

Exploring Physical Information Cloth on a Multitouch Table

Kimberly Mikulecky¹, Mark Hancock², John Brosz¹, Sheelagh Carpendale¹

¹University of Calgary, Canada
Department of Computer Science
{kmmikule, jdlbrosz, sheelagh}@ucalgary.ca

²University of Waterloo, Canada
Department of Management Sciences
mark.hancock@uwaterloo.ca

ABSTRACT

We expand multitouch tabletop information exploration by placing 2D information on a physically-based cloth in a shallow 3D viewing environment. Instead of offering 2D information on a rigid window or screen, we place our information on a soft flexible cloth that can be draped, pulled, stretched, and folded with multiple fingers and hands, supporting any number of information views. Combining our multitouch flexible information cloth with simple manipulable objects provides a physically-based information viewing environment that offers similar advantages to complex detail-in-context viewing. Previous detail-in-context views can be re-created by draping cloth over virtual objects in this physics simulation, thereby approximating many of the existing techniques by providing zoomed-in information in the context of zoomed-out information. These detail-in-context views are approximated because, rather than use distortion, the draped cloth naturally drapes and folds showing magnified regions within a physically understandable context. In addition, the information cloth remains flexibly responsive, allowing one to tweak, unfold, and smooth out regions as desired.

ACM Classification: H.5.2 [Information interfaces and presentation]: Interaction Styles.

General terms: Design, Human Factors

Keywords: Physics, cloth, multi-touch, tabletop display, simulation, information visualization, detail in context

INTRODUCTION

The increasing variety of multitouch technology [9, 13, 17] is allowing people to experience computing in a new way. Instead of a single point of contact, many fingers and hands can be used to manipulate virtual artefacts simultaneously. Perhaps because these devices allow one to directly touch the virtual information, or perhaps because more of one's physical body is in use, this technology encourages a more physical engagement with the digital world. This technology has led to many interaction techniques that are physically-based [19, 26] or that make use of a physics engine directly [1, 16, 31].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference ITS'11, November 13–16, Kobe, Japan.
Copyright 2011 ACM 978-1-4503-0871-7/11/11...\$10.00.




Figure 1: Simply draping a virtual information cloth containing a map over a virtual object approximates a detail-in-context lens in a physical way.

On the other hand, information exploration techniques that have been made possible by digital technology often rely on abstractions such as graphs, node-link diagrams, or other 2D representations that we may not think of as being physical. Indeed, applications for exploring information on multitouch devices often opt for a less physically-based mapping [11, 15, 18]. However, in the physical world, 2D information is often represented on paper or other malleable surfaces, and these surfaces have a physical component that can be useful for information exploration. For example, when exploring a physical map, people find creative ways of viewing the interesting information when the space is limited, such as rolling its edges when the table is too small or folding up the uninteresting parts when in a car. When knitting or sewing, people frequently stretch out and bring closer the part of the material that they are working with, letting the remainder fall out of the way or over a table's edge.

This malleability of information has been a large motivating factor in the detail-in-context work in the domain of information visualization [6]. However, the use of malleable surfaces has mostly been to develop the mathematics behind the techniques; there has been much less attention on the interactions possible. For example, Forlines et al. developed a technique to interact directly with such lenses [11], however they used multitouch for lens placement rather than to control the shape of the distortion.

In this paper, we introduce *information cloth*—virtual cloth containing 2D information that can be manipulated using multitouch interaction and reacts to objects using a physics




Figure 2: This image demonstrates two people cooperating to manipulate the cloth and highlight regions of interest.

engine (Figures 1 & 2). Presenting information in this way allows people to leverage their knowledge of physical cloth behaviour to creatively explore 2D virtual information. We also introduce rigid, manipulable artefacts into this environment that the cloth can “drape” over. Specifically, we demonstrate that the use of a physically-based cloth metaphor to represent 2D information can provide detail-in-context when exploring information on a multitouch table.

The paper is organized as follows. We first discuss the related work, and then introduce the concept of the 2D information cloth. Next we describe examples of how this information cloth can be used both to approximate existing detail-in-context techniques, and to explore information in new and interesting ways. We then provide implementation details, including a description of our object and cloth manipulation.

RELATED WORK

This research draws upon and contributes to multitouch interactions in general and physics-based interaction in particular. Similarly, it relates to and extends information interaction techniques such as detail-in-context approaches and lens based interactions. Thus we focus our discussion of the related literature on these two aspects.

Multitouch Interaction

A significant amount of work has explored the movement, rotation, and scale of virtual artefacts using one’s hands on a multitouch digital surface [5, 14, 19, 23, 27, 30, 31, 32]. Movement and rotation in 2D can be accomplished with a single finger [19, 27], or multiple fingers can be used to move, rotate, and scale [21, 30]. Many fingers can also be used to move, rotate, and scale a virtual 3D artefact [14, 23]. Other approaches attempt to manipulate objects by more closely replicating actions from the physical world, such as the ability to knock over and push objects around [5, 31], or the ability to bend an edge to put objects underneath [5].

Our work extends the notion that virtual objects can react in a physical way to the touches on the surface of a table.

While most of this work explored the movement, rotation, and scaling of rigid virtual bodies, this paper describes the result of presenting 2D information on a 3D virtual cloth-like surface. Some of the previous work has mentioned the ability to control cloth-like 3D virtual objects [5, 16, 31], however we make cloth interaction our primary focus and look more specifically at presenting information on this malleable surface. Thus we address new issues specific to designing interaction with an information cloth.

Detail-in-Context Information Exploration Techniques

In this paper, we show how information cloth can provide detail in context, and so we provide a brief overview of the most closely related detail-in-context approaches.

Since the concept of detail-in-context was introduced close to thirty years ago, a great variety of possible variations have been suggested. Spence and Apperley [28] suggested coping with simultaneous need for context and magnification by providing local magnification and compression algorithms. Furnas [12] suggested that selective filtering of items of lesser importance could create space for items of current interest, based on a degree-of-interest (DOI) function. These ideas provided two basic approaches: one spatially-based where areas of the display were compressed to make room for others; and one content-based in that some content was filtered to make room for others.

Approximately eight years later, Sarkar and Brown [25] demonstrated how Furnas’ DOI could be interpreted spatially creating Graphical Fisheyes, which provide local magnification set in a context which was compressed increasingly in relationship to the distance from the region of interest. Variations on this theme explored the idea of thinking of the region of local magnification as a lens, posing different visual and algorithmic variations. The better known of these include Perspective Wall [20], Document Lens [24], and Gaussian lenses (also known as Pliable Surfaces [8]). Subsequently Carpendale and Montagnese [6] provided The Elastic Presentation Framework (EPF) that gave an algorithmic basis for unifying these approaches, making it possible to include any number of them in a given interface.

DTLens [11] introduced local magnification lenses, algorithmically based on EPF, to multitouch tables. More recently Sigma Lenses [22] combined the concept of Magic Lenses [2] with EPF to create a greater variety of integrated visual lens types. Mélange [10], also looked at the EPF basis but considered the interaction metaphor of making region smaller or pushing regions away from view. This made it possible to juxtapose different regions of information that might initially be quite far apart.

Our work extends previous literature by demonstrating not only that previous detail-in-context techniques can be approximated but also that using a physically-based information cloth supports additional interaction variations. In addition, using multiple fingers and hands to manipulate this cloth has the potential to leverage some of people’s prior knowledge of how physical cloth-like material (cloth, paper, cardboard, etc.) interacts with physical artefacts. Instead of controlling lenses through the somewhat abstract parameters

involved in many of these systems, people can arrange and adjust virtual objects. The physical simulation of how this cloth interacts with other virtual artefacts approximates and extends earlier detail-in-context techniques.

THE CONCEPT: THE 2D INFORMATION CLOTH

Our main contribution is the marriage of physics-based cloth interaction on a multitouch table with information exploration. While there exists a multitude of possibilities within this combination, we have developed interactions with information cloth with the following design considerations:

Replicate existing detail-in-context functionality: In order to demonstrate the power of this idea, we show that this method is capable of replicating most of the existing detail-in-context functionality. Specifically, by rendering 2D information visualizations on a virtual cloth, it is possible to approximate the effect of using lenses such as Graphical Fisheyes [25], Perspective Wall [20] Document Lens [24], Gaussian Lenses [7], Manhattan lenses [6], DTLens [11] and Mélange [10]. We chose this direction because the use of cloth provides a physically understandable mechanism for recreating these somewhat abstract views. Our intention is to reproduce the visual magnification effects rather than the mathematics behind the techniques. By providing an approximate recreation of the effect of each distortion, while maintaining the visual and interactive aesthetic of cloth, the accompanying distortions are understandable in the same way that folded and draped cloth is understandable.

Manipulate virtual artefacts instead of parameters: Our design uses the concept that a person need only manipulate (move, rotate, and resize) virtual artefacts. This direct interaction approach is in contrast to many existing detail-in-context solutions, which (to create the desired information view) requires the adjustment of several parameters, such as scale, distance metric, and drop-off of a lens. E.g., with DTLens [11] one can interact directly with touch to move the lens' location. However, on-screen widgets are used to adjust many other of DTLens' viewing parameters such as degree of magnification, amount of the information space that is distorted, and the degree of the distortion. These widgets occlude information and successful adjustments often require deep understanding of the algorithmic parameters. In contrast, we re-create these viewing effects through the physical manipulation of an information cloth and by draping the cloth over a few simple objects. By focusing on this physical behaviour, we can leverage a person's knowledge of how cloth and objects behave in the physical world.

Provide multitouch manipulation of the cloth: One of the features of malleable surfaces in the physical world, such as cloth, is that a person can control many parts of the surface independently and simultaneously. For instance, a person can stretch the centre of a cloth tight with fingers from both hands, while letting the edges crumple up, or a person can pinch two points together so they can see them in the same space. Our design allows multiple fingers on both hands to control any part of the cloth. Our intention is to simulate as closely as possible how a physical cloth would react.

Provide multitouch manipulation of objects underneath the cloth: One of the benefits of cloth-like surfaces is that they can "drape" over other objects. For example, when quilting, the ability to drape allows the quilter to bring a part of a large material closer for concentrated work, while maintaining the context of where that particular part belongs in the bigger pattern. When moving to the digital world, it becomes possible to manipulate an object that the material is draped over with one's hands, despite it being covered. By enabling people to move the objects underneath the cloth, as well as the cloth itself, our system can provide the choice between moving the data itself versus moving the area(s) of interest.

EXPLORING INFORMATION ON VIRTUAL CLOTH

While there are many ways to manipulate information placed on a virtual cloth, initially investigating possible uses and advantages of an information cloth in a well researched area such as detail-in-context lenses is a good starting point. Thus, we reproduce popular lens effects with information cloth. The idea relates to that of Perspective Wall [20] and EPF [6] lenses that use the third dimension and perspective projection to create magnification. We drape our information cloth over simple objects that have appropriate height in the 3D virtual world. This produces relative magnification of the area of cloth on top of the object (Fig. 1). The information cloth combines lens effects with multitouch object interaction to offer direct manipulation of most lens properties.

By using differently shaped objects under the cloth, we can produce different lenses. Through touching and dragging, the objects can be moved while under the cloth. This movement allows exploration of the cloth's information, emphasizing or magnifying different areas as the object's position is changed. During this movement the cloth becomes slightly transparent, allowing objects to be easily seen and adjusted. Subsequently when the interaction pauses, the cloth is draped once more over the newly arranged objects.

Detail-in-Context with Information Cloth

To illustrate how our virtual cloth approach can simulate detail-in-context views, we step through the better-known lenses showing the structure of the lens via a simple grid and placing this image next to one where the structure of the lens is provided by an object to show comparable virtual cloth approximation of this lens using a checked tablecloth to show the structure. Then, with a somewhat larger image, we show the virtual cloth detail-in-context view with information on the cloth¹. Specifically, we demonstrate the potential for generality by simulating the following lenses: Perspective Wall [20], Graphical Fisheyes [25], Document Lens [24], Constrained Gaussian Lenses [7], Manhattan Lenses [6], and the more recent Mélange [10].

Perspective Wall The Perspective Wall [20] is a technique that provides detailed information in the central region of a

¹The maps we use as information in these examples are Chataways's Map of Greater Winnipeg Enlarged and Revised Edition 1917, and Map of Western Canada Manitoba, Alberta, Assiniboia, Saskatchewan and Part of British Columbia, Showing System of Land Survey and Lines of the Canadian Pacific Railway Company (1893). Both images were obtained through a Creative Commons attribution license (<http://www.flickr.com/photos/manitobamap>).




Figure 3: Perspective Wall




Figure 4: Graphical Fisheye




Figure 5: Document Lens

display and contextual information to the right and left of this central area by bending a 2D layout over a 3D “wall”. Indeed, when working out their ideas for Bifocal Views [28] (a precursor of Perspective Wall), Spence and Apperley created a physical model of their lens idea by using two sticks to create a frame and running a strip of paper over these sticks to form a lens of this type with a single large central region

and quick perspective drop-off on each side. This particular mock-up has been made quite well known as an example of a 3D interaction sketch [4].

Our system can replicate Perspective Wall by dropping an information cloth onto a virtual object that resembles Spence and Apperley’s physical setup (i.e., a polyhedron with four of its edges as the vertical sticks, as in Figure 3). The magnification is provided using the same trick of 3D perspective, and the interaction of the information cloth even works in the same manner. That is, a person can scroll information to the right or left by dragging the cloth in that direction; the cloth would then move off one side as it slides up the other.

Graphical Fisheye Another technique that provides detail-in-context views is Sarker and Brown’s Graphical Fisheye [25]. This technique was used to distort node-link graphs so that nodes of interest appeared larger and more spread out, while compressing the remaining nodes into the surrounding space. Their implementation used a degree-of-interest function [12] that adjusts the size and position of nodes based on their distance to a point considered to be the “focus” of the visualization, as well as a parameter called “a priori importance”. By adjusting these functions the visual output of their system can produce a variety of subtly varying results. Commonly this type of lens offers a single focus point with increasing compression as distance from the focus increases. This compression becomes extreme at the corners of the image. Our reproduction of the single focus scenario places the information cloth on top of a cone with a smoothed top; this reproduction, while very similar, lacks the aforementioned extreme compression at the image corners (Fig. 4).

Because of this physical equivalence, the variety of distortions to size and position provided in their system can be approximated in our system with a variety of virtual object shapes onto which an information cloth can be dropped. The nodes landing on the top of the object are those in “focus” and the ones landing on the sides and edges are the smooth transition provided by the graphical fisheye.

Document Lens The Document Lens [24] provides a technique for viewing a large document whose pages are laid out in an array. It is similar to Perspective Wall in that a central rectangular area is used for magnified information, but with Document lens information drops off all four sides (including the top and bottom). In addition, a person can grab the rectangular lens and move it around to focus on different parts of the document.

Document Lens can be approximated in our system by dropping an information cloth onto a truncated square pyramid. Figure 5 shows (top-left) a Document Lens from EPF; a Document Lens object being used with an information cloth (top-right, and bottom with information). A person can navigate the information in this approximation in a similar way to Document Lens by moving the virtual object that is underneath the information cloth. In addition, a person can drag the cloth across this truncated pyramid to have it slide on and off the edges (an interaction not possible with the original document lens).




Figure 6: A constrained Gaussian EPF lens (left) is approximated by a cloth-covered flattened hemisphere.



Figure 7: An EPF lens that is not readily reproduced with cloth is the Manhattan lens where the physicality of the cloth prevents the extremely sharp transition at the edge of the lens.

Constrained Gaussian Lenses Constrained lenses were originally introduced by Carpendale et al. [7] as Gaussian Lenses and were central to Carpendale and Montagnese’s framework, EPF [6]. Constrained lenses minimize the spread of the distortion into the context regions of the contextual information. By both providing the requested focal magnification and leaving as much as possible of the context undistorted, these lenses opened the doors to multiple focal points, many variations in the types of distortion used, and possibilities of higher magnification in the focal regions.

In our system, we can approximate constrained lenses with a variety of shapes and offer similar smooth transitions between lenses and context by simply allowing the parameters of the physics engine (e.g., stiffness of the cloth and force of gravity) to deal with the transition regions. For example, we can approximate a Gaussian lens by dropping an information cloth onto a hemispheroid object (Figure 6). The stiffness of the cloth can vary from tight wrapping with relatively sharp transitions to the more gradual draped transitions used in our images. The cloth containing the contextual information provides an understandable transition from the object-draped lens and ripples slightly around the bottom of the object as we would expect physical cloth to do. These ripples can be smoothed through touch interactions.

Manhattan Lens Offsets and Magic Lenses [2] (magnified regions of maps and images often displayed in margins), DragMag [29] (offsets with simple lines from the corners




Figure 8: These images show an information cloth version of melange.

to indicate the region that has been magnified) and Manhattan lenses [6] (with abrupt straight direct connections between focus and context) are a group of a detail-in-context methods that provide magnified information, with little or no transition to the contextual information. These types of lenses perhaps most closely resemble physical magnifying glasses, if they were square in shape. However, in terms of an information cloth reproducing these effects, cutting or tearing the cloth would likely be necessary. Instead we approximate the Manhattan Lens by dropping a cloth onto a virtual cube (or more generally a rectangular parallelepiped), as in Figure 7. However, the physicality of the cloth does not offer the extremely sharp transitions at the edge of the virtual object. Instead the cloth drapes around the bottom of the lens. This type of lens still does offer the ability to provide more magnification for small regions. This limitation is a general characteristic of information cloth; i.e., that is unable to exactly reproduce EPF dropoff curves with sharp features.

Mélange Another detail-in-context implementation is provided in Mélange [10]. Mélange supports comparison of separate parts of the image by, in effect, pushing away the parts that are not currently of interest. We produce this type of detail-in-context viewing in our system by using multiple virtual objects and draping an information cloth across them. Notice how in (Figure 8), regions of information that were farther apart can be brought closer together.

EPF Lens Settings Similarly to Elastic Presentation Frameworks [6], which unifies most detail-in-context lens through the use of many parameters, cloth can offer a great variety of lenses. With cloth, the control of EPF style options, such as the shape of the lens and the magnification dropoff, are achieved by the shape of the virtual object under the cloth. Objects different from those we have included can be imported from 3D modeling software as obj meshes. Inside our system, the size of the object and its height control the



Figure 9: In this image several lens objects have been placed out magnify the rivers passing through the city.

equivalent of EPF lens' size and magnification. The main difference is the comprehensible natural draping of the cloth as opposed to the more mathematically precise EPF lenses. For instance, EPF can guarantee visibility via parameter control while cloth interaction cannot. The EPF concept of folding, which in EPF is simply the re-positioning of the tops of lenses independently from their bases, is an idea that can be greatly extended in cloth. The repositionable benefits of information folding can be achieved by lifting, bunching, stretching, and folding the cloth. Lastly, moving the object or dragging the cloth over the object provides the ability to change the focus area/lens position.

Beyond Lenses

In addition to the support of physically-based versions of existing detail-in-context techniques, our system is capable of providing additional ways of manipulating visual information. By changing the relative spacing between objects, as well as their height/magnification, a number of effects useful for exploring the data set can be achieved, such as the *bridging* of data. Similarly, by moving an information cloth around with one's hands and fingers, the data can be *stretched*, *folded*, and *bunched* in interesting ways.

Bridging Interesting draping effects can be created by using multiple objects. One such effect is creating a bridge of cloth between two objects. By placing two objects and pulling the cloth taut between them, a bridge is formed between the two lenses. This brings the context between the lenses into greater magnification than the background as shown in Figure 10. By draping the cloth over two objects and then adjusting the cloth and pushing the lenses towards each other, the cloth between the lenses will go slack and drop out of focus. In this way, a person can focus on only the data presented on the lenses. These sorts of techniques, techniques and various combinations of them are possible with not only



Figure 10: A bridge between a smaller square lens and a larger rectangular lens.




Figure 11: Multiple touches stretch for a closer look.

two, but many objects. Figure 9 shows a complex detail-in-context view created with a couple of rectangular objects that are bringing the rivers in Winnipeg into focus.

Stretching Due to the stiffness property of an information cloth within the physics engine, when grabbing multiple points on the cloth and then pulling them apart, the data can be stretched (Figure 11). This stretching provides a different means of magnifying the information that also has a physical counterpart. This type of magnification can be particularly beneficial for temporary glances at specific parts of the data, for instance to read a label or identify a node in a diagram.

Folding Folding of an information cloth is also possible by grabbing a part of the cloth and moving it over other parts of the cloth. This technique works best when grabbing near the edge of the cloth with one finger, or grabbing a sequence

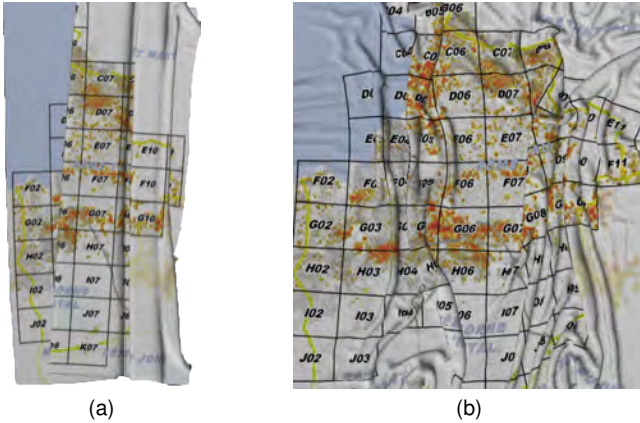


Figure 12: Information cloth also allows (a) folding, and (b) bunching.




Figure 13: In the left image the cloth has settled into a wrinkled configuration. The right image displays the resulting of touching the cloth above the cube and stretching the cloth to smooth a local region.

of points that spreads to two edges of the cloth (e.g., along the centre-line at the top, middle, and bottom of a map). Figure 12a shows an example of folding that was achieved by repeatedly using this latter method to fold the cloth like one might fold a map in a car. This technique is useful when it is desirable to see two distant parts of the same information space next to each other.

Bunching & Smoothing Another technique for bringing distant parts of the information space together is bunching (Figure 12b). Bunching can be achieved by repeatedly moving parts of the cloth with one's fingers to create ripples in the cloth. In contrast to folding, this technique can be used to prevent the hiding of the contextual information that lies between distant parts of the cloth. Alternatively, images can be interactively smoothed (Figure 13).

INFORMATION CLOTH IMPLEMENTATION

In this section, we describe the implementation details for the cloth and the objects, as well as how we enable multitouch interaction for each. We also describe how we recreate magnification effects similar to detail-in-context lenses using the interaction between the cloth and these objects. We make use of NVIDIA's PhysX physics engine² for simulation of cloth and objects in our environment. Our system was implemented on a SMART Table powered by a Core2Duo




Figure 14: Cloth interaction

6550 2.33 GHz CPU with 4 GB RAM and an NVIDIA 9600 GT graphics card that maintains frames rates of 30 Hz.

Information Cloth

The physics engine provides realistic collision detection and reaction between the information cloth, the ground surface, and any objects underneath the cloth.

The cloth parameters can be adjusted for several specific effects (e.g., elasticity, friction, thickness), however the exact parameter values are specific to the physics engine and virtual coordinate scale of our environment; consequently these parameters should be adjusted according to the needs of a specific application. However, we provide the following guidelines for determining suitable parameters. The first is that we specify relatively high values of stiffness and thickness of the cloth to prevent objects from penetrating the cloth. We additionally perform simulation with a relatively low resolution of cloth vertices (20,000), to assure interactive framerates but also to reduce the elasticity and distortion in the cloth's mass-spring system. We also have minimal friction in the virtual environment to reduce tenting and provide the opportunity for cloth to closely fit the shape of objects.

Data is displayed on the cloth with texture mapping [3]. The resulting visual effect is the appearance that the data has been printed on the cloth. In this fashion any 2D data representation can be turned into a 2D image and used on the information cloth. In general these textures need to be quite large (double the display's resolution or greater) in order to have sufficient resolution for magnification to reveal detail.

Cloth Manipulation A variety of techniques for cloth manipulation are possible. Wilson et al. [31] demonstrate that individual proxies can be used to control different parts of the cloth. This technique is consistent with our design goal of allowing different parts of the cloth to be independently controllable, however their technique does not provide a clear method of lifting the cloth to cover another object. Hancock et al. [16] use an alternative approach that allows the entire cloth to be moved in 6DOF as a single unit. This supports the ability to lift and cover objects, but does not provide fine control of different regions of the cloth.

²NVIDIA Corp. http://www.nvidia.com/object/nvidia_physx.html

To interact with the cloth each touch controls a different part of the cloth, but upon contact, the part touched is lifted a small, fixed distance above its current z position. As each touch is moved, that part of the cloth moves with the touch and is lifted over any objects that the touch passes over (Figure 14). The first design iteration treated each touch as a single point of contact, moving a point rather than an area of the cloth. This design allowed for the independent control of the different parts of the cloth and the ability to cover objects. However, lifting the single point created difficulty for the cloth’s mass-spring based calculations because the single points moved much more elastically than the rest of the cloth, providing an unnatural-seeming result and causing the cloth to sling-shot past the point in response. To achieve a realistic cloth action, we treated each contact as an area that included several points that were treated in the same fashion as we had previously treated the single point. The size of the area, approximately a finger’s imprint, does not change once the contact is made and all the points that compose this region move in unison.

Objects

The objects placed under the cloth to distort the presentation of the cloth’s data are simple geometric shapes that are treated as rigid bodies by the physics simulation. An object can be placed in the environment by selecting it from a panel and dragging the newly created object to its desired position. We make use of the shapes shown in Figure 15 to recreate the effects of common detail-in-context lenses, though other shapes described as triangles could easily be chosen. Our implementation uses *obj* triangle meshes, which can be created in most modeling software.



Figure 15: The object shapes that can be used as lenses in the current system.

Object Manipulation Each object can be moved by dragging it around with one finger, lifted with two fingers using Hancock et al.’s sticky fingers technique [16], and scaled in x and y using three or more fingers (described in more detail below). In the case where an object is covered by the cloth, upon contact, the cloth is raised above the objects (to the height of the tallest object) and removed from the physics engine so that the objects can be moved around underneath. The cloth is also made semi-transparent, so that objects are still visible behind the cloth. Once the fingers are lifted from the table, the cloth is then dropped back on top of the objects. The cloth can still be moved by touching anywhere upon the cloth where there is no object underneath.

The scaling in x and y is achieved by growing the object to encompass the bounding box of three or more touches. Thus, holding the object with two fingers and dragging to the right will stretch the object along the x axis by the amount it is pulled. Other gestures are also possible, such as starting with one’s fingers bunched together and spreading them out, to grow the object in both dimensions. Objects were also placed in kinematic mode (described in [31]) so that they did

not react to collisions with one another, but the cloth would still react to collisions with them.

Local Magnification

By dropping an information cloth onto virtual objects in our system, we can achieve local magnification. That is, the use of a perspective projection ensures that upward perturbations in the cloth caused by the objects’ presence will appear magnified relative to the parts of the cloth that fall to the ground, and thus provide physically-based detail-in-context.

While the physical nature of our technique is new, the idea of using perspective geometry to provide detail-in-context is that of the lens framework provided by Carpendale and Montagnese [6]. In this previous work, special attention was required to keep the lenses pointed at the camera to prevent occlusion of information. As can be seen in Figure 16 and Figure 17, our implementation also requires attention to this issue, but some modifications to their technique were needed to deal with the use of our physics simulation.




Figure 16: Comparison between standard perspective and tilting to achieve a top-down view.

Symmetric Drop-Off An important consideration noted by Carpendale and Montagnese [6] is that, because of the use of perspective projection to achieve magnification, 3D objects become skewed as they move away from the centre of the display (see Figure 16). This causes objects/lenses to function differently based on where they are positioned, in particular leading to occlusion of information. To make lens behavior position-independent, their solution is to skew the shape of the objects, shearing the lens in order to tilt them towards the center of projection as shown in Figure 17.

However, while it is possible to warp the objects in the same way in our system, the physics simulation behaves differently

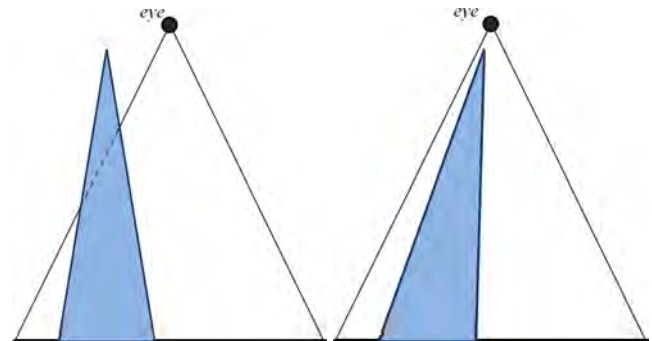


Figure 17: Left: a lens object (blue), when placed away from the center of the display in a perspective projection, can cause occlusion. Right: lens object has been tilted to make the object appear as if it is at the center of projection.




Figure 18: The blue object on the far plane is intersected by a vector from the eye at a non-perpendicular angle, causing the object to appear tilted. By placing the red object inside a large sphere this object's up-axis is aligned with the eye position, making the object appear as though it is viewed from directly above.

than one might expect. That is, if the objects are tilted in the physics simulation, although the objects will appear to be pointed directly at the viewer, the cloth would “hang” according to the leaning imposed by the correction. Thus, a different approach is required.

Instead of placing objects on a plane, our solution is to place them on the inside of a large sphere centered at the eye position as shown in Figure 18. We then manipulate the effect of gravity so that its direction is along the radius of this sphere, instead of being constantly downward. This aligns the vector from the eye to the object with the object's up axis making it appear as though the object is viewed from directly above, removing the perception of tilting due to the perspective projection. However, it is necessary to have the right balance between the projections field of view and the radius of the sphere. In order to minimize possible distortion when objects get close to the near plane we narrowed the projection's field-of-view and increased the sphere's radius.

This technique requires objects to adjust to the shape of the sphere when moved, rather than simply along a ground plane. We perform this adjustment by applying a transformation that maps movement along a plane to movement along this sphere; this transformation can be applied to any planar multi-touch object manipulation technique. The transformation is a combination of a translation that lifts the object from the plane to the sphere and a rotation that points the object's top toward the centre of the sphere. Mathematically, if p is a point at the centre of the object's base, e is the position of the eye, U is the object's model-space up vector, and c is the point on the sphere along the vector between p and e . With these gives the resulting translation becomes $c - p$ and the rotation is $\arccos(U \cdot (e - c))$ around the vector $U \times (e - c)$.

DISCUSSION

In this section, we describe some limitations of information cloth and provide two use cases.

Limitations

Instead of the precision of algorithmically-generated lenses, information cloth can have folds and stretching. These features maintain the feel and appearance of working with actual cloth but have the potential to be distracting or to obstruct useful contextual information. While these features can be interactively smoothed, this requires additional effort. Note

that shading used in displaying the cloth is also necessary with existing techniques [33]. The physical considerations of the cloth also present difficulties in simulating sharp and discontinuous drop-offs such as in Manhattan lenses. Lastly, the particularities of the simulation limit the height that objects can obtain before the simulation breaks and the cloth is pulled through the object. This limit varies depending on object shape, but in general this upper magnification limit is 3-5 times magnification.

Another constraint in the current system is that the cloth's physical properties that cannot be changed (e.g., elasticity, mass, friction, etc). The properties we make use of are the result of balancing the physics engine's capabilities against computational power and responsive, controllable interaction. With future increases in computational power, as well as expanded physic engine capabilities, future versions of information cloth will ideally present the opportunity to choose between different pieces of cloth with varying physical properties (e.g., silk, rubber, burlap, etc); we can then make use of the cloth that best suits the task.

Usage Scenarios

To further illustrate information cloth's strengths and limitations, we discuss how it can be used to accomplish two example tasks.

To start, consider finding a specific node in a large node-link diagram on information cloth. There are two ways one could do this. The first would be entirely by touching the cloth; one could simply use an expanding pinching gesture with both hands to magnify small areas to search, then release, and then pinch again to search another area. One or two fingers can be used to drag the cloth up and down or side to side to “pan” off-screen parts of the cloth into view as necessary. Alternatively one could place an object in the center of this window. This provides smoother and more uniform magnification. Then one or two fingers can be used to drag the cloth across the screen and the object, bringing different areas of the map onto the object and into magnification.

Next, consider trying to find driving directions between two cities. One might begin by placing large objects on the starting and destination cities. To do so, select an object from the object panel and place it underneath the areas of the cloth where two cities are located. To increase the magnification of the cities (for a better view of how to get in and out) use an expanding pinch to make the objects taller. Then to change their width and height to match the city, place three (or more) touches on the object and move them inwards to shrink or outward to expand it to properly cover the city. This can be repeated with additional objects to magnify enroute locations of interest, such as rest stops or areas of tricky navigation. The cloth will naturally form “bridges” between the objects providing some additional magnification of highways/roads between the cities. Areas of no interest will not be magnified.

CONCLUSION

We have described information cloth, an implementation that places 2D information on a virtual cloth and provides multitouch interaction in a physically simulated environment. This expands tabletop interaction techniques to include infor-

mation exploration. Specifically we show how information cloth can approximate previous detail-in-context views that are directly manipulable through multitouch.

Simple objects and the information cloth can produce a great variety of readily changeable lens effects. The cloth recreations are able to express the overall behaviour of detail-in-context lenses using physically-based interaction.

Because people are very familiar with fabric, and their interactions with it, our system may be readily understandable without much training. This familiarity might make it more possible to understand information exploration as it happens through manipulating cloth. Furthermore, using a physics simulation for the movement of the cloth, may make it more clear how and why things are changing the way they are.

ACKNOWLEDGEMENTS

We would like to thank the Natural Science and Engineering Research Council of Canada, the Network of Centre of Excellence (GRAND), the Canadian Foundation of Innovation, and SMART Technologies for research support. We also thank the reviewers for their helpful comments on this work.

REFERENCES

1. Agarawala, A. and Balakrishnan, R. Keepin' it real: pushing the desktop metaphor with physics, piles and the pen. In *Proc. CHI*, 1283–1292. ACM, 2006.
2. Bier, E. A., Stone, M. C., Pier, K., Buxton, W., and DeRose, T. D. Toolglass and magic lenses: The see-through interface. In *Proc. SIGGRAPH*, 73–80. ACM, 1993.
3. Blinn, J. F. and Newell, M. E. Texture and reflection in computer generated images. *Commun. ACM*, 19:542–547, October 1976.
4. Buxton, W. *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan Kaufmann, 2007.
5. Cao, X., Wilson, A. D., Balakrishnan, R., Hinckley, K., and Hudson, S. E. ShapeTouch: leveraging contact shape on interactive surfaces. In *Proc. of Tabletop*, 129–136. IEEE, 2008.
6. Carpendale, M. S. T. and Montagnese, C. A framework for unifying presentation space. In *Proc. UIST*, 61–70. ACM, 2001.
7. Carpendale, S., Cowperthwaite, D. J., and Fracchia, F. D. 3-dimensional pliable surfaces: For effective presentation of visual information. In *Proc. UIST*, 217–226. ACM, 1995.
8. Carpendale, S., Cowperthwaite, D. J., Fracchia, F. D., and Shermer, T. Graph folding: Extending detail and context viewing into a tool for subgraph comparisons. *Lect. Notes in Comp. Science*, 1027:127–139, 1995.
9. Dietz, P. and Leigh, D. Diamondtouch: a multi-user touch technology. In *Proc. UIST*, 219–226. ACM, 2001.
10. Elmqvist, N., Henry, N., Riche, Y., and Fekete, J.-D. Melange: space folding for multi-focus interaction. In *Proc. CHI*, 1333–1342. ACM, 2008.
11. Forlines, C. and Shen, C. Dtlens: multi-user tabletop spatial data exploration. In *Proc. UIST*, 119–122. ACM, 2005.
12. Furnas, G. Generalized fisheye views. In *Proc. CHI*, 16–23. ACM, 1986.
13. Han, J. Y. Low-cost multi-touch sensing through frustrated total internal reflection. In *Proc. UIST*, 115–118. ACM, 2005.
14. Hancock, M., Carpendale, S., and Cockburn, A. Shallow-depth 3D interaction: Design and evaluation of one-, two- and three-touch techniques. In *Proc. CHI*, 1147–1156. ACM, 2007.
15. Hancock, M., Hilliges, O., Collins, C., Baur, D., and Carpendale, S. Exploring tangible and direct touch interfaces for manipulating 2d and 3d information on a digital table. In *Proc. ITS*, 77–84. ACM, 2009.
16. Hancock, M., ten Cate, T., and Carpendale, S. Sticky tools: Full 6DOF force-based interaction for multi-touch tables. In *Proc. ITS*, 133–140. ACM, 2009.
17. Hodges, S., Izadi, S., Butler, A., Rrustemi, A., and Buxton, B. Thinsight: versatile multi-touch sensing for thin form-factor displays. In *Proc. UIST*, 259–268. ACM, 2007.
18. Isenberg, P. and Fisher, D. Collaborative brushing and linking for co-located visual analytics of document collections. In *Proc. EuroVis*, volume 28, 1031–1038. IEEE, 2009.
19. Kruger, R., Carpendale, S., Scott, S. D., and Tang, A. Fluid integration of rotation and translation. In *Proc. CHI*, 601–610. ACM, 2005.
20. Mackinlay, J. D., Robertson, G. G., and Card, S. K. The perspective wall: Detail and context smoothly integrated. In *Proc. CHI*, 173–180. ACM, 1991.
21. Nacenta, M., Baudisch, P., Benko, H., and Wilson, A. Separability of spatial manipulations in multi-touch interfaces. In *Graphics Interface*, 175–182, 2009.
22. Pietriga, E. and Appert, C. Sigma lenses: focus-context transitions combining space, time and translucence. In *Proc. CHI*, 1343–1352. ACM, 2008.
23. Reisman, J. L., Davidson, P. L., and Han, J. Y. A screen-space formulation for 2d and 3d direct manipulation. In *Proc. UIST*, 69–78. ACM, 2009.
24. Robertson, G. G. and Mackinlay, J. D. The document lens. In *Proc. UIST*, 101–108. ACM, 1993.
25. Sarkar, M. and Brown, M. H. Graphical fisheye views. *Comm. of the ACM*, 37(12):73–84, 1994.
26. Shen, C., Lesh, N., and Vernier, F. Personal digital historian: story sharing around the table. *Interactions*, 10(2):15–22, 2003.
27. Shen, C., Vernier, F. D., Forlines, C., and Ringel, M. Diamondspin: an extensible toolkit for around-the-table interaction. In *Proc. CHI*, 167–174. ACM, 2004.
28. Spence, R. and Apperley, M. *Readings in Information Visualization: Using Vision to Think*, chapter Data base navigation: an office environment for the professional, 333–340. Morgan Kaufmann, 1999.
29. Ware, C. and Lewis, M. The dragmag image magnifier. In *Proc. CHI*, 407–408. ACM, 1995.
30. Westerman, W. *Hand Tracking, Finger Identification, and Chordic Manipulation on a Multi-Touch Surface*. PhD thesis, University of Delaware, 1998.
31. Wilson, A. D., Izadi, S., Hilliges, O., Garcia-Mendoza, A., and Kirk, D. Bringing physics to the surface. In *Proc. UIST*, 67–76. ACM, 2008.
32. Wu, M. and Balakrishnan, R. Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. In *Proc. UIST*, 193–202. ACM, 2003.
33. Zanella, A., Carpendale, S., and Rounding, M. On the effects of viewing cues in comprehending distortions. In *Proceedings of ACM Nordi-CHI*, 119–128. ACM, 2002.