

Providing Dynamically Changeable Physical Buttons on a Visual Display

Chris Harrison Scott E. Hudson

Human-Computer Interaction Institute

Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213

{chris.harrison, scott.hudson}@cs.cmu.edu

ABSTRACT

Physical buttons have the unique ability to provide low-attention and vision-free interactions through their intuitive tactile clues. Unfortunately, the physicality of these interfaces makes them static, limiting the number and types of user interfaces they can support. On the other hand, touch screen technologies provide the ultimate interface flexibility, but offer no inherent tactile qualities. In this paper, we describe a technique that seeks to occupy the space between these two extremes – offering some of the flexibility of touch screens, while retaining the beneficial tactile properties of physical interfaces.

The outcome of our investigations is a visual display that contains deformable areas, able to produce physical buttons and other interface elements. These tactile features can be dynamically brought into and out of the interface, and otherwise manipulated under program control. The surfaces we describe provide the full dynamics of a visual display (through rear projection) as well as allowing for multitouch input (though an infrared lighting and camera setup behind the display). To illustrate the tactile capabilities of the surfaces, we describe a number of variations we uncovered in our exploration and prototyping. These go beyond simple on/off actuation and can be combined to provide a range of different possible tactile expressions. A preliminary user study indicates that our dynamic buttons perform much like physical buttons in tactile search tasks.

Author Keywords

Shape displays, tactile, haptic, pneumatic, dynamic buttons, programmatically controlled, multitouch, rear projection, barometric, pressure, input, dashboard, eyes-free, physical interfaces.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

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, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2009, April 4–9, 2009, Boston, Massachusetts, USA.
Copyright 2009 ACM 978-1-60558-246-7/09/04...\$5.00.

INTRODUCTION

The tactile sensations produced by physical buttons often make them easier to find and use than their purely visual analogs, allowing them to e.g., be employed with less attention demand [24]. However, the physical nature of these interfaces typically requires them to remain static in configuration, appearance, and tactile expression. For many interfaces, this means a single button must be overloaded with different functions at different times or places. At the other extreme are displays with touch sensitive surfaces. These can present interactive elements that are fully adapted to the state of the interface. However, they normally are precluded from providing any physical "feel" for those elements - presenting instead only a uniform flat surface. (Partial exceptions to this include vibrotactile feedback during activation [9, 21] and Nokia's forthcoming Hapticos variable-friction touch screen technology [15])

In this paper, we describe a technique for creating dynamic physical buttons using pneumatic actuation. This technique allows aspects of physical form and appearance to be dynamically modified under program control - buttons can be brought into and taken out of an interface as needed, and it allows a small set of distinct interface elements to occupy the same physical space at different times. Further, the characteristics of this tactile display allow it to be constructed using transparent and translucent materials. This allows it to accommodate a visual display (via rear projection) as well as multitouch input sensing (though infrared lighting and a camera behind the surface). This means that devices constructed using this technique can occupy an interesting new middle ground between highly physical, but static conventional buttons, and the highly dynamic, but entirely virtual form of conventional touch screens.

RELATED WORK

Shape Memory Materials (SMMs) are perhaps the most promising technology for tactile displays. These materials, currently alloys or polymers, actuate themselves into different lengths or shapes when they undergo a phase change (e.g., by heating or cooling) [22]. Presently, most varieties require high actuation temperatures (e.g., 50°C), which is problematic, as users need to touch the material. However, recent breakthroughs have introduced types that react to particular frequencies of light, chemical changes, and magnetic fields. Unfortunately, these materials are either expen-

sive or experimental. Slow switching times and modest deformation potential further limit their practicality. Finally, surfaces relying on SMMs (e.g., [6]) have yet to develop a method for rendering graphics without occlusion or voids, and capture user input – both significant challenges.

Pneumatics, the form of actuation we employ, has been used for tactile purposes in the past. For example, Enriquez et al. describe a pneumatically augmented steering wheel for alerting drivers [8]. Hashimoto et al. propose a magnetically driven pneumatic system able to create sensations like that of liquids or organic materials [11]. Lastly, the inflatable mouse, presented by Kim et al., enables several interesting tactile display and pressure input opportunities [13].

Many other tactile actuation/display methods exist (see, e.g., [2] for a much more complete discussion). These include for example: pin matrices [20, 23, 30], vibration [4], piezo-electric actuation [15, 24], ciliated surfaces [5], and electromagnetic actuation with voice coils [17] or a simple solenoid [16]. Many of these technologies have been employed to assist in target finding on touch screen displays (e.g., [9, 15, 21]).

Our technique stands in contrast to all aforementioned methods in several notable ways. Foremost, our technique is an integrated display, input and tactile solution. Second, our displays are inexpensive. This is primarily because the materials we employ are simple and each tactile feature does not require a dedicated actuator. Additionally, the physical actuation of elements does not require placing motors, wires, conduits, or other items behind the display, which precludes rear projection and increases build complexity. Simultaneously, diffused illumination multitouch sensing in concert with rear projection enables considerable scale at little extra cost or complexity.

Finally, Poupyrev et al. defined a class of technologies called shape displays [26]. These surfaces embody two distinct qualities – the ability to physically displace their surfaces and combine these deformations with images. The method we present in the paper fits well under this definition. Examples of previous systems include Kodama’s Flow and Protude installations [14], as well as FEELEX [12] and Lumen [25].

PROTOTYPES AND EXAMPLE CONFIGURATIONS

As a part of early work, we identified five distinct goals for our dynamic tactile displays:

- 1) Inexpensively manufactured (i.e., simple construction, cheap materials).
- 2) Easily actuated.
- 3) Able to display graphics without occlusion from hands or pneumatic/control elements.
- 4) Able to sense user input (and not occlude the display).
- 5) Provide support for tactile expression beyond simple on/off state changes.

Using these goals as requirements, we were able to rapidly eliminate many potential techniques and technologies. We ultimately selected a simple design, one that creates an air chamber by layering several specially cut pieces of clear acrylic. A thin sheet of translucent latex is draped on top of this to act as a deformable projection surface.

This construction offered several unique benefits. Foremost, fabrication was straightforward – using our lab’s laser cutter, we were able to assemble working prototypes with complex features in under an hour. Secondly, the displays rely on inexpensive materials – acrylic, glue, and latex. Third, the air chamber can be negatively or positively pressurized with a small pump, allowing for easy actuation. Finally, using clear acrylic allowed the display to be rear projected, and thus not suffer occlusion from user input. Simultaneously, this offered the ability to employ diffused infrared illumination and an infrared camera for multitouch sensing [3, 18, 19].

However, it was unclear whether such a simple design would have the necessary expressiveness to achieve the final requirement: to support a variety of rich tactile states. To investigate this potential, as well as to evaluate the general robustness and feasibility of the design, we constructed numerous prototype pneumatic tactile displays. Through an iterative design strategy, we identified numerous methods that enable the displays to assume several tactile states. Many of these techniques serve as “building blocks” from which complex, multi-state, tactile displays can be constructed. We include figures that illustrate the construction of each type; a simple color-coding scheme is used:

- Acrylic elements are shown in various shades of grey.
- Areas where adhesive is applied are shown with a textured blue.
- The thin latex layer shown as a translucent green.

Simple Deformations

The simplest design employs an acrylic layer with cut-out areas, on top of which lies a latex layer attached with adhesive. When the air chamber is negatively pressurized, the latex deforms inward, forming concave features (our dynamic buttons). Conversely, when a positive pressure is applied, the latex stretches outwards, forming convex features. The shape of these deformations is determined by the openings cut into the acrylic backing layer. The display is simply flat when no pressure is applied. Figure 1 illustrates the construction of this form, as well as photographs of the actual display in the three tactile states (We emphasize the tactile features through oblique lighting, which creates shadows and specular highlights).

Complex shapes can be readily achieved by producing a backing layer with complex cutouts. Figure 2 offers some examples, including an arrow, car and person.

Decoupling Positive and Negative Forms

A drawback to the previously mentioned design is that the negatively and positively pressurized deformations are

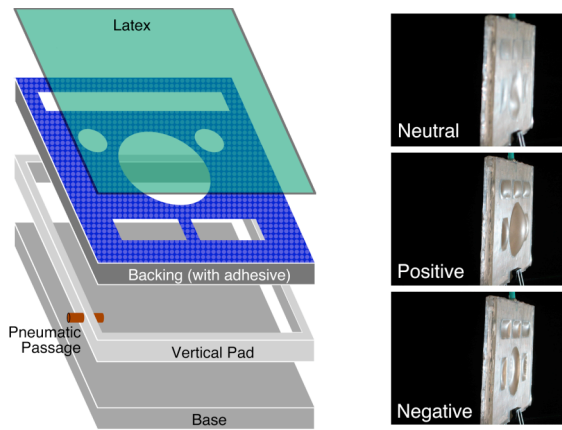


Figure 1. Full adhesive covering on the backing layer produces equivalent negative and positive tactile states.

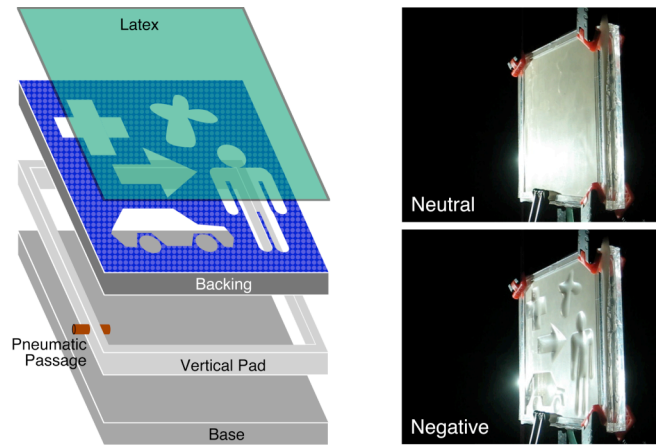


Figure 2. The backing layer can feature complex shapes.

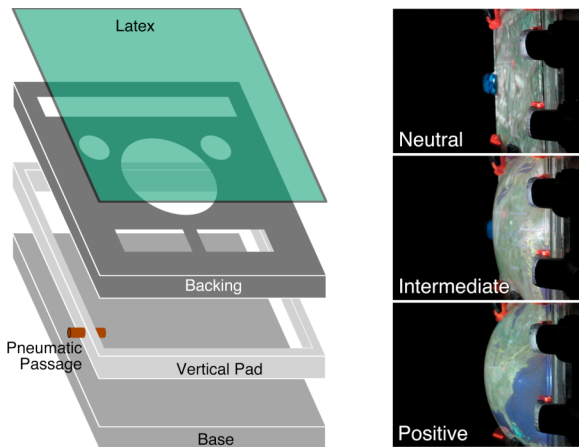


Figure 3. Without adhesive to define positive tactile features, the display can expand into a hemispherical shape, shown here with a projected Earth.

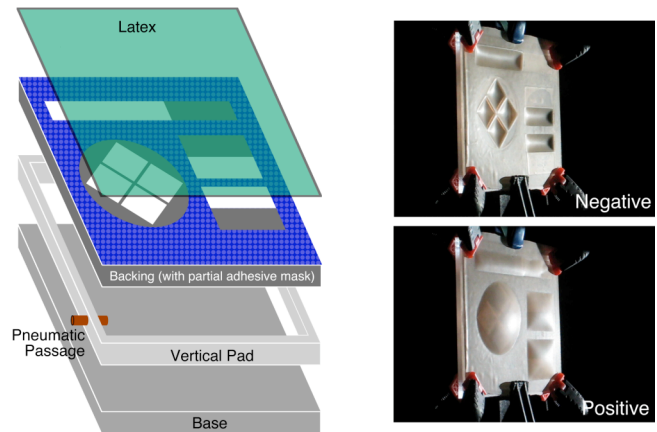


Figure 4. Complex masks can be used to create distinct positive and negative states.

roughly identical. However, it is possible to support independent positive shapes by modifying where adhesive is applied. Figure 3 illustrates an example of this construction where almost all adhesive is removed, allowing the positive deformation to take on a roughly hemispherical shape. This dramatic deformation might have use in viewing contexts with oblique viewing angles (i.e., offering viewers at even 180° off-center the ability to see some content). Other uses might include visualizing geo-spatial data and shared content similar to [3].

Through the use of more intricate adhesive masks, it is possible to produce other variations between positive and negative states. This is illustrated in Figure 4, where for example, a circular positive area covers the same space as a grouping of diamond-shaped negative elements. Positive shapes must contain all of the negative features they overlap with. It is possible to produce positive features without any (discernable) negative features by using small holes in the backing layer to allow air to reach masked areas (Figure 5). This simultaneously solves the problem of areas with cutouts having compliance (i.e., still depressible) when

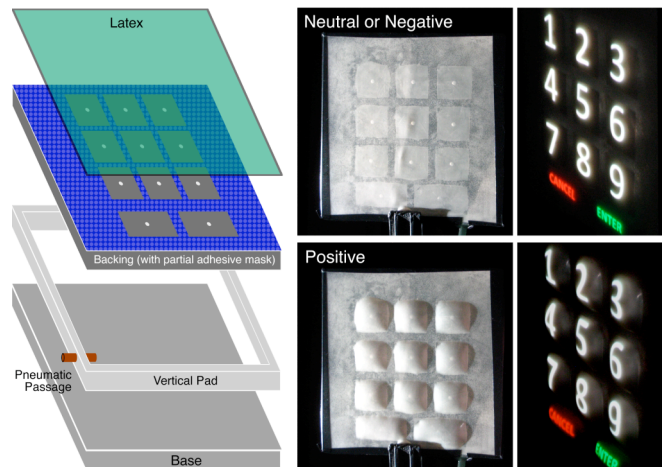


Figure 5. Small holes allow adhesive-masked tactile features to be actuated, but provide the sensation of a contiguous hard surface when neutrally or negatively pressurized.




Figure 6. Stilts in the vertical pad allow areas of the display to be surrounded by “floating” tactile features.

in the neutral state. Instead, users only feel the rigid backing when the display is neutrally or negatively pressurized.

Floating Features

Complex shapes, discussed above, consisted of areas cut out from the acrylic backing layer. However, this method precludes the ability to surround flat elements with tactile features (both negative and positive). One option is to employ stencil-like shapes, which contain thin supports in the backing layer to support “floating” features (like those required in letters like ‘B’ when spray painting). However, in our tests, we found these to unacceptably weak for a touch surface. Far more robust is to stilt features with matching elements in the vertical pad. Figure 6 illustrates this construction.

An interesting variation on this technique can be achieved by stiling elements similar to floating features, but with no surrounding backing layer. When a negative pressure is applied, the latex retracts, exposing positive tactile features surrounded by an indentation (Figure 7).

Inter-Latex Bladders

Stacked latex with interior pneumatic bladders offers another compelling technique for producing multi-state tactile displays. The simplest construction consists of two latex sheets. One sheet has an adhesive applied to one side - a mask in the shape of the desired tactile configuration is used (Figure 8). The second layer of latex is applied on top of this. It is necessary to mask thin channels running between the various features, as well as to the edge, so that air can be pumped in and circulate (only positive pressure is possible). With sufficiently thin latex (<0.2mm), it is possible to stack several uniquely configured layers (each requires its own pneumatic control). Furthermore, these bladder-augmented latex layers can be used on top of all previously discussed methods, adding to their already supported states and potentially creating complex tactile combinations (e.g., negative shapes with positive interior features).




Figure 7. Two levels of feature stilts produce protruding shapes within negative surroundings.

Dynamic Hard Buttons

It is possible to achieve a radically different tactile effect by reversing the latex and backing layers. This is achieved by cutting out features in the backing layer, but leaving the pieces in situ. A clear latex sheet is attached to the underside with adhesive. When a positive pressure is applied below this sheet, the latex deforms outwards, which in turn, pushes the acrylic elements suspended on top of the layer outwards as well (Figure 9). The result feels like raised, physical, plastic buttons that depress when pushed. A material, with matching cutout features, is needed on top to act as a projection surface.

Chambers

Finally, individually controlled pneumatic chambers, potentially as many as one per button, provide an obvious avenue for increasing the expressiveness of pneumatic tactile displays (Figure 10). Chambers can be combined with all the aforementioned methods, yielding scores of possible tactile states in a single enclosure.

Limitations and Downsides

Our displays are robust, inexpensive and scalable. However, several important limitations exist. Foremost, their tactile features are statically arranged. A graphical user interface wishing to take advantage of the tactile capability




Figure 8. Pneumatic bladders can be created in between latex layers with adhesive masks.

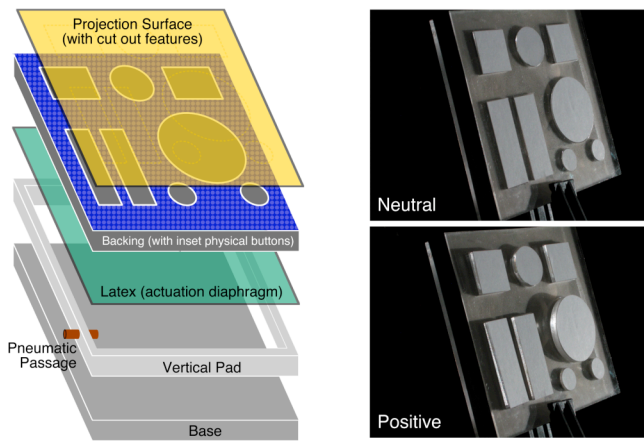


Figure 9. Physical features can be actuated by a latex diaphragm behind the backing layer. For clarity, only movable elements are covered with a projection surface (right).

must conform to what the display offers, invariably linking the two. This stands in stark contrast to previous efforts that offer per-pixel actuation, and thus greater flexibility (albeit with other downsides, such as cost and size) [15, 23, 30].

Additionally, our technique has considerable footprint, primarily due to its reliance on rear projection. Using modern, high definition, rear projection televisions as a rough guide, a 42-inch display would require around 40cm of depth. Integrating a flat touch screen display behind the pneumatically active area may be an alternative. This could allow the technique to see use in mobile devices. However, integrating a pneumatic pump or similar actuator into such a small form factor would still pose a significant obstacle. Thus, the technique is likely confined to large, semi-permanent installations (e.g., digital whiteboards, ATMs, check-in kiosks, car dashboards).

Lastly, our current implementation uses a latex sheet as the deformable surface. However, the durability of this surface over protracted use is unknown. Additionally, pneumatic actuation (like hydraulics) is somewhat “messy”. Tubes, pumps, valves and chambers may become prone to leaking, and degrade in performance over time.

SENSING INPUT

Several sensing techniques are possible with the pneumatic tactile displays we describe. Embedding small physical buttons or capacitive areas [7] beneath potential targets is possible and would be fairly inexpensive. However, these techniques limit interaction to predefined areas – large screens with several tactile states might become prohibitively complex. Simultaneously, these circuits would have to be sufficiently small so as to not interfere with rear projection. To sidestep these complexities, we employed a camera-based sensing technique.

Multitouch

Our touch sensing relies on a well-known technique employing two components operating behind the display: an




Figure 10. Pneumatic chambers allow different areas of the display to be actuated independently.

infrared camera and an infrared light source. Fingers approaching (and especially touching) the display reflect some infrared light, which is seen by the infrared camera as a bright spot. A computer program processes the live video, extracting blobs of infrared light (Figure 11). These can then be mapped to finger input events for use interactive applications. The transparent pneumatic elements introduce little interference (which, in any case, are static and can readily be subtracted from the image), allowing the entire pneumatic tactile display to act as a multitouch surface.

This technique [18], called Diffused Illumination Multitouch sensing, can be employed because the display is entirely constructed from infrared-transmissible materials: latex and acrylic. Furthermore, deployed setups have shown the technique to be highly robust and accurate [3, 19]. Frustrated Total Internal Reflection (FTIR) [10], an alternative multitouch approach, is not possible because the acrylic layer is not contiguous.

Barometry

Pneumatic tactile displays offer another unique sensing dimension: pressure. When a user depresses a pneumatic tactile feature (either neutral, negative, or positive in state), there is a jump in air pressure in the cavity behind the backing layer. Although this feature is in an early stage of development, we have identified several ways barometric data could be utilized.




Figure 11. Multitouch input is achieved with camera sensing of diffused infrared illumination. Here, three frames of video are shown with blob tracking identifying finger inputs.

The amplitude of pressure change is directly related to button displacement distance. Shallower presses yield smaller pressure increases, while deeper ones produce higher pressure changes (Figure 12, left). The force required to depress the elastic layer increases as it is stretched (Hooke's law), and so displacement distance (and thus pressure) is directly related to press force. This data could be used, for example, in magnitude controls, such as volume adjustment and scrolling velocity. See, e.g., [13, 27, 28] for other pressure-based interaction examples.

In a similar manner to how pressure increases when a tactile feature is depressed, pressure decreases when a tactile feature is released. This occurs because the elastic latex layer pulls back to its relaxed state, re-expanding the air cavity to its original volume (decreasing the internal air pressure). Thus, barometric data can provide both on-click and on-release finger events when interacting with tactile elements (Figure 12, right).

One of the weaknesses of Diffused Illumination sensing is that it is hard to disambiguate between fingers on the surface and those slightly above it (i.e., hovering). Pressure changes resulting from interaction with tactile features offer a straightforward method for rejecting false positives. It should also be noted that in the case of multi-finger input, pressure changes would be combined and ambiguous. If independent, multi-finger barometric readings are needed, several chambers will have to be employed, each with a dedicated sensor.

For barometric sensing, we use a SPC1000-D01 Absolute Pressure Sensor manufactured by VTI Technology. This provides a resolution of $\sim 3\text{Pa}$ with a sampling rate of 9Hz.

PNEUMATIC ACTUATION

Several air delivery systems are possible. We relied on small, electrically driven, fan-based pumps. These were simple to operate, portable, and safe (no external moving parts, low pressure). However, centrifugal pumps, vacuum pumps, high/low pressure reservoirs, and many compressor varieties are all strong candidates. Several of these technologies are available in "silent" varieties (15-30db, equivalent to that a soft whisper).

Some motors have the ability to produce negative and positive pressures. However, a vast majority can provide only one or the other. In these cases, an electronically controlled solenoid valve can be used to switch between negatively and positively pressured air lines. These could be run to separate pneumatic systems, or to the intake/outlet areas of a single motor. Although we have not yet experimented with this, the fast switching time of solenoid valves opens additional possibilities for expressive display. For example it should be possible to create pressure fluctuations such that buttons could pulsate when needing attention, and make other dynamic changes in response to input - for example, to create buttons that push back when pressed, or "snap in" when depressed.

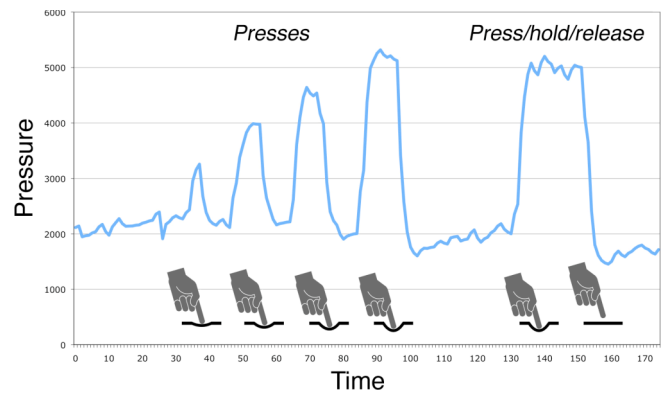


Figure 12. Air pressure in the chamber behind the display changes in response to a series of momentary presses of increasing force and a two-second press-and-hold action.

EXAMPLE APPLICATIONS

To help illustrate the utility of pneumatic tactile displays, we present three example domains where the technology could be deployed.

Automobiles

Touch screens are becoming increasingly prevalent in automobile consoles. This trend has been fueled by the need to incorporate several distinct interfaces into a limited surface area. A single screen, for example, might feature interfaces for GPS navigation, music selection, in-car entertainment, phone call management, and air conditioning. However, these screens, although providing significant flexibility, also remove many tactile clues that equivalent physical button interfaces would have provided. The result is increased need to look at the display to find and press targets, which clearly impacts driver safety [1, 29].

Pneumatic tactile displays offer a compelling alternative. Not only do they provide essentially equivalent graphical capability to that of a flat screen, but also feature sufficient tactile expressiveness to accommodate to the limited number of interfaces a dashboard would need to display.

As part of our investigations into the scalability of pneumatically actuated tactile displays, we built a large, single chambered, car dashboard mockup, seen in Figure 13. Actuation was achieved with a high-powered, commercial vacuum motor, which was able to switch states in about 3 seconds (internal air volume was ~ 8 liters). Although successful, it is likely that several chambers would be desirable for displays of this size.

Existing Multitouch or Large Screens

Most interaction on large screens (e.g., [7, 10, 19]) calls for uninterrupted, flat surfaces (e.g., drawing, viewing photos, exploring maps). It is for this reason that buttons are rarely seen integrated with the surface, as they inappropriately consume most of the time, and would thus consume useful input space. Contemporary systems typically employ soft (e.g., projected image) buttons, to avoid this dilemma. However, in some situations, buttons with physicality could be useful.

For example, when text entry is required, a physical keyboard could be raised up from the flat surface to aid in typing, and retract when complete.

Kiosks and ATMs

Like automobile dashboards, kiosks and ATMs often have simple and limited numbers of interface combinations. For example, the number keypad featured in touch screen ATMs (used for entering PINs and monetary amounts) is static in both size and location, making it an obvious candidate for transient, tactile augmentation, of the type that pneumatic tactile displays can provide. Not only could this produce an easier-to-use and potentially richer user experience, but also significantly help those with impaired vision or motor skills (where tactile cues could assist in target location).

We use the example of an ATM interface to showcase how our “building block” methods, discussed previously, can be combined in a simple way to produce a complex and useful tactile display. Figure 14 illustrates the display’s construction. There are three chambers, one for each button group. The two peripheral button groups are entirely cut out – allowing for them to be deformed inwards or outwards. To keep the center of the screen available for a variety of functions (including conventional interaction with virtual projected buttons using Diffused Illumination multitouch input), there are no cutout features. However, a number keypad can be inflated from the display surface through small holes in the backing layer to assist in numerical input tasks.

Figure 15 displays a series of example interface screens. An image of the tactile state alone (lit from the left) is shown below each for clarity. Note: The display was not designed to represent a functional ATM interface (hence missing elements like the number 0). Its purpose is simply to illustrate different tactile possibilities - 5 out of the 18 possible configurations are pictured.




Figure 13. A car dashboard was created as a large-scale test.




Figure 14. Construction of example multi-chambered, multi-state, ATM-like display.




Figure 15. Example ATM-like tactile display in various interface states. Tactile configuration only shown below for clarity.

An interaction might occur as follows: User enters PIN on raised number keypad (A). After four digits are entered, the enter and cancel buttons rise from surface, indicating the user can proceed (B). Cancel could be pressed at any time, regardless of actuation state (or given a dedicated chamber). Upon successful login, all current buttons retract and three new buttons rise from the screen (C). These allow users to select from one of three functions: check balance, withdraw and transfer. User selects withdraw. A keypad rises again from the surface to aid in number input (D). Peripheral buttons could be used for other functions, such as selecting currency type or which account to withdraw from. When an item is selected, the button group could retract. An alternative “quick cash” interface might require selection before the enter button is depressible (E).

USER STUDY

The primary goal of our tactile displays was to provide the flexibility of touch screens, with the tactile benefits of physical buttons. Our exploration of possible designs yielded a display technique with equivalent graphical capabilities to that of conventional flat screen technology (through the use of rear-projection). However, it remained unclear how air-driven, latex-based tactile features compared to conventional, physical buttons and flat touch screens – the two most prevalent technologies. To answer this simple, but central question, we devised an experiment that engaged users in a visually demanding task. This encouraged participants to use their sense of touch to locate a series of requested buttons. We recruited 15 participants (9 female, mean age 29), who received \$10 for their time.

Setup and Procedure

Four surfaces were created for the experiment (Figure 16):

- 1) Flat – a sheet of acrylic; designed to mimic a touch screen interface.
- 2) Physical Buttons – a sheet of acrylic with physical, protrusions; designed to mimic buttons.
- 3) Negative Pneumatic Buttons – pneumatic tactile display with concave features.
- 4) Positive Pneumatic Buttons – pneumatic tactile display with convex features.

All four surfaces had a topmost latex layer. This was necessary as a projection surface, as well as to make all four conditions texturally similar. Buttons were of identical size and layout in the four conditions.

Participants were seated in front of a laptop computer running a custom, full-screen application (Figure 17). A full-sized mouse was provided for input. The right side of the screen featured a layout matching the test surfaces. After a random period of time between 2 and 14 seconds, a button would activate by “lighting up” (see top-right-most button in Figure 17). This indicated to the participant that they should press the corresponding button on the test surface. Each button was activated, in a random order, four times for each of the four test surfaces (which were also presented in a random order).

To provide a measure of the attention necessary to find and press a button, we employed an attention saturating dual task framework [31]. In this framework, an attention saturating “primary” task is provided. Attention, which is necessary to complete actions in a secondary task (in our case finding and pressing a particular button), is then measured as a performance deficit in the primary task. In this way, a fine-grained measure of attention demand can be obtained.

For this experiment, our “primary” attention saturating task is placed on the left of the screen. It featured a circle moving randomly according to a two-dimensional Perlin noise function. Participants could counteract this force by manipulating a rubber-band-like leash, attached to their mouse cursor. They were asked to keep the circle centered over the crosshairs as best they could. The average distance of the ball from the crosshair served as the performance metric.

The test surfaces were oriented horizontally (i.e., parallel with the ground), and placed 25cm to the left and 7cm above of the laptop keyboard, which was centered in front of participants. A Mitsubishi PK20 short throw “pocket projector” was placed below the surface facing upwards, and projected a static image of the interface onto the test surfaces. Participants were asked to keep their hand no less than 15cm away from the test surfaces. This was necessary to simulate a reaching and target acquisition task, as might be seen while e.g., driving an automobile. Participants were told that keeping the circle centered on the crosshair was the most important task, and that reaction time to the button requests was not important.




Figure 16. The four surfaces evaluated in the user study.



Figure 17. The experiment interface.

When a button target was active (e.g., “lit up”), participants were allowed to move their fingers over the display surface to take advantage, if desired, of tactile features. Buttons on the flat and physical button surfaces were rigid. The ability to be depressed was not needed because the task sought to evaluate tactile search performance, which does not require buttons to be pressed (e.g., when one feels for the volume knob in car radio system, it is not necessary to press other buttons along the way). Additionally, visual attention is not needed to press a button once it has been located. However, to properly duplicate the “feel” of our pneumatic buttons, actual tactile displays were created for the experiment. This meant buttons on these displays were depressible.

When the correct button was located, participants were asked to provide a firm and distinct press. Once confirmed correct by a human observer, the experimental interface was advanced to the next random button (which activated after a random delay). Button press error rates were found to be negligible across the four conditions and not amenable to further analysis.

Several metrics were collected during the study. The laptop’s integrated, 640x480 webcam was used to film the participant’s face. These videos were captured to investigate participants’ eye movement behaviors, such as glance frequency and gaze time. The application running on the laptop recorded participant performance on the attention saturating task (i.e., how well the participants kept the circle centered over a target). The application also recorded how long participants’ took to locate and press the target button (i.e., reaction time). At the conclusion of the study, participants were asked to rank the difficulty of the four surfaces from easiest (1) to most difficult (4).

Results

As expected, flat screens performed the worst. Figure 18 (left) illustrates how much additional visual attention is required in the absence of tactile clues (two-sided t-test yields $p < .001$ for all conditions vs. flat). This reliance on vision for locating buttons clearly degrades performance on

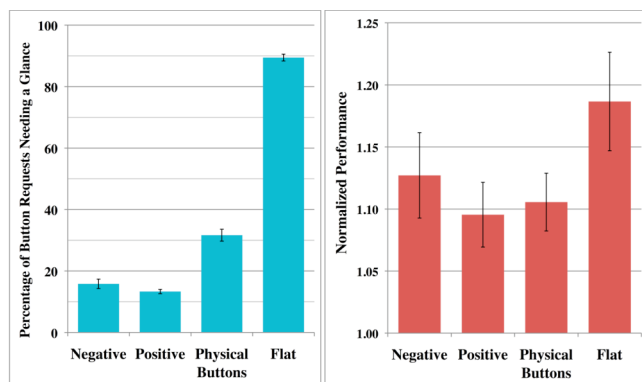


Figure 18. Left: Percentage of button requests in which participants glanced at the test surface to assist them in locating the correct button. Right: Averaged within-subject normalized performance scores on the circle centering task. Best performance fixed at 1.0, with other, worse scores relative.

the circle centering task (Figure 18, right). Encouragingly, our positively pressurized pneumatic buttons performed the best. A two-sided t-test comparing within-subject normalized scores indicates near significance between positive and flat conditions, with $p = .057$.

Interestingly, our two pneumatic displays appear to perform as well as the physical buttons (Figure 18, right). Subjectively, participants agree, ranking the three conditions favorably and fairly equivalent (flat was uniformly ranked the most difficult). This subjective ranking strongly aligns with the performance results - a relationship that can be readily seen in Figure 19, which plots performance on the circle centering task (objective difficulty) against participants’ average ranking of the four surfaces (subjective difficulty).

This result was fairly surprising. When we first began our investigations, we believed displays occupying the space between physical buttons and flat screens would have to compromise on graphical flexibility, tactile performance, or both simultaneously. In a way, we have achieved the best of both worlds - the graphical flexibility of touch screens with the tactile benefits of physical buttons, albeit with some unique limitations.

There also appears to be a slight positive performance trend towards raised tactile features. Positive pneumatic features and physical buttons both perform about the same, outperforming negative pneumatic features objectively and subjectively (although not significantly so) (Figure 19). We believe this slight advantage stems from superior tactile clues provided by outward facing features, such as convex deformations and edges. In the inward-facing negative condition, participants had to slide their fingers down into negative features to sense important tactile cues, such as depth and curvature.

Curiously, the physical buttons condition tended to require about twice as many glances, on average, as the two pneumatic conditions. This did not appear, however, to significantly degrade performance on the circle centering task. One possible explanation for this effect is that participants needed to confirm their location before committing to pressing a particular button. These “confirmation” glances may have required less attention than “searching” glances (like those needed in the flat condition). This effect may have resulted from a lack of tactile clues on the buttons




Figure 19. Average perceived difficulty plotted against average normalized performance for the four test surfaces.

themselves – only edges provided any tactile information, as all of the tops were flat. Pneumatic buttons, on the other hand, exhibited different curvatures based on the button shape and size, allowing greater confidence in button selection with fewer glances. Overall we believe this effect is linked to the particular design we employed for the physical buttons condition, and so caution against generalizing from it to conclude that our dynamic pneumatic buttons require less glances than conventional physical buttons.

Although participants tended to glance more frequently in the flat condition, and to a lesser degree with the physical buttons, there was no significant difference in glance duration across the four conditions. It appears all four conditions are equally accessible when vision is used to assist targeting. Additionally, there was no significant difference in button request reaction times across the four conditions.

CONCLUSION

In this paper, we have presented a range of methods for producing pneumatically actuated, dynamic, physical buttons on a visual display. The use of clear acrylic allows the displays to be rear projected, avoiding occlusion from user input. Multitouch sensing can be achieved with an infrared-based computer vision system. Additionally, we outline several interesting ways barometric features could be used to enhance input functionality. To help motivate the technology, we detailed three example applications where pneumatic tactile displays could be usefully deployed. We concluded with the results from our preliminary user study, which suggests our dynamic button displays have tactile qualities similar to that of physical buttons.

ACKNOWLEDGMENTS

This work was supported in part by grants from the Intel Research Council, General Motors, and the National Science Foundation under Grant IIS-0713509.

REFERENCES

1. Bah, K., Jæger, M., Skov, M. B., and Thomassen, N. You can touch, but you can't look: interacting with in-vehicle systems. In *Proc. CHI '08*. ACM Press (2008), 1139-1148.
2. Benali-Khoudja, M., Hafez, M., Alexandre, J. M., Kheddar, A. Tactile Interfaces: A State of the Art Survey. In *Proc. ISR '04*. 721-726.
3. Benko, H., Wilson, A. D., Balakrishnan, R. Sphere: multi-touch interactions on a spherical display. In *Proc. UIST '08*. ACM Press (2008), 77-86.
4. Brewster, S. and Brown, L. M. Tactons: structured tactile messages for non-visual information display. In *Proc. AUIC '04*. Australian Computer Society (2004), 15-23.
5. Coelho, M. and Maes, P. Sprout I/O: a texturally rich interface. In *Proc. TEI '08*. ACM Press (2008), 221-222.
6. Coelho, M., Ishii, H., and Maes, P. Surfex: a programmable surface for the design of tangible interfaces. In *CHI '08 Ext. Abst.* ACM Press (2008), 3429-3434.
7. Dietz, P. and Leigh, D. DiamondTouch: a multi-user touch technology. In *Proc. UIST '01*. ACM Press (2001), 219-226.
8. Enriquez, M., Afonin, O., Yager, B., and Maclean, K. 2001. A pneumatic tactile alerting system for the driving environment. In *Proc. PUI '01*. ACM Press (2001), 1-7.
9. Fukumoto, M. and Sugimura, T. Active click: tactile feedback for touch panels. In *CHI '01 Ext. Abst.* ACM Press (2001), 121-122.
10. Han, J. Y. Low-Cost Multi-Touch Sensing through Frustrated Total Internal Reflection. In *Proc. UIST '05*. ACM Press (2005), 115-118.
11. Hashimoto, Y. and Kajimoto, H. A novel interface to present emotional tactile sensation to a palm using air pressure. In *CHI '08 Ext. Abst.* ACM Press (2008), 2703-2708.
12. Iwata, H., Yano, H., Nakaizumi, F., and Kawamura, R. Project FEELEX: adding haptic surface to graphics. In *Proc. SIGGRAPH '01*. ACM Press (2001), 469-476.
13. Kim, S., Kim, H., Lee, B., Nam, T., and Lee, W. Inflatable mouse: volume-adjustable mouse with air-pressure-sensitive input and haptic feedback. In *Proc. CHI '08*. ACM Press (2008), 211-224.
14. Kodama, S. and M. Takeno. Protrude, Flow. (Electronic Arts and Animation Catalogue) In *Proc. SIGGRAPH '01*. ACM Press (2001), 138.
15. Laitinen, P. Tactile Touch Screen. International Application PCT/EP2006/009377, filed 27 September 2006.
16. Lee, J. C., Dietz, P. H., Leigh, D., Yerazunis, W. S., and Hudson, S. E. Haptic pen: a tactile feedback stylus for touch screens. In *Proc. UIST '04*. ACM Press (2004), 291-294.
17. Li, K.A. Baudisch, P., Griswold, W.G., Hollan, J.D. Tapping and rubbing: exploring new dimensions of tactile feedback with voice coil motors. In *Proc. UIST '08*. ACM Press (2008).
18. Matsushita, N. and Rekimoto, J. HoloWall: designing a finger, hand, body, and object sensitive wall. In *Proc. UIST '97*. ACM Press (1997), 209-210.
19. Microsoft surface. <http://www.microsoft.com/surface/>
20. Nakatani, M., Kajimoto, H., Kawakami, N., and Tachi, S. Tactile sensation with high-density pin-matrix. In *Proc. APOV '05*. ACM Press (2005), 169-169.
21. Nashel, A. and Razaque, S. Tactile virtual buttons for mobile devices. In *CHI '03 Ext. Abst.* ACM Press (2003), 854-855.
22. Otsuka, K. and Wayman, C. M. Eds. Shape Memory Materials. Cambridge University Press, New York, 1999.
23. Overholt, D., Pasztor, E., Mazalek, A. A Multipurpose Array of Tactile Rods for Interactive sXpression (Abstracts and Applications), In *Proc. SIGGRAPH '01*. ACM Press (2001).
24. Poupirev, I. Maruyama, S. and Rekimoto, J. Ambient touch: designing tactile interfaces for handheld devices. In *Proc. UIST '02*. ACM Press (2002), 51-60.
25. Poupirev, I., Nashida, T., Maruyama, S., Rekimoto, J., and Yamaji, Y. Lumen: interactive visual and shape display for calm computing (Emerging Technologies). In *Proc. SIGGRAPH '04*. ACM Press (2004), 17.
26. Poupirev, I., Nashida, T., and Okabe, M. Actuation and tangible user interfaces: the Vaucanson duck, robots, and shape displays. In *Proc. TEI '07*. ACM Press (2007), 205-212.
27. Ramos, G., Boulos, M., and Balakrishnan, R. Pressure widgets. In *Proc. CHI '04*. ACM Press, 487-494.
28. Rekimoto, J. and Schwesig, C. PreSenseII: bi-directional touch and pressure sensing interactions with tactile feedback. In *CHI '06 Ext. Abst.* ACM Press (2006), 1253-1258.
29. Rockwell, T.H. Spare Visual Capacity in Driving Revisited: New Empirical Results of an Old Idea, in *Vision in Vehicles II*. Elsevier (1988), 317-324.
30. Wagner, C. R., Lederman, S. J., Howe, R.D. A Tactile Shape Display Using RC Servomotors. *Electronic Journal of Haptics Research*, 3, 4.
31. Wickens, C.D. and Hollands, J.G. *Engineering Psychology and Human Performance*. Prentice Hall, 2000.