

T3: Rapid Prototyping of High-Resolution and Mixed-Presence Tabletop Applications

Philip Tuddenham and Peter Robinson
University of Cambridge Computer Laboratory
15 JJ Thomson Avenue, Cambridge CB3 0FD, UK
{firstname}.{lastname}@cl.cam.ac.uk

Abstract

Multi-person tabletop applications that require a high display resolution, such as collaborative web-browsing, are currently very difficult to create. Tabletop systems that support mixed-presence collaboration, where some collaborators are remote, are also hard to build. As a consequence, investigation of some important tabletop applications has been rather limited. In this paper, we present T3, a software toolkit that addresses these challenges. T3 allows researchers to rapidly create high-resolution multi-person tabletop applications for co-located or remote collaborators. It uses multiple projectors to create a single seamless high-resolution tabletop display, and allows multiple table-tops to be connected together to support mixed-presence collaboration. This engineering is hidden behind a simple, flexible API. T3 also supports existing user interface components, including buttons and spreadsheets, allowing the rapid creation of complex tabletop applications.

1. Introduction

In recent years, interactive tabletop interfaces have emerged as a key tool for co-located collaboration over digital artifacts. Yet, in spite of much promising research, there has been little investigation of tabletop interfaces to support the collaborative tasks for which people currently use their desktop computers, such as collaborative web browsing, spreadsheets and document review. These are compelling applications to which tabletop interfaces may bring significant benefits. However, this area remains largely unexplored because, with very few exceptions, the display resolution of today's tabletop interfaces is too low to support such applications.

A further area that remains largely unexplored is the extension of tabletop interfaces to remote groups of collaborators. In this mixed-presence setting, each remote group

would sit at its own tabletop. As shown in Figure 1, all the table-tops would then be linked together to provide a shared workspace for collaboration in which collaborators could interact with, position and orient digital artifacts. All the table-tops would show the same artifacts, along with remote embodiments (such as arm shadows) of the participants. Such a system could well offer remote collaborators some of the benefits of tabletop interaction, such as a greater awareness of each other's actions, and space to explore both personal and group work [13], both longstanding problems in conventional groupware. However, despite promising early results, investigation of these mixed-presence systems has been rather limited [16, 7, 4].

The significance of these two gaps in the research should not be underestimated. If we are truly to believe that these interfaces will be adopted then we must begin to explore the benefits that tabletop interfaces can offer over and above conventional physical table-tops.

Recent tabletop research has been fuelled by software toolkits that address the core engineering involved, allowing researchers to concentrate on interaction techniques and applications. As we shall show, the reason for these two gaps in the research is that these kinds of systems pose unique engineering challenges that cannot easily be solved using today's toolkits. From our own experiences, and from discussions at a workshop [17], we know that it is presently very difficult for researchers to overcome these issues.

In this paper we present T3, a software toolkit that we have implemented to address these problems. It uses multiple projectors to create a single seamless high-resolution tabletop display, and allows multiple table-tops to be connected together. T3 allows researchers to rapidly prototype high-resolution tabletop interfaces for applications such as collaborative spreadsheets and collaborative web-browsing, both for co-located and mixed-presence collaboration. It is freely available for academic research and will allow rapid exploration of this field. Furthermore, through our experiences, we have been able to investigate the challenges in creating these interfaces.

In the next section, we review recent work to identify the challenges in engineering these applications and establish design goals for our work. We then present T3, showing how it meets these goals, and outline a simple worked example. We illustrate its utility through five novel research projects and conclude by discussing its limitations.

2. Background and Design Goals

2.1. Tabletop Collaboration

A great number of projects have investigated various aspects of tabletop interfaces for co-located collaboration [e.g. 15]. Much of the work has been possible because of reusable software toolkits, notably DiamondSpin [14], and the more recent Buffer Framework [6], that handle the core engineering such as:

- Allowing collaborators to arbitrarily position and orient digital artifacts, and groups of digital artifacts.
- Allowing multiple users to interact concurrently.
- Supporting direct interaction using bare-hands or stylus.

DiamondSpin is particularly useful because it allows researchers to rapidly prototype complex tabletop applications by reusing existing Java Swing user interface components, such as buttons and file choosers. This feature is essential if we are to investigate complex applications such as collaborative spreadsheets or web-browsing, because it is not within the scope of a research project to engineer such applications from scratch.

The Buffer Framework demonstrates a novel architecture that provides responsive performance when hundreds of digital artifacts appear on the tabletop. Unfortunately, it does not allow the reuse of existing user interface components, and so creating new applications is difficult.

2.2. Higher-Resolution Displays

The key problem in supporting applications such as collaborative web-browsing or spreadsheets, is in creating a




Figure 1: Mixed-presence collaboration.

display surface with a sufficiently high resolution. In order to use the unique affordances of tabletop collaboration, we wish each web-page or spreadsheet to appear no larger than an ordinary sheet of paper, so that they can be passed between collaborators as one might with paper documents. However, to accomplish this, we need to be able to legibly display small text and user interface components, which is impossible on most of today's tabletop displays.

We aimed to display 12pt text legibly (i.e. text that appears the same size as 12pt text that comes out of the printer). We have found that this requires a resolution of at least 60dpi, so that a fairly modest 85cm \times 85cm table requires a 4 Megapixel display. By contrast, almost all today's tabletop displays provide at most 2.6 Megapixels using at most two projectors.

Higher-resolution projectors are extremely expensive and unsuitable. By far the easiest way to create a higher-resolution display using is to tile multiple projectors. For example, we have tiled 6 projectors to create a 4.7 Megapixel display using modest, inexpensive equipment. Such multi-projector designs are commonplace in the display walls used in visualisation research. However, tiling introduces further problems.

Firstly, the huge number of pixels and multiple graphics outputs can lead to unresponsive applications if a naïve rendering architecture is used. Secondly, it can be very difficult to align the projectors to the degree of precision required for a seamless display, and so these displays can suffer from small overlaps, mismatches and keystoning. Even using the precision mechanical alignment mountings that have been designed for this purpose, aligning 6 projectors to the required accuracy is time-consuming, requires careful engineering, and often relies on rear-projection, which precludes the use of some direct-touch and stylus technologies.

Large multi-projector display wall systems address both of these issues. Hardware-acceleration and selective updating speeds up the rendering process, using commodity graphics cards and interfaces such as OpenGL. These systems also automatically compensate for small projector misalignments by applying small adjustment transformations and blending masks to each frame before it is sent to the projectors [9, 18, 1]. These adjustments are usually calculated using a quick calibration procedure, avoiding the need for precise mechanical alignment. Unfortunately, these display wall toolkits are not designed to afford tabletop interaction or rapid prototyping; they typically do not, for example, support rotation of artifacts.

The DiamondSpin toolkit is unlikely to extend easily to a multi-projector tiled display. It does not yet use hardware-accelerated rendering, so performance would be limited, and there is not scope to add adjustment transformations or blending masks. By contrast, the Buffer Framework uses

OpenGL rendering, performed well on a 4 projector tiled display using mechanical alignment, and could easily be modified to support adjustment transformations and masks.

2.3. Mixed-Presence Collaboration

We also wish to investigate mixed-presence collaboration, whereby two geographically-separated tabletops are linked together to allow two remote groups to collaborate as though co-located around the same tabletop. Prior research in this area has investigated mixed-presence drawing surfaces [e.g. 16] and tangible interfaces [e.g. 2, 19].

There has, however, been little mixed-presence investigation of applications in which collaborators interact with, position and orient digital artifacts like spreadsheets or text documents. TIDL [7] and RemoteDT [4] both allow mixed-presence collaborators to position and interact with digital artifacts on large horizontal displays (much as GroupKit [10] does for remote collaborators on conventional desktop computer interfaces). However, none of these systems allow collaborators to reorient artifacts. Orientation serves several important roles in co-located tabletop collaboration, such as allowing transitions between personal and group work [13, 8], and it should not be overlooked when designing mixed-presence tabletop systems.

Creating a mixed presence system that allows participants to reorient artifacts is desirable but technically difficult: firstly because standard remote display protocols are not designed to handle artifact rotation; and also because the existing tabletop toolkits, DiamondSpin and Buffer Framework, cannot easily be extended to support mixed-presence collaboration. The very recent DigiTable work [3] supports mixed-presence tabletop collaboration, including orientation, remote arm shadows and existing Java Swing components but is not widely available and, to our knowledge, does not support multiple projectors, precluding high-resolution applications.

2.4. Design Goals

Having reviewed the literature, we now establish several design goals. The system should:

- provide abstractions to support the core tabletop interaction functionality provided by other tabletop toolkits, such as rotation of artifacts. Such requirements have been discussed in prior work [11, 14], and in Section 2.1 we have enumerated the most salient given the space available.
- allow creation of higher-resolution tabletops by supporting multiple projectors in a tiled array, using hardware-accelerated rendering for performance, and applying small transformations and masks to create the illusion of a seamless display.

	Diamond Spin	Buffer F'work	Display walls	TIDL/ RemoteDT	DigiTables	T3
Tabletop interaction (e.g. rotation, etc.)	✓	✓			✓	✓
Higher-resolution		✓	✓			✓
Two linked surfaces (mixed-presence)			Some	✓	✓	✓
Reuse existing UI components	✓		✓	✓	✓	✓

Table 1: Comparing tools and design goals.

- allow geographically-separated tabletops to be connected together to create a shared workspace for mixed-presence tabletop applications.
- allow researchers to rapidly create complex tabletop applications like spreadsheets by reusing existing user interface components.

Table 1 summarises our design goals and the limitations of existing tools.

3. The T3 Toolkit

Having established these design goals, we now provide a brief overview of the T3 architecture, followed by a description of each component, example applications, and a discussion of the implementation.

Figure 2 outlines the T3 architecture. Applications are created using a simple, well-documented Java API. The API is based around the notion of a single seamless large display, in which the application programmer can create rectangular tiles to represent interactive draggable digital artifacts such as spreadsheets or web-pages. The application programmer uses T3 to arbitrarily position and orient the tiles within a single large coordinate space.

Importantly, tiles can be filled with existing Java Swing user interface components, which function as expected without any extra code required on the part of the applications programmer. This includes buttons and text boxes, and even third-party components such as spreadsheets and web-browsers, allowing rapid creation of complex applications. T3 works behind the scenes to ensure that the application programmer is never required to consider the effects of tile rotation or scale, nor how to distribute the information to multiple tabletops, nor multi-projector blending techniques, nor simultaneous input event streams.

Each multi-projector tabletop display is controlled by a local computer running a T3 client. This client communicates with the application, receiving tile updates and sending back user input events. If desired, multiple clients can connect to the same application via the Internet to allow

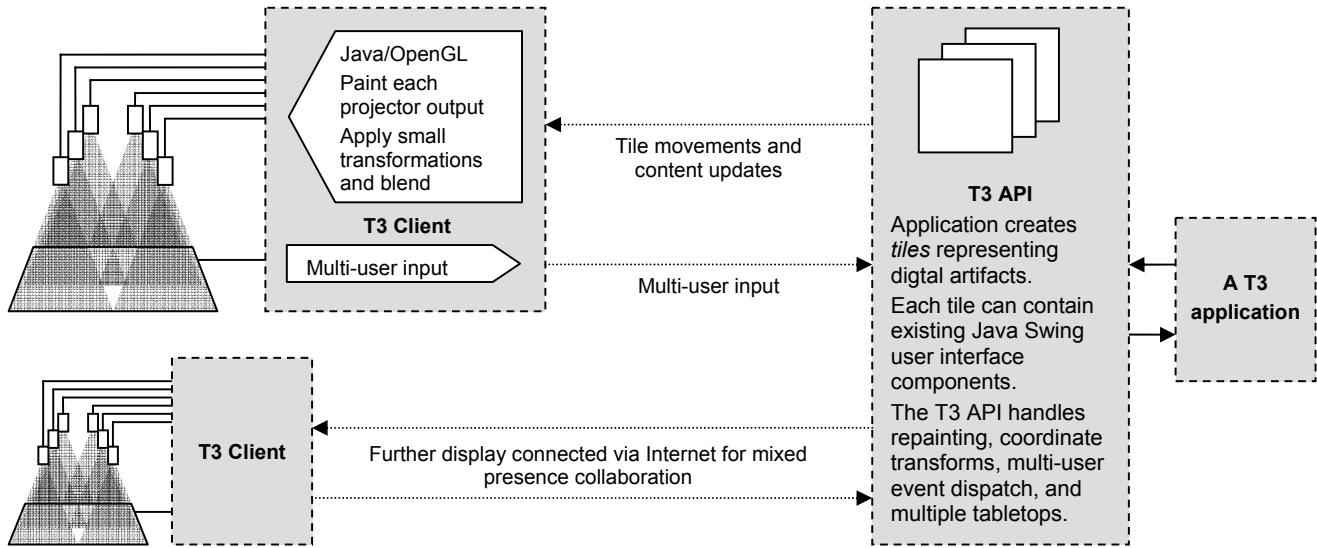


Figure 2: T3 system architecture to support multi-projector tabletops and mixed-presence collaboration.

mixed-presence collaboration, as proposed in the Introduction. The client automatically creates the correct images to send to each of the projectors in the multi-projector display. It positions and orients the tiles, applying small warps and masks to correct for projector misalignment and create the illusion of a single seamless display.

3.1. Physical Apparatus

T3 is easily configurable to support any number of projectors connected to a single PC. We have created three tabletop displays:

- Six projectors and a single 2.4GHz PC with three NVIDIA GeForce 7600 dual-head graphics cards to create a 4.7 Megapixel display.
- Four projectors and a single PC with two dual-head graphics cards to create a 5.7 Megapixel display.
- One projector to create a 0.8 Megapixel display.

All the components are available off-the-shelf, and neither multi-projector display cost more than \$13,000, including mountings. We have tested the toolkit using 3 brands of graphics cards. For multi-user input, we use standard graphics tablet styluses and Anoto streaming styluses, and the system could easily be extended to multi-touch surfaces.

3.2. API and Swing Applications

The application programmer creates tiles by instantiating T3's tile class. The programmer specifies how each tile should be positioned, rotated and scaled on the tabletop by specifying coordinates and dimensions in millimetres, and

an angle in radians. Collaborators move tiles around the display surface by dragging with their stylus using Rotate 'N' Translate [8]. The application programmer determines which tiles are draggable and can group tiles so that they are then dragged together, allowing creation of mobile container elements like Storage Bins [12].

Each tile functions as a Java Swing window in which existing user interface components can be used without any modification required, as we illustrate later with a worked example (Section 4). Multiple collaborators can interact simultaneously to drag these tiles and manipulate the components within. The vast majority of Swing components work as expected without any modification, though components that use popup windows currently require special attention.

This mechanism works by opening the necessary Swing windows off-screen. Swing repaint events are then trapped by the T3 toolkit which renders the windows into images that are sent to the T3 clients. Similarly, user input events received from the T3 clients are translated into Swing input events and dispatched to the appropriate Swing window. T3 uses geometric transforms to determine the tile immediately "underneath" the event on the tabletop and then transforms the event from tabletop coordinates (in millimetres) into Swing window coordinates, "undoing" the effects of tile rotation, translation and scaling, similar to Diamond-Spin's transformation engine.

T3 also provides an alternative API to support more customised applications that avoid the constraints of Java Swing. This allows implementation of more complex designs like Storage Bins [12], and both APIs can be mixed within the same application to produce, for example, a Storage Bins design that supports web browsing using a Java Swing web browser component. In this alternative API,

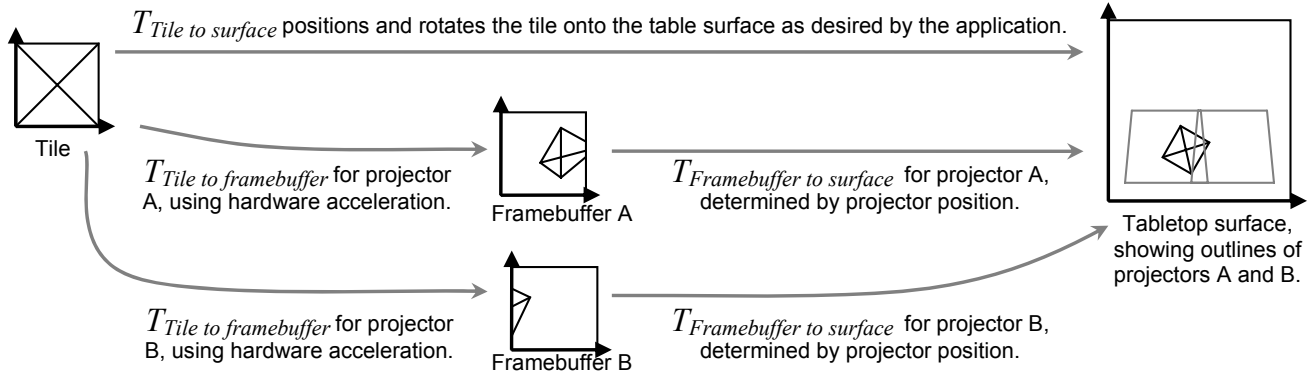


Figure 3: Transformations between coordinate spaces in the T3 client.

the application programmer simply overrides a paint routine and an input event processing routine for each tile.

T3 uses a multi-threaded architecture to handle the simultaneous input event streams from multiple users and multiple tabletops. This provides responsive performance and allows multiple collaborators to interact simultaneously to drag or manipulate tiles. The multi-threading and object locking are handled automatically by the toolkit.

Tiles are rectangular by default, but other shapes can be created by painting parts of the tile with transparent pixels. An ordering for the tiles determines occlusion when tiles overlap. T3 also provides a mechanism to smoothly animate groups of tiles between two specified positions and orientations, which can be used, for example, to zoom out to a thumbnail overview. A further mechanism allows the creation of translucent lines that join different tiles, which can be used, for example, to illustrate dependencies between artifacts.

3.3. Rendering and Performance

Each multi-projector display is controlled by a local T3 client, which receives tile information from the application, via the Internet if necessary, and performs the actual rendering.

The client uses OpenGL (via the JOGL library) to exploit hardware-accelerated rendering provided by modern graphics cards. Tile images are stored in the texture memory within each graphics card. The client then renders each frame for each projector by transforming tile images into the framebuffer. Figure 3 illustrates the transformations between the tile coordinate space, the frame buffer, and the display surface. The transformation, $T_{Tile\ to\ framebuffer}$ consists of two parts:

$$T_{Tile\ to\ framebuffer} = T_{Framebuffer\ to\ surface}^{-1} \cdot T_{Tile\ to\ surface}$$

The first part, $T_{Tile\ to\ surface}$, positions, scales and rotates the tile onto the display surface as desired by the ap-

plication. The second part, $T_{Framebuffer\ to\ surface}^{-1}$, is a small transformation that compensates for projector misalignment. It is obtained using a short calibration procedure, which calculates the relationship between points in the framebuffer and points on the display surface, and then performs a matrix inversion. As described previously, the misalignment problem is difficult to solve mechanically for high-resolution displays, and such software solutions are well understood in the display-wall community.

We have tested T3 using a range of applications (Section 4) using the 6-projector apparatus. In all cases, T3 performed responsively and achieved a frame rate of 60fps (the maximum rate supported by the projectors).

We have also tested T3 to the limits of its performance by filling the display with large numbers of draggable, partially overlapping, 100x100px images, all moving simultaneously along predetermined paths. Using the 6-projector apparatus, T3 achieved a frame rate of 40fps when animating 400 such images, and 30fps when animating 800 such images, and remained responsive throughout both tests. These results are similar to the reported performance of the Buffer Framework [6] in comparable tests, though precise comparisons are not possible.

3.4. Multiple Tabletops

The client and the application can run on the same computer to create a single high-resolution multi-projector tabletop interface. Alternatively, as described earlier, multiple clients can connect to the application via the Internet for mixed-presence tabletop collaboration. In this case, the clients and the server use a protocol to communicate tile content updates, tile movement and user input events.

This protocol is optimised so that the most frequent operations, such as dragging tiles, use low bandwidth and provide responsive performance. The system also uses a basic adaptive scheme to avoid overwhelming lower-performing clients. We have tested our system using both a high speed

```

public static void main(String[] a) throws Exception {

// Create a new tabletop, d.
PortfolioServer d = new PortfolioServer(
    new ServerSocket(2000),
    new ServerSocket(2001), false);

// create a new tile of 500*500 pixels on d
SwingFramePortfolio myTile = new SwingFramePortfolio(
    d, d.rootPortfolio, new RotateNTranslate(),
    500,500);

// add a JSpreadsheet and menu bar to the tile
myTile.getFrame().add( new JSpreadsheet(80,40) );
myTile.getFrame().setJMenuBar(menuBar);

// make the tile appear a physical size 200mm*200mm
myTile.setTileWidthAndHeightInPORT(200.0, 200.0);

// position the tile 300mm from the top and left of
// the tabletop, rotate it by 0.17 radians and set
// scale factor 1.00.
myTile.setPORTtoPPORT(300.0, 300.0, 0.17, 1.00);

// make the tile visible
myTile.setVisibleWhenParentVisible(true);
}

```




Figure 4: Complete sample code (top) using T3 to create an editable tabletop spreadsheet (bottom) that appears on the multi-projector tabletop, can be passed using Rotate 'N' Translate and can be used in mixed-presence collaboration.

LAN and a lower bandwidth wireless network, and it performed responsively in both cases.

In mixed-presence collaboration, we display remote embodiments on the tabletops in order to convey presence and to allow remote participants to gesture to each other. We currently use telepointer traces [5] that follow each participant's stylus. Traces are a starting point in our investigation of remote embodiments: they are easy to implement, allow rich gestures, and are robust to network jitter. They allow participants to convey shapes, routes and indicate groupings. However, they may not convey presence as well as alternative embodiments such as arm shadows [16].

4. T3 Applications

We now describe a range of projects that rely on T3 to explore new tabletop applications. The projects are all being undertaken by five students in our Laboratory.

Worked Example. Figure 4 illustrates a short program




Figure 5: Collaborative web-browsing. Pages appear small yet legible, can be passed around the table (top) and browsed in a tree (bottom).

that creates a tile and fills it with a third party Swing JSpreadsheet component. The result is a working tabletop spreadsheet application. The spreadsheet appears as a legible rectangular 20cm × 20cm tile on the table. It is sufficiently small that multiple spreadsheets can be viewed simultaneously by collaborators around the table, and can be passed between collaborators using the Rotate 'N' Translate technique. Columns and rows can be selected, and formulae can be entered into cells. Without requiring modification of this program, mixed presence collaborators can connect over the Internet and can similarly interact with the spreadsheet. Further development will allow participants to collaborate over multiple interdependent worksheets to perform different analyses of the same data set.

Collaborative Web-Browsing. A second project investigates finding and sharing information from the web (Figure 5). The high-resolution capability of T3 allows web pages to appear legible yet sufficiently small that multiple collaborators can each read their own pages and pass them around. T3 also allows us to reuse a third party Java Swing web browser component for rapid development; the basic tabletop web browsing application was implemented by a student in 60 lines of Java in around 1 hour. Collaborators can open web pages, follow links, and pass pages between each other.

Remote Document Review Meetings. Our third project investigates mixed-presence document review meetings, in which remote collaborators discuss and collaboratively an-




Figure 6: Remote review meetings using T3. Documents containing size 12pt text can be read (top-left), browsed (top-right), and used for remote collaboration (bottom).



Figure 7: Mixed-presence command and control interfaces using T3. Two co-located participants (left) collaborate with a remote participant (right).

notate draft text documents. Our interface allows multi-page text documents to appear rather like an open book on the tabletop (Figure 6). We use the high-resolution feature of T3 to project legible text at size 12pt. The interface allows remote collaborators to use styluses to pass documents to each other, to browse within a document using a thumbnail view, and to annotate pages. T3 allowed us to rapidly create the interface, which was implemented by a graduate student in 650 lines of Java in around three days.

Command and Control Interfaces. A further student project uses T3 to investigate map-based command and control tasks, comparing co-located collaboration to mixed-presence collaboration (Figure 7). T3 has allowed us to create a single interface and then switch easily between co-located collaboration and mixed-presence collaboration to conduct the study.

Large-Format Collaborative Programming Interfaces.




Figure 8: Interactive UML sequence diagrams with annotations using T3.

Students have used T3’s high-resolution capabilities to generate interactive tabletop UML diagrams annotated with source code (Figure 8).

5. Discussion and Conclusions

At the outset of this paper, we highlighted two important areas of tabletop research that remain largely unexplored: applications that require higher-resolution displays, such as collaborative web-browsing; and mixed-presence applications, whereby geographically-separated tabletops are linked together. We showed that such applications are particularly difficult for researchers to create, because they present engineering problems that are not addressed by existing tools. By integrating recent research from co-located tabletop interfaces, multi-projector display walls and distributed groupware, T3 goes a significant way towards addressing these problems.

T3 addresses the core engineering required to use multi-projector displays and to connect remote tabletops together. However, it nevertheless provides a simple, flexible API that allows researchers to rapidly prototype new tabletop applications, without having to consider the engineering realities of multi-projector transformations or distributed rendering. Existing Java Swing user interface components, including buttons and spreadsheets, can be used without modification, allowing researchers to rapidly create complex applications. We have shown that T3 achieves responsive performance in a variety of applications, even when animating hundreds of digital artifacts.

We believe that the combination of these features will enable novel tabletop research that has not previously been possible. For example, the combination of high-resolution, tabletop interaction and complex applications allows rapid prototyping of tabletop web browsing interfaces. Similarly, mixed-presence collaboration and tabletop interaction permits interfaces in which mixed-presence collaborators

can reorient artifacts, a behaviour which serves important awareness roles in co-located collaboration.

The ability to reuse existing Java Swing user interface components, such as buttons and spreadsheets, is not without its limitations. The components are designed only for a single user and so, for example, in our tabletop spreadsheets it is not possible for two users to simultaneously select two different columns within the same spreadsheet. Furthermore, we do not believe that applications designed for a single user at a desktop PC are by themselves sufficient to support tabletop collaboration. Nevertheless, the ability to use existing components will undoubtedly lead to rapid development of more complex tabletop applications, such as collaborative web-browsing. T3 also supports more radical application designs by providing an alternative API that is not constrained by Java Swing.

Tabletop research has advanced considerably over the last 8 years, and we would be foolish to think that the abstractions provided by T3 are by themselves sufficient to support the tabletop interaction techniques of the future. Nevertheless, we have illustrated the versatility of T3 with a wide range of applications, and we believe that it will be of utility to researchers wishing to explore these new areas. We believe that T3 could be used to implement the vast majority of today's tabletop applications, and we expect the abstractions to develop over time to reflect new research.

T3 is freely available to academic researchers, along with documentation and example applications, at <http://www.cl.cam.ac.uk/t3/>.

Acknowledgements

We gratefully acknowledge the support of Rowan Hill, Tom Matthews, Richard Russell and Stephen Williams, who are undertaking projects using T3; members of the Rainbow Group and the anonymous reviewers, for useful comments; and Thales Research and Technology (UK) and the EPSRC, who jointly funded this work.

References

- [1] Ashdown, M., Flagg, M., Sukthankar, R., and Rehg, J. A flexible projector-camera system for multi-planar displays. In *Proc. Conf. Computer Vision and Pattern Recognition (CVPR) 2004*, vol. 2, 165–172.
- [2] Brave, S., Ishii, H., and Dahley, A. Tangible Interfaces for Remote Collaboration and Communication. In *Proc. CSCW 1998*, 169–178.
- [3] Coldefy, F. and Louis-dit-Picard, S. DigiTable: an interactive multiuser table for collocated and remote collaboration enabling remote gesture visualization. In *Proc. PROCAMS 2007*.
- [4] Esenther, A., and Ryall, K., RemotedT: Support for Multi-Site Table Collaboration. In *Proc. Int. Conf. Collaboration Technologies (CollabTech)*, 2006.
- [5] Gutwin, C. and Penner, R. Improving interpretation of remote gestures with telepointer traces. In *Proc. CSCW 2002*, 49–57.
- [6] Isenberg, T., Miede, A., and Carpendale, S. A buffer framework for supporting responsive interaction in information visualization interfaces. In *Proc. Int. Conf. Creating, Connecting and Collaborating through Computing (C5'06)*, IEEE Computer Society (2006), 262–269.
- [7] Hutterer, P., Close, B. S., and Thomas, B. H. Supporting Mixed Presence Groupware in Tabletop Applications. In *Proc. TABLETOP'06*, 63–70.
- [8] Kruger, R., Carpendale, S., Scott, S.D., and Tang, A. Fluid integration of rotation and translation. In *Proc. CHI 2005*, 601–610.
- [9] Li, K., Chen, H., Chen, Y., Clark, D.W., Cook, P., Damianakis, S., Essl, G., Finkelstein, A., Funkhouser, T., Housel, T., Klein, A., Liu, Z., Praun, E., Samanta, R., Shedd, B., Singh, J.P., Tzanetakis, G., and Zheng, J. Building and using a scalable display wall system. *IEEE Comp. Graphics and Applications* 20, 4 (2000), 29–37.
- [10] Roseman, M. and Greenberg, S. Building Real Time Groupware with GroupKit, A Groupware Toolkit. *ACM Trans. Comp. Human Interaction* 3, 1 (1996), 66–106.
- [11] Scott, S., Grant, K., and Mandryk, R. System guidelines for co-located collaborative work on a tabletop display. In *Proc. ECSCW 2003*, 2003.
- [12] Scott, S.D., Carpendale, M.S.T., and Habelski, S. Storage bins: Mobile storage for collaborative tabletop displays. *IEEE Comp. Graphics and Applications* 25, 4 (2005), 58–65.
- [13] Scott, S.D., Carpendale, M.S.T., and Inkpen, K.M. Territoriality in collaborative tabletop workspaces. In *Proc. CSCW 2004*, 294–303.
- [14] Shen, C., Vernier, F.D., Forlines, C., and Ringel, M. Diamondspin: an extensible toolkit for around-the-table interaction. In *Proc. CHI 2004*, 167–174.
- [15] Shen, C., Ryall, K., Forlines, C., Esenther, A., Vernier, F.D., Everitt, K., Wu, M., Wigdor, D., Ringel Morris, M., Hancock, M., and Tse, E. Informing the Design of Direct-Touch Tabletops. *IEEE Comp. Graphics and Applications* 26, 5 (2006), 36–46.
- [16] Tang, A., Boyle, M., and Greenberg, S. Understanding and mitigating display and presence disparity in mixed presence groupware. *J. Res. and Practice in Inf. Tech.* 37, 2.
- [17] Terrenghi, L., May, R., Baudisch, P., MacKay, W., Paterno, F., Thomas, J., and Billingham, M. Information visualization and interaction techniques for collaboration across multiple displays. In *Ext. Abstr. CHI '06*, 1643–1646.
- [18] Wallace, G., Anshus, O.J., Bi, P., Chen, H., Chen, Y., Clark, D., Cook, P., Finkelstein, A., Funkhouser, T., Gupta, A., Hibbs, M., Li, K., Liu, Z., Samanta, R., Sukthankar, R., and Troyanskaya, O. Tools and applications for large-scale display walls. *IEEE Comp. Graphics and Applications* 25, 4 (2005), 24–33.
- [19] Wilson, A. D., and Robbins, D. C. PlayTogether: Playing Games across Multiple Interactive Tabletops. *IUI'07 workshop on Tangible Play*.