Tangible 3D Tabletops: Combining Tangible Tabletop Interaction and 3D Projection

Peter Dalsgaard and Kim Halskov
CAVI & PIT
Aarhus University
{dalsgaard,halskov}@cavi.dk

ABSTRACT
In this paper we present the tangible 3D tabletop and discuss the design potential of this novel interface. The tangible 3D tabletop combines tangible tabletop interaction with 3D projection in such a way that the tangible objects may be augmented with visual material corresponding to their physical shapes, positions, and orientation on the tabletop. In practice, this means that both the tabletop and the tangibles can serve as displays. We present the basic design principles for this interface, particularly concerning the interplay between 2D on the tabletop and 3D for the tangibles, and present examples of how this kind of interface might be used in the domain of maps and geolocalized data. We then discuss three central design considerations concerning 1) the combination and connection of content and functions of the tangibles and tabletop surface, 2) the use of tangibles as dynamic displays and input devices, and 3) the visual effects facilitated by the combination of the 2D tabletop surface and the 3D tangibles.

Author Keywords
Interaction design, projection, 3D, tabletops, tangible interaction, Augmented Reality.

ACM Classification Keywords
H.5.1 Multimedia Information Systems.

INTRODUCTION
Three areas of interface research that have garnered much attention in recent years are tangible interaction, tabletop interfaces, and 3D projection. In this paper, we present a novel interface, the tangible 3D tabletop, which draws upon and combines elements from these three areas. The tangible 3D tabletop thus combines a rear-projected tabletop interface with a 3D engine, which enables precise projection of content from multiple top-mounted projectors onto tangible objects placed on the table.

Since the tangible 3D tabletop is a novel type of interface, it may be approached from a number of research angles. In this paper we approach it from a design perspective and focus on the challenges as well as the potentials for interaction designers who wish to develop applications for the tangible 3D tabletop.

We start out from a presentation of related research in the two areas, which tangible 3D tabletops build upon: tangible tabletop interaction and 3D projection. This is followed by a brief description of the technical setup for tangible 3D tabletops.

Subsequently we outline a series of general design principles concerning the characteristics of, and interplay among the components of the system. Next we present a number of applications, which we have developed for the tangible 3D tabletop in order to explore the underlying design principles and potentials. Seven of these applications are demonstrated in the video accompanying this paper, and we will refer to them as follows: [see the accompanying video 1:11].

In the last part of the paper we outline three overall design considerations: 1) combining and connecting content and functions of tangibles and tabletop surface; 2) potential for employing dynamic content on tangibles that can simultaneously serve as displays and input devices; and 3) the particular visual effects afforded by the combination of the 2D tabletop surface and the 3D tangibles. On this basis, we outline potential avenues for future work, including potential use scenarios in a variety of domains.

RELATED WORK
Milgram and Kishino's reality/virtuality continuum offers a starting point for positioning the tangible 3D tabletop in relation to other academic contributions. The reality/virtuality continuum places the real, physical environment at one end of a spectrum, followed by augmented reality, augmented virtuality, and, finally, the virtual environment. With reference to the continuum, tangible 3D tabletops are positioned in the area of augmented reality, the core of which is the integration of digital information into the physical environment. This entails a broad collection of strategies for the integration of physical and digital elements. Tangible 3D tabletops build upon and combine three means of developing augmented

NordiCHI '12, October 14-17, 2012 Copenhagen, Denmark

reality, namely, tangible interaction, tabletop interfaces, and 3D augmented spaces, which we briefly introduce next.

Tabletop Interfaces and Tangible Interaction

Tabletop interfaces represent a rapidly growing area of interface technology, which was originally based on a horizontal projection surface integrated into a tabletop, in combination with tangible interaction devices, such as styluses or physical objects positioned on top of a projection surface [20]. The idea of tangible interaction was coined by Ishii and Ullmer [16] to describe interactive physical objects that can serve as input and output devices. Jordà and colleagues [18] trace the roots of tangible interaction even further back, to Shneiderman's [32] conception of direct manipulation. Tangible interaction eventually developed into a dedicated research domain, and has yielded commercial products such as Siftables [22]. Some of the early explorations of interactive tabletops include the Active Desk [7], which combines a sensor camera and a projector, and the Sensetable [26], which combines projection with electro-magnetic sensors embedded into the table. Because of the rapid development in smartphones and tablets, we have recently have seen a dramatic increase in the application of multi-touch interfaces, where several users can simultaneously interact directly with a flat surface, using their fingertips [31;39].

Williams et al [38] have explored multi-touch interfaces based on three-dimensional objects made from an acrylic material, which makes it possible to extend the touch-sensitive part of a flat surface into the third dimension, for instance, by placing a pyramid-shaped object on a regular table surface. A different strategy for enabling interaction above the 2D surface is based on layers of infrared light and the use of high-speed cameras, enabling the identification of the position and angle of a single finger [35]. Recently, Marquardt and his colleagues [21] have explored the continuum between direct touch and gesture interaction with tabletop installations, which they call 'the continuous interaction space'. With respect to three-dimensional aspects of tabletops, Hancock et al. [15] addresses various aspects of rendering 3D models onto a 2D surface, and discusses issues related to multiple viewing perspectives. Yoshida and his colleagues [40] have developed a setup based on 96 micro-LCD projectors enabling the creation of a 3D image that may be viewed through 120 degrees.

The most direct source of inspiration for the tangible 3D tabletop presented in this paper is the Reactable [5], which also relies on the combination of camera and projector. While the Reactable was developed for music performances, interactive tabletops have generally been explored in a number of other settings, including education [31], museums [9], gaming [37], marketing [24], and design [8]. The Reactable employs the custom-developed Reactivision software to camera-track the position and orientation of objects on a surface, through the use of so-called 'fiducial markers'. Baudisch, Becker, and Rudeck [4] have extended the use of fiducial markers into the third dimension. By using physical building block containing glass fibre bundles, their Lumino system supports the capacity to track cubes being stacked on top of one another', but does not – as we do – use 3D projection onto the tangible object.

Tangible Magic Lenses [33] extends a two dimensional tabletop into the space above the table by dividing the space into discrete parallel layers stacked upon each other. By moving a two dimensional plate of cardboard parallel to the table top surface the user gets access to a part of the current layer. Whereas Tangible Magic Lenses achieves the effect of displaying on two dimensional plate above a conventional tabletop having a projector above the table, SecondLight [17] achieves a similar effect by having two projection sources below the table surface in combination with a special kind of liquid crystal based screen material for the table top display. By integrating a prims into a physical object, for instance a cylinder, it become possible to display on the sides, but not the top of the cylinder.

With regard to input, tabletop interaction styles have a variety of forms, ranging from tangible interaction, to (multi-)touch to gestures above the surface enabled by a variety technologies, including visual markers, capacitive sensing, infrared light, and cameras.

In our work with the tangible 3D tabletop, we employ a combination of rear- and top-mounted projectors for displaying content on a tabletop surface and tangibles, while we employ a camera and fiducial markers on tangibles, to track the position of the tangibles.

3D Augmented Space

While tangible interaction and tabletop interfaces are fairly well established areas of research, 3D projection is a relatively new domain of research. 3D projection may be construed as type of augmented space, which, in turn, is a particular kind of augmented reality, in which part of physical reality is augmented by a display technology. One strategy is to use a semitransparent display, on which the visual content is aligned in such a way that it matches the physical space or object. Another strategy is to project visual content directly onto objects in physical space [6;30].

As a subset of augmented space, the basic principle of 3D augmented space is to start with a physical object, and then create a 3D model of the object. Using 3D software, the model may be modified, by changing one of the surfaces, for example, and subsequently it is projected onto the physical object, thereby augmenting the object. Raskar and his research team [28] were among the pioneers in the area of 2D and 3D projection, and have extensively explored the use of projectors to graphically texturize and animate physical objects, as well as to address technical issues, such as aligning the projection in relation to the physical object [2;6;28]. The basic technique of projecting on three-dimensional physical objects has been used for prototyping by Akaoka, Ginn and Vertegaal [1].

Piper, Ratti, and Ishii [2002] have developed a system, Illuminating Clay, a 3D tangible interface for landscape analysis. Their system incorporates a ceiling-mounted laser scanner, which captures the changing geometry of a clay landscape model in real time. In other words, they project onto a single, shape-changing object that is not related to a table top. pCube [34], a case also working with only a single object, is a cube-shaped display created by arranging five LED panels into a box shape, which for this particular shape provides a 3D augmented object without the use of projection technology.

In several cases, 3D projection has been used for artistic purposes. In this domain, the works of Pablo Valbuena are among the more prominent examples. His best known works comprise the series of installations entitled Augmented Sculptures [36], in which the conjunction of angular and clear-cut geometrical shapes and 3D projections enables the painting of an outline of the installation's edges, or the creation of the illusion of light sources moving across the installation's surfaces.

3D projected space has also been used for cultural heritage communication. For instance, by projecting visuals onto historical artefacts, facts and narratives associated with the artefacts may be integrated into the exhibition of the artefacts themselves [3]. In the area of architecture and design, 3D projection has been explored on a scale that ranges from small-scale, tangible objects to media architecture [11].

In our work, we have combined the basic technique of tracking the position of physical object on a 2D surface well-known from tangible tabletops with 3D projection, in order to create a tangible 3D tabletop, which we will present and discuss in the remainder of this paper.

THE TANGIBLE 3D TABLETOP

The tangible 3D tabletop developed by our research laboratory, CAVI [14], consists of a translucent table surface (80 cm × 107 cm) under which a projector (1) and a camera (2) are mounted, see Figure 1. Above the table, two additional projectors (4 + 5) are mounted. The projector beneath the table displays visuals on the table, while the projectors mounted around the table project content onto tangibles (3), which are fitted with fiducial markers beneath their bases.

The fiducial markers are tracked by the camera (2) connected to a computer, which, using the Reactivision software [19], identifies the position and rotation of each tangible object. This computer renders the image to be displayed by projector 1 onto the table surface, and sends the data on a bus to two separate computers, each of which uses the commercially available 3D game engine, UNITY, to render images projected onto the tangibles by the projectors mounted above the table (4 + 5). For a more detailed presentation of the technical set up, including calibration, see [11].

The current setup employs two HD projectors, but the infrastructure of our systems allows for extending the set up to more projectors. In practice, the resolution of the projection on the table surface is perceived as significantly higher than on the tangible objects, owing to the fact that pixels projected onto the tangibles must be stretched, to compensate for projection angles (see the examples on the accompanying video, for the difference in resolution). This is a circumstance designers must consider, when designing content for the tangible objects. One strategy is to use abstract graphics for the tangible objects, and show detailed content on the table surface. An advantage of the multi-projector setup with a dedicated project for the table surface is that users handling the tangible objects do not cast shadows on the table, when moving tangible objects. Perhaps the most significant quality of the tangible 3D tabletop is that it supports multiple users' views into a 3D world, without the need for a sweet spot, as is required in most other 3D projection setups.

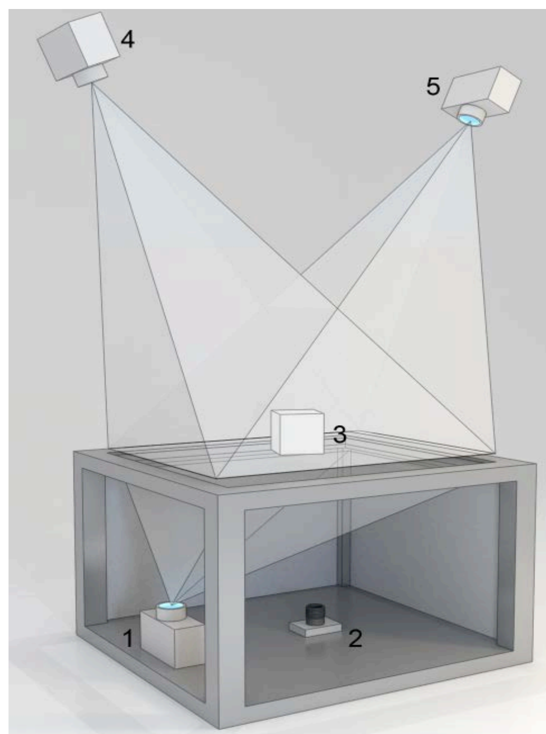

Figure 1. The main components of the tangible 3D tabletop.

RESEARCH APPROACH

Using the technical setup presented above, we have applied what can be characterized as a research through design approach [41] for exploring not only the technical feasibility and issues of the system, but also the design space and potentials that it affords. Since Tangible 3D Tabletops is new kind of interface, our research process have been of an explorative nature in which we have conducted a large number of design experiments in collaboration with a design firm over the course of six months, in addition to exploring a number of use cases. Throughout the process we have used a collaborative online system, the Process Reflection Tool (PRT) [12], for both documenting the design process and sharing reflections and thereby generating thorough data for ongoing and subsequent analysis. In addition to ongoing evaluations from our design team and the collaborating design firm, our explorative prototypes have been evaluated during a half day workshop with 20 twenty participants from industry, including architectural firms and design companies. Our analysis and findings concerning Tangible 3D Tabletops is a result of this process.

DESIGN PRINCIPLES

We will now outline the general features and design principles of the tangible 3D tabletop. Subsequently, we will illustrate how these principles may be implemented in the domain of maps and geodata.

Basic features and functionality of tangibles

The fiducial markers on the tangibles are tracked by the camera, and enable the system to identify each tangible, and its position and orientation. This serves two main purposes: first, it enables the system to match each tangible with a 3D model, and project unique content onto it; second, it allows users to employ the tangibles as input devices for the system.

The tangibles may have any shape, as long as their bases are large enough for a fiducial marker to be mounted. Since the Reactivision software can track a large number of fiducial markers simultaneously, the question of how many tangibles may be employed in practice is primarily limited by the physical size of the tangibles and the table. The tangibles may be combined and connected in a number of ways. As each tangible may be uniquely identified by its fiducial marker it is possible to develop sets of functions and behaviours for individual tangibles, as well as to determine how the system responds to combinations of tangibles. In conjunction with orientation and position data, this presents designers with a number of options. For example, tangibles may be linked, when placed in close proximity, or when they face each other; they may serve as 'screens' onto which content is projected, and yet others may serve as input devices or 'handles' for manipulating the content displayed on the table surface, or on other tangibles. Since the content projected onto a tangible can change, the connections and relationships among tangibles may be visualized not only on the tabletop, but also on the tangibles themselves [see the accompanying video 2:34]

Combining content on tangibles & surfaces

The fundamental aspect of the tangible 3D tabletop is that content can be displayed on both the surface and the tangibles. In our version of the tangible 3D tabletop, we have experimented with a number of combinations of different types of content on the tabletop and tangibles, which may be summarized by the following principles:

Displaying 3D content on a tangible integrated with 2D surface content: In this mode, the content displayed on the surface is matched with the position and shape of the tangible, and 3D content is projected onto it in such a way that it appears to wrap around the tangible (see Figure 2).

Displaying 3D content on a tangible that stands out from 2D surface content: In this mode, one type of 2D content is displayed on the table surface, while a different type of content is projected onto the tangible (see Figure 3).

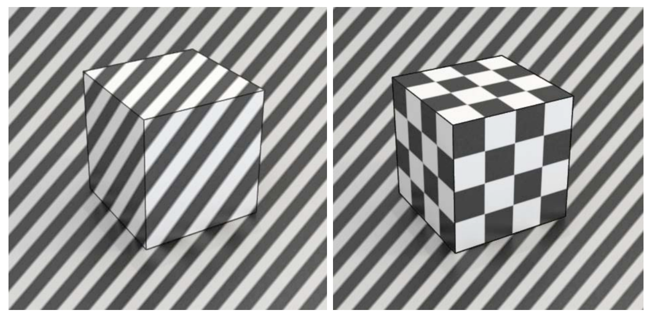

Figures 2 & 3. Displaying 3D content on a tangible that is either integrated with, or stands out from the content of the tabletop surface display.

Displaying 3D content on a flat tangible, integrated with 2D surface content: In this mode, the 2D content on the table is supplemented by content projected onto a flat surface of a tangible (see Figure 4).

Displaying 2D content on a flat tangible that stands out from 2D surface content: In this mode, the 2D content on the table is supplemented by content projected onto a flat surface of a tangible (see Figure 5).

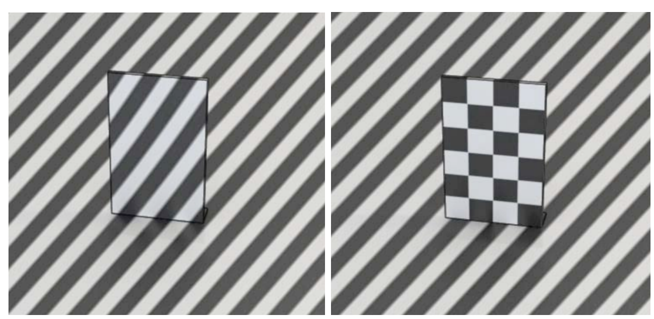

Figures 4 & 5. Displaying content on a tangible that is either integrated with or stands out from the content of the tabletop surface display.

We do not consider this to be an exhaustive list; rather, we regard it as set of basic principles describing the relationships between different types of content projected onto the tangibles and the tabletop surface, in the applications we have developed for the tangible 3D tabletop, so far. The four modes are not mutually exclusive, but may be combined in different ways. For example, one tangible may be designed to have 3D content integrated with the 2D surface content, while another tangible on the same table may have 2D content projected onto it. Table 1 gives an overview of the relations between integrated and separated 2D and 3D content:

	Integrated	Separate
3D	Figure 2	Figure 3
2D	Figure 4	Figure 5

Table 1. Integration of 2D and 3D content and tangibles

In the following section, we will show how these principles may be employed and combined in a specific domain, namely that of maps and geo-localized data visualization.

A USE CASE EXAMPLE: 3D MAPS AND GEODATA

The above-mentioned design principles are intentionally abstract, since they must encompass the general characteristics of the interface. In order to exemplify how the principles may be employed and combined in practice, we have developed a series of examples to show how a tangible 3D tabletop installation may be used for visualizing and exploring maps and geodata. A large part of our research concerns participatory design in urban spaces, and all the examples touch on how the tangible 3D tabletop may support exploration of urban data and joint design processes. The examples are the result of our collaboration with a design firm and have all been implemented as functional applications of the system in order to evaluate them on an ongoing basis. The examples follows the structure of the aforementioned design principles, and they are all represented in the video accompanying this paper.

Displaying 3D content on a tangible integrated with 2D surface content: A traditional flat map of a city is displayed on the table surface, and a number of tangibles have been modeled in the form of buildings in the city. When they are placed on the surface, the façade textures of the buildings are projected onto them. This can scaffold joint discussions about future projects, among citizens and urban planners [see Figure 6 and the accompanying video 0:06].

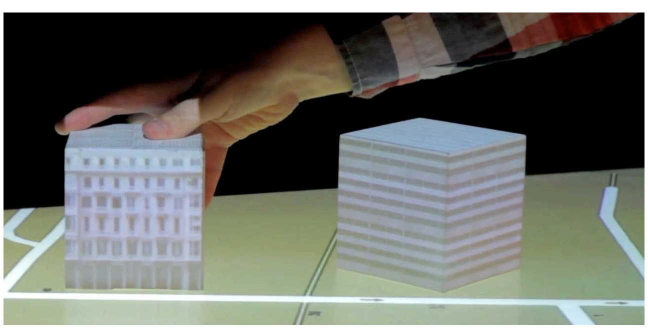

Figure 6. Displaying 3D building façades on tangibles placed on a map.

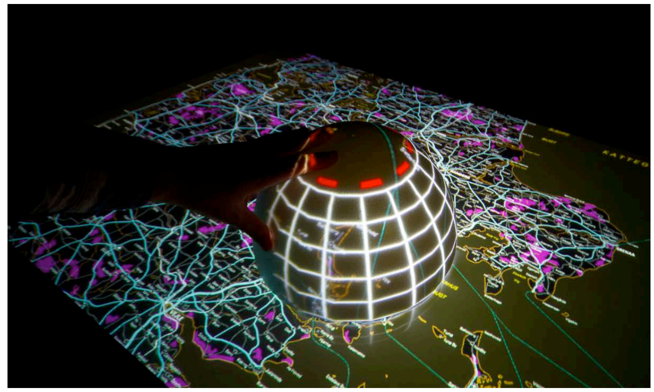

Figure 7. A tangible placed on a map may be used as a magnifying glass; by turning the tangible, users can zoom in and out of a selection on the map.

Another example of this principle in use is the employment of a tangible placed on a traditional flat map, as a combined filter and magnifying glass. When placed on a section of the map, it displays a magnified satellite image of the section. Rotating the tangible controls the zoom level [see Figure 7 and the accompanying video 2:19].

Displaying 3D content on tangible that stands out from 2D surface content: A traditional flat map is displayed on the table surface. A square tangible is available to the user. Each side of the tangible represents a specific type of data related to the region displayed on the map, for example, median income, median age, pollution levels, or traffic volume. When the tangible is placed on the table, the data for that particular location is displayed on the sides of the tangible, allowing the user to simultaneously explore how different statistical data are related, in different parts of the city. A number of tangibles from different categories may be available, allowing multiple users to explore the region at the same time, and discuss the interrelations between different statistics in various parts of the city [see Figure 8 and the accompanying video 0:45].

Displaying 3D content on a flat tangible integrated with 2D surface content: A number of white flat tangibles offers a window into a 3D model of the area represented on a map on the tabletop surface.

Figure 8. A set of tangibles may be employed simultaneously to display statistical data related to specific locations.

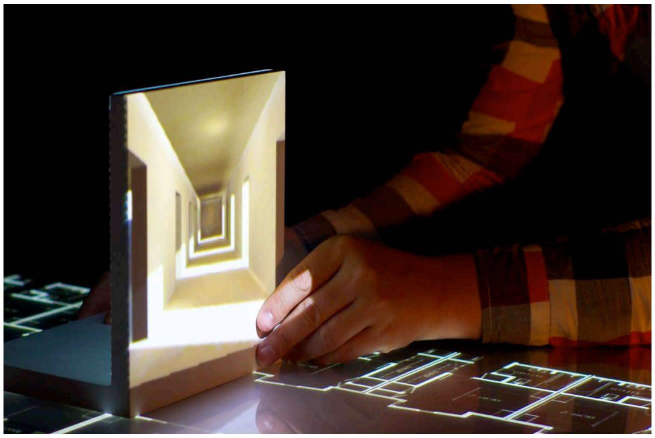

Figure 9. A blueprint of a building is displayed on the tabletop surface, and a 3D view of the building is projected onto a tangible.

When the tangible is moved, the perspective changes correspondingly. By placing multiple tangibles on the table, several users can explore different portions of a city or building at the same time. Different tangibles may represent different aspects, for example, one may display the selected area in bright daylight, while another shows a night-time view. In the example in the accompanying video, a blueprint of a planned building is displayed on the tabletop, while the tangible offers a 3D view into a model of the building, enabling users to experience its spatial characteristics [see Figure 9 and the accompanying video 1:19]. The example here is an extension of DeskRama (http://cat2.mit.edu/deskrama), which enables a view into a 3D model by positioning a 2D display perpendicular to printout of a plan drawing.

Displaying 2D content on a flat tangible that stands out from 2D surface content: Tangibles with similar flat shapes, as in the above-mentioned example, are placed on the table. When they are placed onto the flat map of the city, geo-tagged images from the specific locations are projected onto the tangibles. This may have several uses: a user can browse images from his own geo-tagged photo collection and show them to friends and family around the table; a family can explore an upcoming holiday destination by drawing upon geo-tagged data from web-based social image services such as Flickr; a user may explore a traditional Google Map in combination with Google Street View images projected onto the tangibles, and so forth. In the example in the video, users explore events hosted by the city's cultural institutions, such as art museums and concert halls [see Figure 10 and the accompanying video 1:41].

	Integrated	Separate
3D	Figures 6+7	Figure 8
2D	Figure 9	Figure 10

Table 2: Integration of 2D and 3D content and tangibles using the examples of 3D maps and geo-data visualization

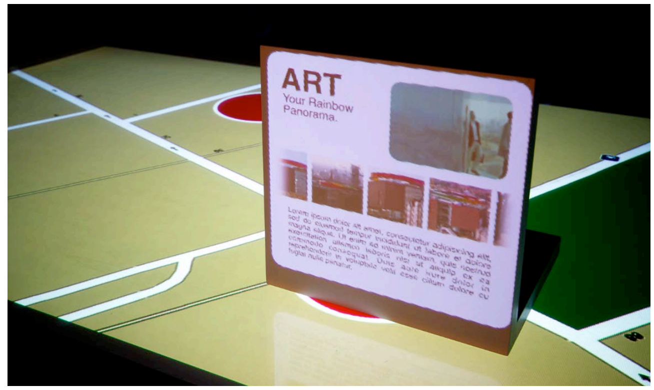

Figure 10. A tangible is used to display information about cultural institutions highlighted on map displayed on the tabletop surface.

Revisiting our categorization of content integration, the examples above can be plotted into Table 2.

DISCUSSION: DESIGN CONSIDERATIONS FOR TANGIBLE 3D TABLETOPS

Above, we presented the general principles of the tangible 3D tabletop interface, and demonstrated a series of applications in the domains of maps and geodata. We will now discuss three design considerations that we have found to be salient, based on our design experiments with the tangible 3D tabletop, so far. These considerations include 1) the combination and connection of content and functions of tangibles and tabletop surface; 2) the potential for employing dynamic content on tangibles, which simultaneously serve as displays and input devices; and 3) the specific visual effects afforded by the combination of 2D and 3D on tabletop surface and tangibles.

Combining and connecting content and functions of tangibles and tabletop surface

The examples of tangible 3D tabletop interaction in the domain of maps and geodata each represent quite specific uses of the tabletop display and the individual tangibles. However, one of the major strengths of the tangible 3D tabletop is that these elements may be combined in a number of ways, and the elements can change dynamically, depending on the context. A number of tangibles with different functions may be employed at the same time. Within the domain of maps and geodata, the tabletop might display a Google map, while one tangible might display geotagged images from Picasa and Flickr, or YouTube videos from specific locations on the map, another tangible might offer a 3D view of the city, based on models from Google Earth and Sketchup, a third tangible could present images from Google Streetview, while a fourth tangible could be used as a filter, to display a satellite image of a selected portion of the map. These interface elements may be construed as either *filters* (in the case of the Google Earth/Sketchup view and the satellite imagery) or *displays* (in the case of the Flickr, Picasa, YouTube, and Streetview images). Furthermore other tangibles could serve as more traditional *input components*; for example, one tangible could function as a slider to control zooming on the map, a second tangible could function as a knob for turning the map, a third tangible could switch between different views such as transit and terrain, while a fourth tangible could be used to move the map.

Since we can uniquely identify each interface component – whether a filter, display or input component – and we can determine its position and orientation, we can program specific behaviours, depending on the relationships between them. Furthermore, we can display the results of the interactions between these elements and the tabletop on each element, as well as on the tabletop display. In the accompanying video, we offer a simple example: three square tangibles and one circular tangible are available. When a square tangible is placed on the table, red, green, or blue is projected onto it. When all three tangibles are present on the table, a triangle is drawn on the table, with each tangible representing a corner. The triangle is filled with a solid colour, and when the tangibles are turned, the amount of red, green, and blue changes, and affects the colour of the triangle. Furthermore, the R/G/B ratio is displayed on the tabletop surface, beside each tangible. When the circular tangible is placed on the table, the square tangibles acquire extra features: they now also appear to project their colours onto the circular tangible. If two square tangibles are aimed at the same portion of the circular tangible, their two colours will bleed together on the curved surface of the circular tangible [see Figure 11 and the accompanying video 2:34].

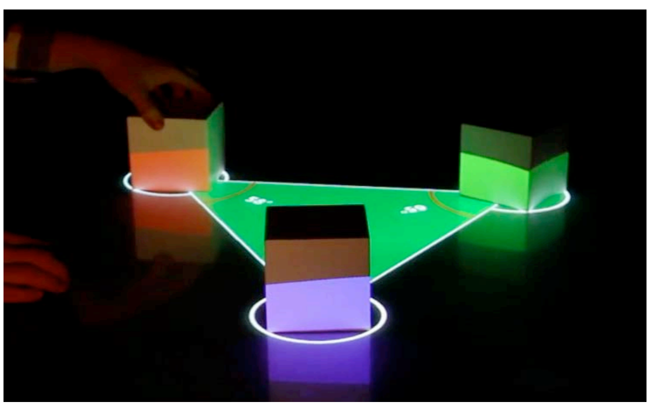

Figure 11. Three tangibles function as R/G/B blenders. The base color is projected onto the tangible, and users can change the hue of the triangle drawn within the tangibles by turning them.

Dynamic content on tangibles that are simultaneously displays and input devices

In addition to illustrating how interface components may be connected and combined, the apparently simple example above also highlights a second central design consideration, which is distinctive of tangible 3D tabletops: Designers must consider not only what information is displayed on the tabletop surface, but also what information should be displayed on the tangibles. In effect, each tangible functions as a dynamic display which can represent its own function in the system, as well as offer feedback when used. Additionally, the tabletop display may be employed to provide feedback about the tangibles (for instance a slowly pulsating circle displayed around a tangible may show the user that the display has detected the presence of the tangible). Compared to the traditional Reactable [19], one of the ways in which the tangible 3D tabletop stands out is in that each tangible may serve not just as an input device whose function affects the content on the tabletop display, but the tangible itself may serve as a dynamic display. In some respects, this resonates with the design rationale behind the Optimus Maximus keyboard (http://www.artlebedev.com/ everything/optimus/), on which a tiny display is embedded in every key, so that the functions of the keys may change according to the use context; however, on the tangible 3D tabletop, the input components may be positioned anywhere on the tabletop surface.

Our initial experiments indicate that users of the tangible 3D tabletop respond to dynamic content on tangibles, and generally accept and understand their dynamic nature. Although we have not carried out systematic tests to draw any definitive conclusions, we speculate that the physical shapes of tangibles may be employed to indicate particular uses, for example, in a given setup, organically shaped tangibles could represent input devices, circular objects could represent filters, and rectangular tangibles could represent displays. This identifies one of the general challenges we have yet to address: Owing to the novelty of the interface, the extent to which we may rely on existing interface heuristics and vocabularies remains to be seen. While some principles, such as consistency and feedback will probably also apply to tangible 3D tabletops, others are likely to need modification and extension. For instance, it would be extremely interesting to explore how well-established concepts such as affordances and constraints [25] can be employed, when dealing with tangibles that have a fixed physical shape, but dynamic digital content.

Visual effects and the combination of 2D and 3D on tangibles and tabletop surface

The content projected from the top-mounted projectors (4 and 5 in Figure 1) is created using 3D software (e.g. 3D Studio Max). Similarly, the content projected onto the flat surface may also be created using 3D software, although this is not the only option. By producing digital content with 3D software, designers can take advantage of a number of existing 3D techniques. For instance, designers may use different lighting angles to create realistic and natural looking shadows on the tabletop surface, corresponding to the physical object. The different angles of the projectors enable the designer to control precisely how shadows and highlights appear. This enables the designer to use moving light sources, which may be employed to both add to the naturalism of a setup as a whole, and to alter the perception of the tangibles. For instance, it may be used to visualize the sun's movement across the sky, as tangibles with building façades cast shadows onto terrain; it may visualize 'magic' light sources, which can shine through a solid physical object and strike the table surface in a natural way on the other side of the object; it may even be used to alter the perception of the physical shape of tangibles, by using false highlights and unnatural shadows to make the objects appear to be of a different size or shape. Similar techniques may be employed to highlight specific aspects of a tangible, or other parts of the interface. In this respect, visual cues can be used to guide users' interaction.

Another common technique afforded by 3D software is the use of particle systems to create visual effects such as fire or smoke, which extend in a natural way across the tangible objects and the table surface. In addition to projection particle systems onto the surface and tangibles, the source of the particle system may also be placed above the table surface, and in the space between the tangible objects, thereby creating realistic effects in terms of reflections on the table surface and the tangible objects. 3D software also offers large numbers of filters and animated texture maps, enabling visual transformations of the surfaces of the physical objects. Yet another technique that may be employed on the tangible 3D tabletop is the use of physics engines from 3D game environments, which, among other things, may be used to simulate collisions, friction, and momentum of digital content.

While we have only begun to explore the potential of these visual effects in our design experiments with the tangible 3D tabletop, related work in 3D projection offers insights into the further potential of such techniques. However, we have found that several of the well-known, existing techniques become even more effective when projected onto three-dimensional physical objects than when displayed on flat surfaces, because the presence of physical objects lends them an air of reality not afforded by flat displays.

There are a number of challenges pertaining to the combination of tangibles and tabletop. In the current setup, there is a slight time lag between the movement of a tangible object, and the recalculated image being displayed on the scene. This results in a lag in the image displayed on the tangible when it is moved. An interesting finding from our initial tests of the setup is that users seem to accept such a delay, and either move the tangible objects slowly, or move them fairly quickly to a new position, remove their hands, and wait while the image stabilizes, within a fraction of a second. We have not explored this finding in depth, but we speculate that users are more forgiving with regards to this type of delay because there is obviously no delay on the tangible itself, so they can tell exactly where the lagging projected image will eventually line up.

We have also faced a challenge that is commonplace for 3D projection in general, namely the alignment of 3D elements produced using the 3D imaging software, and the physical objects [11]. The geometry of the physical objects may be slightly different from representation in the digital 3D model, owing to imperfections in their production, and the lenses in projectors may distort the image being projected, which, taken together, may result in pixels that 'spill over' their designated positions. While it can be hard to avoid this problem entirely, it can be mitigated by avoiding projecting content close to the edges of tangibles and/or by avoiding content that requires pixel precision near the edges.

CONCLUSIONS AND FUTURE WORK

This paper presents the first version of the tangible 3D tabletop. In addition to outlining the technical setup of the interface and the general design principles, we have discussed three particularly salient design considerations that have emerged in our work so far, namely the combination and connection of content and functions of tangibles and tabletop surface, the use of tangibles as dynamic displays and input devices, and the visual effects afforded by the combination of 2D and 3D on tabletop surface and tangibles. While we are encouraged by the findings from this version of the interface, it is also evident that there are many aspects that remain to be addressed. In addition to addressing technical limitations like low resolution of imagery on tangibles, occlusion, and the slight lag in tracking tangible objects we consider the design of applications for specific contexts, studies of real life use, and further development of the interface to be the most important next step.

Figure 12. A tangible in the shape of a milk carton takes on the appearance of different products when placed on highlighted areas of the tabletop, and relevant product information appears on the tabletop.

Regarding the development of applications for the tangible 3D tabletop, in this paper we have focused on presenting applications that demonstrate the basic principles of the interface. For the sake of presentation, we have primarily chosen simple example applications from one domain, namely that of maps and geolocalized data. However, we envision (and have begun experiments into) numerous applications in other domains. In the domain of product presentation, we have developed an application for displaying a series of dairy products (see Figure 12 and the accompanying video 3:09).

Some of the questions that emerge concern which types of content are better displayed on the table, which are better displayed on the tangibles, and how best to exploit the interaction between tabletop and tangibles, when presenting content. In our experiments, we have often been surprised by the results, as it can be difficult to predict how well different types of well-known content and visuals from other interfaces will work on the tangible 3D tabletop, since we are not yet accustomed to considering the combination of interactive 2D and 3D content on flat and polymorphous forms. We plan on entering into partnerships with outside partners, to develop real life applications, to enable us to gather valid use data, and to further examine the benefits and drawbacks of employing this type of interface. As part of this work, we also aim to explore design principles and potential, which have not been addressed in this paper. Another obvious next step would be to conduct a systematic comparison of our tangible 3D tabletop with more traditional tabletops.

We have only just begun to understand the design space for this type of system, and we are very encouraged to explore the further potentials of it. For instance, we have primarily experimented with static and/or relatively simple surface content (e.g. blueprints and maps), in order to examine the basic potential of the interface. When we look at these experiments, the setups we have explored so far may be characterized primarily as *augmented surface* versions of the tangible 3D tabletop. By this, we mean that the content on the table at the centre of attention, and the tangibles primarily serve to add extra layers of information or functionality to the surface content. For instance, visuals displayed on the tangibles in Figure 8 are dependent on the fixed map on the surface. At the other end of the spectrum, we may conceive of a different sort of setup in which the tangibles are at the centre of attention and determine the content displayed on the table. Such a setup may be labeled an *augmented tangible* version of the tangible 3D tabletop. Sticking to the realm of maps and geo-data visualisation, one example would be to have tangibles with rich data added to which additional information could be displayed on the surface. Somewhere between these two polar opposites, we find what is arguably the most interesting combination of tangibles and surface, one in which there is a reciprocal relationship between the tangibles and the surface content. We have begun to explore such a setup, which may be labeled a *symbiotic* version of the tangible 3D tabletop. One example of a symbiotic setup is the R/G/B Blender experiment illustrated in Figure 11. Our experiments and findings thus far have encouraged us to explore such symbiotic setups and experiment with more dynamic forms of content displayed on the tabletop surface in our future work.

ACKNOWLEDGMENTS

We would like to thank our colleagues at CAVI (Rolf Bagge, Peter Friis, Janus Bager Kristensen, Morten Lerivg, Jonas Petersen,) who have been deeply involved in developing the software and prototypes discussed in this paper. Moreover we would like to thank our industry collaborators Andreas Lykke-Olesen and Rune Nielsen from the design firm Kollision. This research has been funded by Center for Digital Urban Living (the Danish Council for Strategic Research, grant 09-063245) and by Aarhus University's interdisciplinary research center Participatory IT, PIT.

REFERENCES

1. Akaoka, E., Ginn, T. & Vertegaal, R. 2010. DisplayObjects: prototyping functional physical interfaces on 3d styrofoam, paper or cardboard models TEI '10. ACM, 49-56.
2. Bandyopadhyay, D., Raskar, R. and Fuchs, H. 2001. Dynamic Shader Lamps: Painting on Movable Objects. ACM International Symposium on Augmented Reality, ACM Press.
3. Basballe, D. and Halskov, K. 2010. Projections on museum exhibits – engaging visitors in the museum setting. Proc. Of OZCHI 2010, 80-87.
4. Baudisch, P. Becker, T. and Frederik Rudeck. 2010. Lumino: tangible building blocks based on glass fiber bundles. Proc. of SIGGRAPH '10, ACM Press.
5. Bencina, R., Kaltenbrunner, M. and Jordà, S. 2005. Improved Topological Fiducial Tracking in the reacTIVision System. Proc. IEEE International Workshop on Projector-Camera Systems.
6. Bimber, O. and Raskar, R. 2004. Spatial Augmented Reality. Merging Real and Virtual Worlds, A.K. Peters Publishing.
7. Buxton, W. 1997. Living in Augmented Reality: Ubiquitous Media and Reactive Environments. Video Mediated Communication, K. Finn, A. Sellen and S. Wilber (Eds.). Erlbaum, Hillsdale, NJ, USA, 363-384.
8. Clifton, P., Mazalek, A., Sanford, J., Rabola, C., Lee, S. and Powell, N. 2011. SketchTop: design collaboration on a multi-touch tabletop. Proc. of TEI '11, ACM Press, 333-336.
9. Correia, N., Mota, T., Nóbrega, R., Silva, L. and Almeida, A. 2010. A multi-touch tabletop for robust multimedia interaction in museums. Proc of ITS '10, ACM, 117-120.
10. Dalsgaard, P. and Halskov, K. 2010. Designing Urban Media Façades: Cases and Challenges. Proc of CHI '10, ACM, 2277-2286.
11. Dalsgaard, P. and Halskov, K. 2011. 3D Projection on Physical Objects: Design Insights from Five Real Life Cases. Proc of CHI '11, ACM.

12. Dalsgaard, P. & Halskov, K. 2012. "Reflective Design Documentation". In Proc. of DIS 2012: ACM conference on Designing interactive systems, ACM,.

13. Halskov, K. and Dalsgaard, P. 2011 Using 3D projection to bring a statue to life. Interactions.

14. Halskov, K. 2011. CAVI - An interaction design research lab. interactions 18(4), 92-95.

15. Hancock, M. Nacenta, M., Gutwin, C., and Carpendale, S. 2009. The effects of changing projection geometry on the interpretation of 3D orientation on tabletops. Proc. Of ITS '09, ACM Press, 157-164.

16. Ishii, H. and Ullmer, B. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. Proc. Of CHI '97, ACM, 234-241.

17. Villar, N., Butler, A. Westhues, J. 2008. Going beyond the display: a surface technology with an electronically switchable diffuser. UIST '08. ACM, New York, NY, USA, 269-278.

18. Jordà, S., Julia, C.F. and Gallardo, D. 2010. Interactive surfaces and tangibles. XRDS 16, 4, 21-28.

19. Kaltenbrunner, M. and Bencina, R. 2007. reacTIVision: a computer-vision framework for table-based tangible interaction. Proc. Of TEI '07, ACM, 69-74.

20. Lucchi, A., Jermann, P., Zufferey, G. & and Dillenbourg, P. 2010. An empirical evaluation of touch and tangible interfaces for tabletop displays. Proc. of TEI '10, ACM.

21. Marquardt, N., Jota, R., Greenberg, S., and Jorge, J.A. 2011.The Continuous Interaction Space: Interaction Techniques Unifying Touch and Gesture on and above a Digital Surface. Proc. Interact 2011, 461-476.

22. Merrill, D., Kalanithi, J. and Maes, P. 2007. Siftables: towards sensor network user interfaces. Proc. of TEI 2007, ACM Press, 75-78.

23. Milgram, P. and Kishino, F.A. 1994. Taxonomy of Mixed Reality Visual Display. IEICE Transactions on Information Systems 12 , 1321-1329.

24. Nielsen, R., Fritsch, J., Halskov, K. and Brynskov, M. 2009. Out of the Box – Exploring the Richness of Children's Use of an Interactive Table. IDC .

25. Norman, D.A. 1988. The design of everyday things, Basic Books, New York, NY, USA.

26. Patten, J., Ishii, H., Hines, J. and Pangaro, G. 2001. Sensetable: A wireless object tracking platform for tangible user interfaces. CHI '01: Proc. SIGCHI Conference on Human Factors in Computing Systems. , 253-260.

27. Piper, B, Ratti, C. & Ishii, H. 2002. Illuminating clay: a 3-D tangible interface for landscape analysis. In Proc. of CHI '02. ACM, New York, NY, USA, 355-362.

28. Raskar, R., Welch, G., Low, K-L. and Bandyopadhyay, D. 2001. Shader Lamps: Animating Real Objects With Image-Based Illumination. Proc. of the 12th Eurographics Workshop on Rendering Techniques, 89-102.

29. Reisman, J.L., Davidson, P.L. and Han, J.Y. 2009. Generalizing multi-touch direct manipulation. SIGGRAPH.

30. Saakes, D. 2010. Shape Does Matter - designing materials in products. Lulu Publishers.

31. Shaer, O., Strait, M., Valdes, C., Feng, T., Lintz, M. and Wang, H. 2011. Enhancing genomic learning through tabletop interaction. Proc. of CHI '11, ACM.

32. Shneiderman, B. 1983. Direct Manipulation: A step beyond programming languages. IEEE Computer, 16, 57-69.

33. Spindler, M., Martsch, M. & Dachselt, R. 2012. Going beyond the surface: studying multi-layer interaction above the tabletop. CHI 2012. ACM, New York, NY, USA, 1277-1286.

34. Stavness, I., Lam, B. & Fels, S. 2010. pCubee: a perspective-corrected handheld cubic display. CHI '10. ACM, New York, NY, USA, 1381-1390.

35. Takeoka, Y., Miyaki, T. and Rekimoto, J. 2010. Z-touch: an infrastructure for 3d gesture interaction in the proximity of tabletop surfaces. In Proc. Of ACM ITS '10.

36. Valbuena, P. 2007. Augmented Sculptures. http://www.pablovalbuena.com/p05.htm .

37. Velsen, M., Williams, J. and Verhulsdonck, G. 2009. Table-Top Gaming Narratology for Digital Interactive Storytelling. Proc. of ICIDS '09, Springer-Verlag, Berlin, Heidelberg, 109-120.

38. Williams, C., Yang, X. D., Partridge, G., Millar-Usiskin, J., Major, A. and Irani, P. 2011. TZee: exploiting the lighting properties of multi-touch tabletops for tangible 3d interactions. Proc. of CHI '11. ACM Press 1363-1372.

39. Wu, W. and Balakrishnan, R.R. 2003. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. Proc. of UIST '03. ACM Press, 193-202.

40. Yoshida. S., Yano, S. and Ando, H. 2010. Implementation of a tabletop 3D display based on light field reproduction. Proc of SIGGRAPH '10, ACM Press, Article 61. 1 page.

41. Zimmerman, J., Forlizzi, J. and Evenson, S. 2007. Research through design as a method for interaction design research in HCI. Proc. of CHI '07, ACM Press, 493-502