on the table. We observe "bumping" and "overlapping" zoomed-in windows quite often. This crowding creates another potential conflict situation among users – for instance, users of TeamTag have been observed to position photos they are looking at on top of widgets that other group members may need to use. Employing techniques such as ZoomScapes [4] and "black-hole views" [15] to alleviate this clutter are vital to successful tabletop interfaces.

3.2.2 Text Input

Bare fingers are insufficient for text input. Our experiences with the table have shown that text input is particularly challenging. Providing virtual keyboards on the tabletop has proved a feasible, but tedious, solution. Graffiti-style input using "finger-ink" such as that provided by the DiamondSpin toolkit [13] is also not a practical solution for large amounts of text entry, because people draw large, clumsy shapes with fingers. Auxiliary input sources, such as wireless keyboards or PDAs with styli might be a promising solution to incorporating serious text-entry into tabletop groupware. However, our experiences with the tabletops suggest that tasks such as the organization, examination, or annotation of digital media are better suited to co-located tabletop collaboration than text-entry-based tasks.

3.2.3 Orientation

For some types of documents, orientation is not a problem. Orientation of information on interactive tabletops is more complex than on traditional, vertical displays. Orientation is relevant for clarity (e.g., the ability to view and understand information on the tabletop) as well as for communicative purposes [6] (e.g., indicating willingness to share a document with other people [10]). However, our observations of TeamTag suggest that proper orientation might not be needed for all types of information – users had no trouble comprehending small chunks of text that were improperly oriented, and did not take advantage of interface mechanisms that would have allowed them to rotate the text. For large amounts of text, orientation for clarity becomes more important.

Several solutions to the orientation issue have been explored, including providing a handle on each item to allow arbitrary re-orienting [13], using specialized multi-point-of-view display hardware [7], having items on the table automatically orient themselves to face the nearest edge [13], having an item automatically orient itself to face the user who most recently touched it [8], replicating critical information around each edge of the table, or using a "magnet" feature to re-orient all of the table's contents toward a particular user [13].

In our experience, the appropriateness of each of these solutions depends on the task and user population. In some applications the automatic techniques for orientation are sufficient for most users; providing them with explicit rotation handles on objects is distracting since they generally don't need that much control over the orientation beyond the four cardinal directions. As a result they then spend a lot of time fidgeting with the handles to get the orientation

“just right.” In other cases, such as with a collage-making application, the arbitrary orientation feature is useful for artistic layouts, but auto-rotation of other objects for viewing and reading is still preferred. While many orientation mechanisms are available for application developers, there isn’t yet a clear choice regarding which orientation techniques are most appropriate for specific application types.

3.2.4 Multi-User Coordination

The actions of multiple people often conflict with one another, both intentionally and accidentally. Although users are working together around the tabletop, we have observed that it is not unusual for a user’s actions to interfere with those of other group members. Typically this interference is accidental, although it can be intentional as well. Some of the types of conflicts we’ve observed include a user placing photos he is looking at on top of information that other users were viewing, a user taking digital documents away from another user who was interacting with them, a user changing the layout or contents of the tabletop while other users were interacting, users issuing conflicting menu commands (such as one user trying to open a new document while another user tries to clear all documents from the table), and a user selecting from another user’s menu. Coordination mechanisms, such as [9], may be desirable for tabletop groupware design.

3.3 Physical Setup

3.3.1 Occlusion

Concerns about shadowing caused by top-projected displays are not a problem in practice. Many of today’s interactive tabletop technologies use projection, typically top-projection, to co-locate a display on the input device. From our experience with both casual and focused usage scenarios, we have not observed any instances where the shadowing that results from top-projection presented a problem in practice, and users immediately forget that this was a concern once they begin to use the device. In practice, the shadowed area is not larger than the area that is naturally occluded by the opacity of people’s arms and hands. In fact, in some respects top-projection is superior to bottom-projected or LCD-type surfaces, since the content is still projected, albeit with reduced legibility, on top of the arm and hand. Wu and Balakrishnan [18], for example, take advantage of this feature of top-projection by intentionally displaying private data on a user’s cupped palm.

3.3.2 Ergonomics

The design of the table’s edge and its height impacted its use. Tabletop users tended to lean on their elbows or arms while interacting with the device. For a touch-sensitive surface, such as DiamondTouch, this tendency was problematic, as it led to false detections of interaction. Providing a wider non-sensitive surface around the outer edge of the interactive surface, or designing software that selectively ignores this class of touches, can improve the usability of interactive tabletops. The tables used for the observations described in this paper (Figures 1-4) had only a minimal bezel of 2-2.25 inches. We are now experimenting with larger offsets, as shown in Figure 5.





Figure 5. New furniture prototypes with larger non-sensitive areas surrounding the touch-sensitive surface.

Changing the height of the tabletop impacts the manner in which users interact with it. A low-set, coffee-table arrangement (such as Figure 1, our “Lobby Table”) is conducive to casual interactions, such as photo browsing, while a desk-height placement (Figures 2, 3, and 4) is more suited for productivity tasks. Table height can also impact reachability and readability of the display, and is an important factor for

developers to keep in mind when designing tabletop groupware.

Two other considerations are the angle of the tabletop, and whether people are seated or standing. To date we have primarily focused on seated tables, with horizontal displays. Figure 6 shows two other configurations we have developed, but of which we do not yet have long-term observations.



Figure 6. Other prototype form factors. Top: A “drafting table” with an angled surface, used as a desktop computer replacement in an office. Bottom: A walk-up public kiosk to support standing interactions.

3.3.3 Mental Models

Users do not view the interactive tabletop as a “computer.” Perhaps most surprisingly of all, many people using the table and its applications were unaware of the fact that it was merely a novel input device for a traditional desktop computer. It seems that interactive tabletops may have already quickly

achieved Mark Weiser’s notion of ubiquity [17]. In some cases, when exiting an application people were surprised to see the Windows desktop appear. Overall, people seemed to find the tabletop and its applications less intimidating, and often more playful, than a traditional computer desktop environment. Most took it for granted that it should support multi-user activities.

In many instances people thought the table itself was the display and were not aware of the projector overhead. Although the display resolution of the table is determined by the quality of the projector used and is not inherent in the table itself, users consistently thought that the table was a display device, even though they often saw us turn on the projector. While the projectors we used had a resolution of either 1024x768 or 1280x1024 pixels (lower than the resolution of most new PC monitors), the physically large size of the table gave users the impression that the pixel resolution was much higher than that of a typical monitor.

For instance, one biologist viewing photos on the tabletop commented that she could view much larger and more detailed blow-ups of her photos on the tabletop than on her monitor, even though the tabletop did not in fact have a higher display resolution. This suggests that physically large tabletops add a perceived value greater than the actual qualities that they possess. Research has shown that large vertical displays offer greater benefits than equi-resolution, but physically smaller displays [14]. Our observations that users perceive the tabletop to be higher resolution suggest that large horizontal displays warrant further investigation in terms of perceptual benefits they may offer.

4. Conclusions

We have presented our observations of and experiences with interactive tables, collected over two years in four broad usage contexts. Although we observed diverse groups of people (i.e., varying in age, culture, gender and technical background) in a range of non-controlled settings (i.e., from casual to focused usage scenarios) across a broad set of applications (i.e., from collaborative to independent to competitive), we discovered a number of common usage patterns. These experiences revealed issues with implications for the design of direct-touch tabletops. We hope that these observations compliment the set of published controlled user study results and offer further insights into our collective understanding of interactive tabletop technology design.

5. References

- [1] Blaine, T. and Forlines, C. JAM-O-WORLD: Evolution of the Jam-O-Drum Multi-player Musical Controller into the Jam-O-Whirl Gaming Interface. *Proceedings of the 2002 Conference on New Instruments for Musical Expression*, Dublin, Ireland, May 2002.
- [2] Dietz, P. and Leigh, D. DiamondTouch: A Multi-User Touch Technology. *Proceedings of UIST 2001*, 219-226.
- [3] Forlines, C. and Shen, C. DTLens: Multi-user Tabletop Spatial Data Exploration. *Proceedings of UIST 2005*, in press.
- [4] Guimbretiere, F., Stone, M., and Winograd, T. Fluid Interaction with High-Resolution Wall-Size Displays. *Proceedings of UIST 2001*, 21-30.
- [5] Hall, E.T. *The Hidden Dimension*. New York: Anchor Books, 1966.
- [6] Kruger, R., Carpendale, M.S.T., Scott, S.D., and Greenberg, S. How People Use Orientation on Tables: Comprehension, Coordination, and Communication. *Proceedings of GROUP 2003*, 369-378.
- [7] Matsushita, M., Iida, M., Ohguro, T., Shirai, Y., Kakehi, Y., and Naemura, T. Lumisight Table: A Face-to-face Collaboration Support System That Optimizes Direction of Projected Information to Each Stakeholder. *Proceedings of CSCW 2004*, 274-283.
- [8] Morris, M.R., Morris, D., and Winograd, T. Individual Audio Channels with Single Display Groupware: Effects on Communication and Task Strategy. *Proceedings of CSCW 2004*, 242-251.
- [9] Morris, M.R., Ryall, K., Shen, C., Forlines, C., and Vernier, F. Beyond "Social Protocols": Multi-User Coordination Policies for Co-located Groupware. *Proceedings of CSCW 2004*, 262-265.
- [10] Ringel, M., Ryall, K., Shen, C., Forlines, C., and Vernier, F. Release, Relocate, Reorient, Resize: Fluid Techniques for Document Sharing on Multi-User Interactive Tables. *Proceedings of CHI 2004 Extended Abstracts*, 1441-1444.
- [11] Ryall, K., Forlines, C., Shen, C., and Morris, M.R. Exploring the Effects of Group Size and Table Size on Interactions with Shared-Display Groupware. *Proceedings of CSCW 2004*, 284-293.
- [12] Scott, S.D., Carpendale, M.S.T., and Inkpen, K. Territoriality in Collaborative Tabletop Workspaces. *Proceedings of CSCW 2004*, 294-303.
- [13] Shen, C., Vernier, F., Forlines, C., and Ringel, M. DiamondSpin: An Extensible Toolkit for Around-the-Table Interaction. *Proceedings of CHI 2004*, 167-174.
- [14] Tan, D., Gergle, D., Scupelli, P., and Pausch, R. With Similar Visual Angles, Larger Displays Improve Spatial Performance. *Proceedings of CHI 2003*, 217-224.
- [15] Vernier, F., Lesh, N., and Shen, C. Visualization Techniques for Circular Tabletop Interfaces. *Proceedings of AVI 2002*, 257-263.
- [16] Vogel, D. and Balakrishnan, R. Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. *Proceedings of UIST 2004*, 137-146.
- [17] Weiser, M. The Computer for the 21st Century. *Scientific American*, 265(30):94-104, September, 1991.
- [18] Wu, M. and Balakrishnan, R. Multi-Finger and Whole Hand Gestural Interaction Techniques for Multi-User Tabletop Displays. *Proceedings of UIST 2003*, 193-202.
- [19] National GeoSpatial Agency (NGA) (<http://www.geoint2004.com/index.cfm>).
- [20] NextFest (<http://www.nextfest.net>).