shiftIO: Reconfigurable Tactile Elements for Dynamic Affordances and Mobile Interaction

Evan Strasnick[1] **Jackie Yang**[1] **Kesler Tanner**[1] **Alex Olwal**[2] **Sean Follmer**[1]
[1]Stanford University, Stanford, CA, USA [2]Google, Inc., Mountain View, CA, USA;
{estrasni, jackiey, keslert, sfollmer}@stanford.edu, olwal@google.com

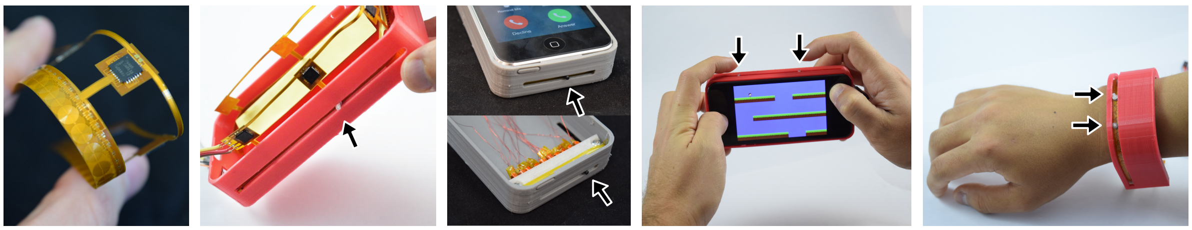

Figure 1: Reconfigurable tactile elements (shown in arrows) for interaction on a mobile device. Left to right: a flexible PCB utilizing micro-coils for magnetic actuation of tactile elements, a mobile prototype using the PCB micro-coil technique, a mobile prototype using low-power switchable permanent magnet actuation, a game played using reconfigurable tactile controls, a wearable band with tactile elements for display and interaction.

ABSTRACT
Currently, virtual (i.e. touchscreen) controls are dynamic, but lack the advantageous tactile feedback of physical controls. Similarly, devices may also have dedicated physical controls, but they lack the flexibility to adapt for different contexts and applications. On mobile and wearable devices in particular, space constraints further limit our input and output capabilities. We propose utilizing reconfigurable tactile elements around the edge of a mobile device to enable dynamic physical controls and feedback. These tactile elements can be used for physical touch input and output, and can reposition according to the application both around the edge of and hidden within the device. We present shiftIO, two implementations of such a system which actuate physical controls around the edge of a mobile device using magnetic locomotion. One version utilizes PCB-manufactured electromagnetic coils, and the other uses switchable permanent magnets. We perform a technical evaluation of these prototypes and compare their advantages in various applications. Finally, we demonstrate several mobile applications which leverage shiftIO to create novel mobile interactions.

CHI '17, May 06 - 11, 2017, Denver, CO, USA

ACM 978-1-4503-4655-9/17/05...$15.00
DOI: http://dx.doi.org/10.1145/3025453.3025988

ACM Classification Keywords
H.5.2. User Interfaces: Graphical user interfaces (GUI), Haptic I/O, Input devices and strategies

Author Keywords
Mobile Haptics; Tactile Display; Dynamic Affordance; Magnetically-actuated Buttons

INTRODUCTION
Current mobile devices allow users to choose from millions of applications. However, due to the convergence of hardware for smartphones and tablets, interaction with these different applications is generally limited to the same means of physical input—a touch screen and a few physical buttons. This greatly limits interaction, especially when the user cannot visually attend to the display, which is a common scenario when users multi-task in a mobile context, but also an everyday reality for the visually impaired. Thus, we seek to expand the capabilities of mobile I/O without radically changing the form factor or functionality of these devices. In particular we asked ourselves: What if the physical interface elements of a mobile device could reconfigure on the device to fit the application and user needs?

In this paper, we propose a new approach to mobile physical interaction: reconfigurable tactile elements (RTEs) which can travel on the exterior of traditional mobile and wearable devices. We specifically explore RTEs on the edges of mobile devices, resulting in shiftIO, a smartphone prototype with dynamic physical controls which can emerge from a hidden reservoir, move to a target location along its edge, and return to a hidden state. These RTEs can both provide haptic feedback and enable expressive input methods utilizing the dominant or non-dominant hand. As discrete, movable tactile elements, RTEs permit a number of interactions on a mobile interface, such as shear input and tactile display. They enable context-dependent physical controls for applications, and introduce new tactile notifications that allow a user to "glance" at information through touch. Because these RTEs operate on the edges of the device, users can interact with them without occluding the graphical display.

In selecting an implementation to realize the tactile elements, we chose to explore magnetic actuation with the goal of a small, lightweight, low cost design with few moving parts which could be integrated into mobile devices (phones and tablets), wearables (e.g. smart watches), and automobile dashboard interfaces or steering wheels. Magnetic actuation utilizing an array of electromagnetic coils and passive RTEs also enables the system to scale towards a high number of elements, in contrast to mechanical actuation methods. This paper examines the strengths and weaknesses of two different magnetic actuation techniques and the prototypes built to explore them. The first utilizes thin electromagnetic micro-coils integrated into flexible printed circuit boards, inspired by previous work [7, 27, 28]. These boards can be designed and fabricated with traditional Printed Circuit Board (PCB) techniques, making it ideal for applications requiring thin form factors and low cost. The second system uses a bi-stable design through switchable permanent magnets, which have a larger footprint and are more rigid, but potentially result in significantly lower power consumption. These systems have been integrated into two form factors. The first positions RTEs on the edges of a mobile device, the second utilizes the thin form factor of the Flexible Micro Coil system to be integrated into a wearable wristband.

Contributions

- The concept of Reconfigurable Tactile Element Interfaces to enable dynamic affordances and haptic feedback on the physical surfaces of devices.
- shiftIO - Two mobile implementations of such an interface which leverage the electromagnetic actuation of neodymium elements:
  - *Flexible Micro-Coil System* based on thin, flexible PCB with micro-coils for arbitrarily curved surfaces and minimal footprint in small devices.
  - *Switchable Permanent Magnet System* that enables bi-stability and low power magnetic actuation.
- A technical evaluation of the electrical characteristics and magnetic properties of shiftIO, discussing their suitability for mobile hardware.
- Novel applications and mobile interaction techniques for a Reconfigurable Tactile Element Interface that enable context-specific controls, self-adjusting interface elements, physical extensions of the GUI, and rich haptic notifications.

RELATED WORK

Mobile Tactile Interaction

Much research attention has focused on expanding the physical output modalities of mobile devices. In addition to actuating the entire device through vibration, there are broadly two classes of approach: 1) systems that provide surface haptics co-located directly on a touch screen display, and 2) tactile feedback on the periphery (e.g. edge or back) of the device. Haptic feedback can be added directly to GUI interactions through vibration [29, 5, 21], electrostatic friction [2] or reconfigurable elements emerging from the display (e.g. by pneumatics [11], hydraulics [8], or actuation [10, 30]).

Other contributions have utilized the periphery of mobile devices, exploring touch on the back [3], sides [4, 13] and surrounding regions [6]. Mobile devices also often include passive physical controls near the edges, and one line of research involves increasing the dynamic nature of these elements. Hemmert et al. created a single dynamic button for mobile interactions [12], and Pasquero et al. created a button with an array of piezo actuators to provide skin stretch directly to the user's thumb [24]. More recently, Jang et al. augmented a mobile device with an array of linear actuators to create dynamic affordances. [16].

Also related is the Eone Bradley tactile watch [34]. This watch uses two magnetic ball bearings in grooves to display the time both visually and tactually. This system has no input capabilities, and is limited to two tactors on different surfaces due to its use of a motorized actuation system.

While mobile tactile feedback as a whole has received much research attention, we believe that RTE Interfaces have a number of key distinctions from prior work, such as the ability to support lateral displacement for feedback and user input.

Reconfigurable and Actuated Input Devices

Researchers have also explored user reconfigurable physical input devices. Some work, such as that of Jansen and colleagues, customizes a traditional device with passive physical widgets that can be sensed for input [17, 38]. Villar and Gellersen used pushpin-style connectors and flexible circuit membranes [35]. In the mobile space, the MagGetz system used magnetic sensing to allow users to reposition physical input elements which were sensed by a mobile device's magnetometer [14]. These systems require the user to manually reconfigure the device, which makes them low-cost.

Outside of a mobile context, there has been much work in creating actuated table top interfaces with reconfigurable tangible elements [1, 20, 23, 26, 31]. Researchers have also explored how users perceive these moving physical affordances by leveraging patterns of motion and shape change [33]. Many of these systems use arrays of electromagnets to induce magnetic fields and move permanent magnets [23, 25, 36, 37]. However, such systems require large electromagnets, making them ill-suited to mobile applications. To address the size and power constraints of mobile devices, we require an alternative design. Furthermore, our system must work in various orientations and configurations. Our goal is to reduce the cost, size, and power consumption of such a system, and to develop meaningful interactions to work in the mobile context.

Magnetic Locomotion

Since at least the 1990s, researchers have been exploring magnetically levitated and controlled micro robots for manufacturing [9, 27]. These devices are fabricated using traditional PCB techniques [28] and utilize diamagnetic materials to levitate the magnets to reduce friction. Many of these systems use a row and column approach to drive the magnetic field. However, this presents challenges for independent control of multiple robots, leading to the exploration of alternative approaches [7, 18]. Inspired by this work, we aim to apply a similar technology to mobile user interfaces, with the additional challenges posed by interaction and display, such as the need for integrated sensing.

RECONFIGURABLE TACTILE ELEMENTS

We propose a new class of I/O for providing dynamic physical affordances on mobile devices called the Reconfigurable Tactile Element Interface. In these interfaces the physical elements can reconfigure their positions on a device to adapt to different applications or provide haptic feedback. While these elements can be perceived visually, their main function is in tactile interaction, and thus we chose not to explore integrated visual display elements such as LEDs.

Design Space

There are a number of parameters to be considered when designing a RTE Interface. These parameters affect the ways in which RTEs are used for display and interaction, as well as the size and power consumption of its components. We explore this design space in the section below.

Size. The size of a RTE changes how the user interacts with it much like a static button. Fitt's Law and ergonomic guidelines should be used to determine the ideal size. Smaller elements could be combined to form larger compound elements.

Number. The number of elements supported by the system has a large impact on its interactions and applications. A single RTE enables simple interactions with a single interface element, such as a scroll bar. With more RTEs, more complex interfaces can be generated, and more expressive tactile display can be achieved. Multiple elements could be attached together to form larger elements, and then split apart.

Dimensionality. This paper focuses on 1D actuation. However, RTEs could operate in 2D on a given surface, or stack to create elements of different heights.

Location. The RTEs can be located on different areas of an interactive device. RTEs could be located on the 2D visual display of a device, or on the back of the device. Here we explore interaction on the edges of the device.

Homogeneous vs. Heterogeneous. RTEs can all be the same shape and size, or they could be comprised of a set of differing geometries. For example, a larger button could be used as a camera shutter, whereas smaller buttons could form zoom controls. This could also enable the use of Phicons [15].

Visibility and Accessibility. RTEs could be exposed at all times. Alternatively, RTEs could be stored out of sight of the user in a reservoir. Because they are physical elements and cannot instantaneously appear/disappear, it is important to help the user to distinguish when the RTE is actively displaying information and when it is moving into a position.

Force. The amount of force a RTE can impart largely affects its use in feedback. Low force suggests that it can mostly be used for locomotion of the RTE. However, a RTE with higher force could impart force on a user and induce a haptic sensation, either by hitting the side of their finger, vibrating underneath it, or even moving the finger.

User Input. RTEs can be touch sensitive, either by integrating sensing into the RTE or elsewhere in the device. This touch could also be pressure sensitive to provide analog input. If the RTEs are backdriveable or loosely coupled to the actuation (i.e. through magnetic fields), it is possible for the user to reposition the RTE. In this mode, the RTE can be used for shear input through its lateral displacement, provided that there is appropriate position sensing.

Interaction

The RTE Interface consists of a number of passive RTEs that can be actuated to assume different positions and roles around the edge of a device. These elements can, for example, act as physical controls, haptic notifications, or tactile displays. They can emerge from a hidden state within the device itself, assume a given function on the device, and then return back into concealment when the interaction completes. Multiple elements can be actuated at once.

RTEs can be interacted with in a static state, wherein they assume a particular form when an application is launched and act like traditional physical controls. Alternatively, RTEs can use their ability to dynamically reposition to provide more active affordances. For example, a button moving quickly in an erratic pattern might imply that users should not touch it.

These RTEs can be controlled in coordination with the primary graphical display of the device. As such, there are different paradigms by which to design these joint interactions: By *mirroring*, RTEs can display the same interface elements as displayed on the graphical display, such as physical buttons near existing virtual buttons. By *complementing*, RTEs can display a spatially relevant interface element in addition to the graphical display, such as a scroll bar for text. Finally, by *extending*, RTEs can render information not represented on the primary display, such as a notification.

We describe below some of the main interaction primitives of the RTE Interface:

- *Buttons.* A RTE becomes a dynamic, touch responsive button on the edge of the device. Because it can be located easily via the sense of touch, the button is more readily recognized than a virtual button when attempting input.

- *Sliders.* A RTE acts as a linear slider, allowing the user to scroll through content by moving a physical control down the side of the device.

- *Toggles.* A RTE becomes a switch, where a tap causes it to toggle the value for a parameter, and correspondingly move to a new position that represents the updated value.
- *Pinch controls.* A pair of RTEs operate to form a pinch-gesture interaction with physical feedback, e.g. for zooming. Users can slide the RTEs closer together or further apart to adjust along a continuous scale.

In addition to generating these input elements, RTEs are also capable of providing feedback to the user through a number of techniques:

- *Haptic notifications.* RTEs can be used to "bump" into the user's hand as it grips the device, alerting them of new information in a discreet fashion.
- *Physical information display.* RTEs can represent discrete chunks of information, such as unread notifications or number of participants in a chat room, which can be perceived both visually and haptically by the user. Further, motion of the RTE can also be used for information display, such as rendering a loading bar, or a playback head for a music player.
- *Haptic detents in lateral interaction.* The device can create regions of varying force, such that a user moving an interface element along the device feels haptic pulls or resistance to their action.

SHIFTIO IMPLEMENTATION

Technical Considerations

There are many actuation approaches that could be used to implement RTE Interfaces, including belt/cable driven tactors, linear actuators (motor driven, S.M.A, pneumatic, hydraulic, etc.), electrostatic actuation, magnetic actuation, or self-actuated elements (such as microrobots). We considered a number of factors in the design of shiftIO, towards the goal of creating a system suitable for mobile and wearable applications:

Size. Size is a chief concern in mobile and wearable devices. Ideally, the actuation technology is thin and light, so as not to add thickness or weight to a mobile or wearable device. The length of the active area and max linear displacement of RTEs was also an important consideration, as we wanted to allow for interaction on multiple sides of the device.

Multiple Elements. Our goal was to support several elements simultaneously. This scaling issue limits the feasible technical solutions, as many technologies would make it difficult to drive RTEs along a single track. For example, if linear actuators were used they may physically collide or not have enough space, unless stacked. Cable driven systems could support more elements, but still run into space constraints. Magnetic drive systems could theoretically support many RTEs simultaneously, based on the number of magnetic coils present.

Force. The force of the RTEs is important both for haptic feedback and for the robustness of the system. Stronger lateral forces allow us to create greater haptic sensations, and a stronger normal force helps to keep the RTE from dislodging from its target position. The amount of force varies greatly with the chosen actuation mechanism. For example, large forces are most easily achieved via motor-driven systems.

Cost. The cost per RTE is also a consideration. Small motors and linear actuators generally have relatively high parts costs, so it is rare to find multiple of them in one consumer product. Often, components that can be easily fabricated with existing PCB technology can be low cost, due to the optimization and scale of this fabrication process.

Reduced number of moving parts. Lowering the number of moving parts that can be broken by users or external forces is a key consideration. Many linear actuators are not compliant or backdrivable, which would cause them to potentially break. Magnetic action provides compliance and robustness to high forces, as the RTE can just slide freely if a high external force exceeds the magnetic force.

Power Consumption. Power consumption for mobile devices is extremely important. Some technologies only use power to move elements (most linear actuators, switchable magnets), however others require power to hold a steady state (e.g. shape-memory alloys and electromagnets).

Implementation Design Selection

Prioritizing the goal of independently controlling a number of RTEs along the edges of a mobile device or wearable, we decided to utilize a magnetic approach to actuation, as opposed to mechanical alternatives such as linear actuators or belts beneath the edge of the device. We chose to explore two different approaches to magnetic actuation - thin and flexible elements that could be low cost to produce, and secondly, lower power consumption using switchable magnets. Both of these systems share a number of benefits including the reduced number of moving parts and overall low cost to actuate many RTEs, with the trade-off of a relatively low output force. Both versions of shiftIO provide enough force to allow for haptic feedback in the form of "taps" on the side of the finger, but not enough to physically displace the user's finger.

In the sections that follow, we refer to our specific implementations of a RTE Interface (shiftIO), noting that other implementations could facilitate somewhat different interactions and applications.

FLEXIBLE MICRO COIL SYSTEM

In the first version of shiftIO, the RTEs are actuated by a flexible multilayer PCB with patterned copper traces. To drive magnetic elements, we require only a thin, flexible strip of small coils, which can be manufactured through standard PCB fabrication techniques. Further, a single strip can be folded around to cover the entire perimeter of the device, simplifying the transition between edges. By continuing this strip beyond the exposed area, we create a reservoir area where RTEs can be stored in a "hidden" state when unused.

Running current through the circuit layers creates magnetic fields and imparts forces on the RTEs, causing the RTEs to move in a controlled direction. In addition to the lateral movement, the effect also creates a strong normal force which keeps the RTEs attached to the circuit surface even when vertical.

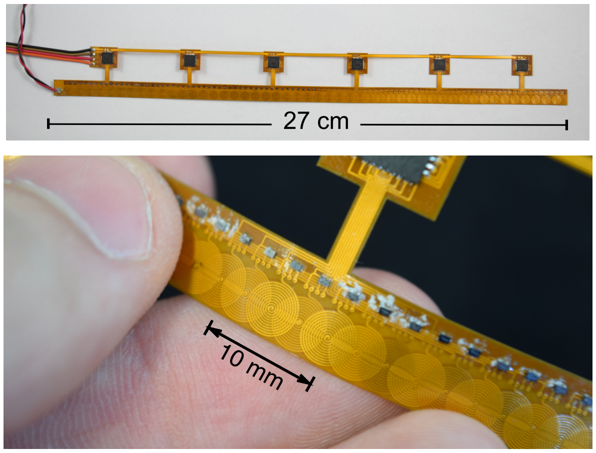

Figure 2: Interleaved micro-coils on a flexible PCB. Top: The total length of the PCB is 27 cm, allowing it to operate on all sides of a mobile device. Bottom: Close up of coils and transistors. Coils have a radius of 2.5 mm. Each transistor controls a single coil.

Each layer of the PCB contains a series of micro-coils (see Figure 2). The circuit layers are patterned identically, but are offset so that as the layers are driven independently, the magnets are pushed and pulled to the next position. The RTEs are driven in open-loop control via microstepping—i.e. the activation of a given coil is increased and decreased by adjusting PWM pulse widths.

In similar systems such as [7], a combination of repulsive and attractive forces have been used to create smooth motion of the travelling magnet. However, creating both attractive and repulsive forces from the same set of coils requires switching with an H-bridge configuration for each coil, drastically increasing the cost, complexity, and size of the circuitry. Instead, by solely leveraging attractive forces, we can switch each coil with a single transistor. The differences in drive patterns in

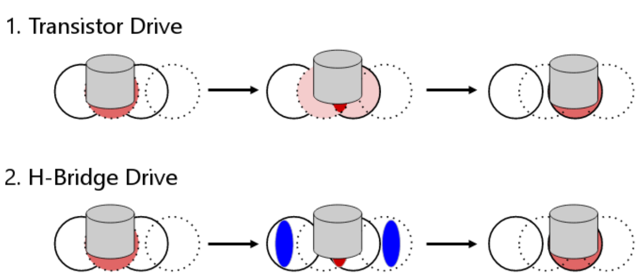

Figure 3: Single-transistor drive versus H-Bridge drive for our interleaved micro-coil setup. With a single transistor, a microstep involves transferring power from one coil to the next, shifting the attractive force. With an H-bridge, coils beneath the RTE produce an attractive force (red), while coils bordering the RTE produce a repelling force (blue).

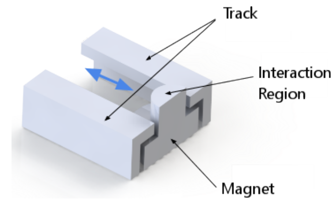

Figure 4: A mechanical track constrains the motion of the RTE, while exposing one end for user interaction.

the single transistor and H-bridge configurations are shown in Figure 3.

While a normal force is generated in the drive process to keep the RTE atop the PCB coils, we additionally mechanically constrain the motion of the RTE by fixing it to travel within a fixed track along the side of the device (as shown in Figure 4). The magnetic base of the RTE slides along the coils beneath the track, and a 3D-printed cap on the magnet protrudes as a region for user interaction. This keeps the RTE from escaping even when powered off, and helps prevent dislodging in the event of bumps or drops.

Coil Design

The first consideration in designing the PCB is the layout and number of coils. By interleaving multiple sets of coils offset in phase, we can achieve a smoother travel in either direction than with a single set of coils of the same radius. However, assuming the same force per coil, powering additional coils results in increased power consumption as well as an increased number of transistors required over the same length. shiftIO uses two sets of coils positioned 180 degrees out of phase, which can optionally be run using just a single set.

Secondly, we consider the design of each individual coil. Based on the work of Cappelleri et al., we utilize a spiral-shaped micro-coil, to maximize the field in the region beneath the cylindrical magnet in the planar PCB layer [7]. Our parameters include the radius of the coil, the thickness of the trace, and the number of turns. To inform our designs, we used finite element analysis (FEMM [22]) to explore the effect of various parameters on the resulting output force. We approximate our spiral shape as a series of concentric circles for the purposes of simulation. The simulation geometry is shown in Figure 5.

Given that the traces are relatively short and contribute minimal resistance, we can vary the trace width and assume a constant current without much error. The results of varying the number of turns in tandem with trace width are shown in Figure 6. As expected, a tightly wound spiral with thin traces produces the largest force on the RTE, so our coils are designed with the minimum trace width/separation per the manufacturer (0.125 mm), and the maximum number of turns (10) in an empirically chosen radius (2.5 mm).

Magnet Selection

Magnet Grade

shiftIO utilizes N52 Neodymium magnets as the RTE base. The N52 grade is one of the highest available, and Neodymium

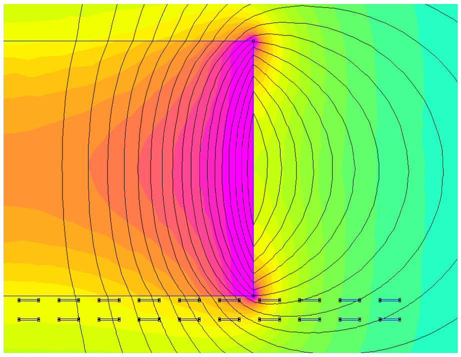

Figure 5: Axisymmetric finite element analysis showing the flux density of our magnetic RTE atop a single powered 2-layer, 10-turn coil.

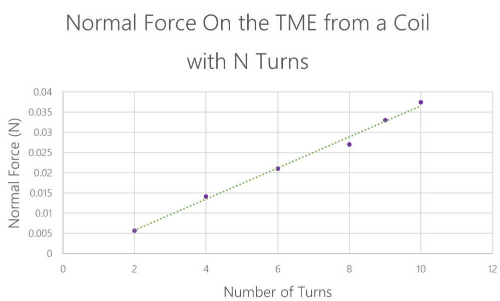

Figure 6: The normal force experienced by a RTE located directly above a single coil powered at 1A. As the number of turns are increased, the trace width is decreased to fill the same 5 mm diameter radius with a constant .125mm separation between turns. Normal forces were calculated using block integral stress tensors.

magnets are particularly suitable for translations on flat circuits given their high surface magnetic field.

Magnet Dimensions

The dimensions of our magnets take into consideration both the constraints of actuation as well as the ideal form for user interaction. While magnetic field strength increases with both thickness and diameter and enables greater attractive forces, increased dimensions also increase the weight of the RTE, which has a net decrease in performance. Thus, the ideal magnet for shiftIO has the smallest dimensions while still being large enough to permit user interactions.

We settled on a 1/16" (1.59 mm) magnet thickness, and a 1/8" (3.18 mm) magnet diameter. The diameter allows us to constrain the motion of the RTE within the mechanical track, while affording sufficient area for touch interaction. The thickness is sufficient to provide stability to the section of the RTE within the track, allowing it to resist torque resulting from the weight of the 3D-printed cap.

The mass of the magnet is 0.09 g, and the printed cap adds an extra 0.04 g, for a total of 0.13 g.

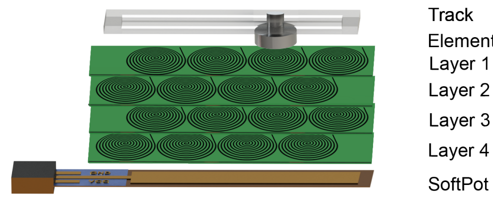

Figure 7: Exploded view of the PCB micro-coil system, consisting of a RTE, a linear track, flexible circuit layers containing interleaved electromagnetic coils, and a soft potentiometer.

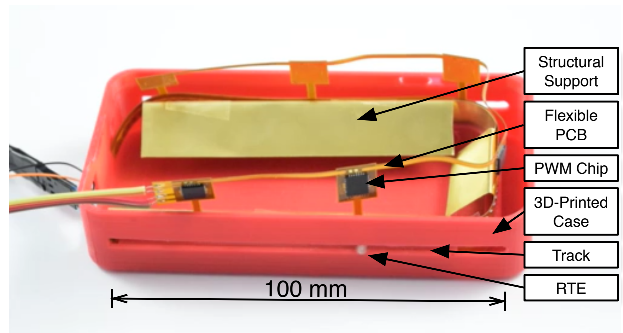

Figure 8: shiftIO prototype using a 4-layer flexible PCB. The 3D printed case has integrated tracks for the RTE to slide in.

System Design

Our final design (Figure 2) uses a four-layer flexible PCB, with one coil set in layers 1 and 3, and the other set in layers 2 and 4. A via connects each of two paired coils between layers to make a continuous trace, effectively increasing the number of turns in the same radius. The two sets are offset 180 degrees out of phase. Each coil has 10 turns in each layer, for a total of 20 turns, equating to roughly 1Ω of resistance.

An Arduino Uno microcontroller controls the coils using PCA9685 PWM ICs over I^2C communication. Each IC is capable of driving 16 coils, and each coil is switched with a CSD13383F4T transistor. In addition, a linear soft potentiometer behind the PCB is used to sense pressure and to calculate the input position when a user pushes a button into the track. An exploded view of the layers of the system is shown in Figure 7.

Mechanical Design

To test shiftIO with an existing mobile device, we 3D printed a case to house the electronics alongside an iPod Touch (Figure 8). The case features a linear track to constrain the motion of the RTEs, and has interior room for RTEs to “disappear” when not used in a given application. We utilize a Bluetooth Low-Energy UART module from Adafruit Industries to communicate with the iPod Touch, enabling interactive applications.

We also designed a wearable form factor utilizing a wristband-style device, see Figure 9. A 3D-printed case encloses the

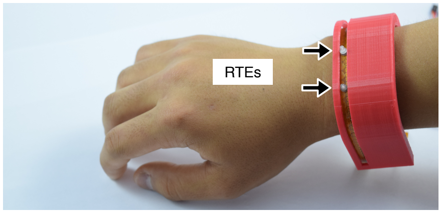

Figure 9: The wristband wearable device with RTEs. This prototype utilizes a 3D printed case and the Flexible Micro-coil array to actuate magnetic RTEs.

device and again has an integrated a linear track. The RTEs can travel around the wrist to display information to the user. While the current prototype does not have a coordinated graphical display or touch sensing, those could be added as for the mobile device.

Technical Evaluation

The first shiftIO prototype runs at 1.1V and draws a steady 0.5 A of current per RTE, regardless of whether the RTE is moving or stationary. RTEs can be actuated at speeds up to 80 mm/s. An individual RTE can be positioned to a resolution of ≈1 mm. Because of magnetic interactions between RTEs in close proximity, a minimum separation of 15 mm is required between adjacent RTEs.

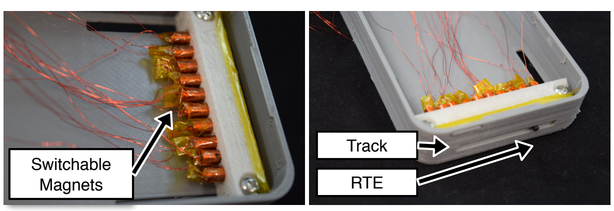

Figure 10: A prototype system utilizing a switchable permanent magnet drive.

SWITCHABLE PERMANENT MAGNET SYSTEM

In many applications, RTEs remain in static positions for significant periods of time. In our current design, holding this position requires continuous power. To demonstrate an alternate design with greater power efficiency, we prototyped a second version of shiftIO (Figure 10) which leverages switchable electromagnet actuators to generate the magnetic field. We use a magnet design similar to that described by Strasnick and Follmer, with Grade 6 AlNiCo magnets wrapped in a solenoid [32]. Because of the low coercivity of the AlNiCo magnet, when current is briefly pulsed through the wire, the polarity of the magnet is permanently changed. This allows us to maintain an attractive force on the RTE without continuous power.

We use the same N52 neodymium magnets for the RTEs. Switchable permanent magnets are lined up along the edge of the device (Figure 10). Each AlNiCo magnet is wound with $N = 140$ turns of 36 AWG wire, yielding a radius of 1.9 mm. Rather than traveling directly atop the AlNiCo magnets, RTEs travel along a spacing surface. The width of this spacer (2.275 mm) was chosen to be thick enough for the AlNiCo magnets to switch easily in the presence of the RTEs magnetic field, yet thin enough to still impart significant forces on the RTE. As in the previous system design, we add a linear track to the case to additionally constrain the motion of the RTEs.

We also use a similar PWM-based microstepping approach to generate smooth motion of the micro-robot. However, since we cannot interleave coils as in the flexible PCB variant, we utilize MOSFETs in an H-bridge configuration on each AlNiCo magnet to create variable strength attractive and repulsive forces (Figure 11). The combination of repulsive and attractive forces to create a motion vector is similar to the translation motion primitive described in [7]. Thus, when the magnet is moving, at most 2 AlNiCo magnets are being switched. When the magnet is at rest, power is completely disconnected.

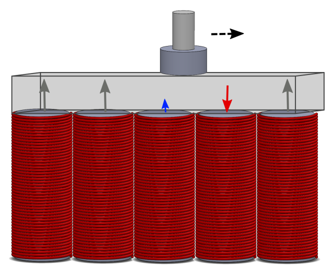

Figure 11: Microstepping a tactile element towards the right across a surface by using switchable permanent magnet actuators. Grey arrows represent magnets which are not being switched, but remain in the maximally repulsive state. The blue arrow (an increasing repulsive force) and the red arrow (a decreasing attractive force) constitute the microstep.

Though the system operates at 29V to allow for large current spikes, these spikes are brief and infrequent, resulting in a low amortized current. 100uF capacitors are charged up and discharged when switching a magnet to prevent large current draws from the power supply. Our prototype was driven atwith a 100 μs pulse length and a PWM frequency of 500kHz, empirically determined to be the minimal pulse length to reliably switch the permanent magnet using our components.

Technical Evaluation

The switchable permanent magnet variation exhibits a linear average power response with respect to the traveling speed, controlled by adjusting the delay between microsteps (Figure 12). While the system draws more power when moving RTEs than the coil-based prototype, we can amortize its cost over time spent with RTEs in a static position to find a break-even point at which the switchable permanent magnet variation becomes efficient. Examining the single RTE case, let T_m be the time spent moving the RTE, and T_s be the time spent with the RTE stationary. P_s is the power consumption for a stationary RTE, P_m is the consumption for a moving RTE, and P is

Figure 12: Average power consumption of the switchable permanent magnet variant while moving a single RTE, as a function of its speed.

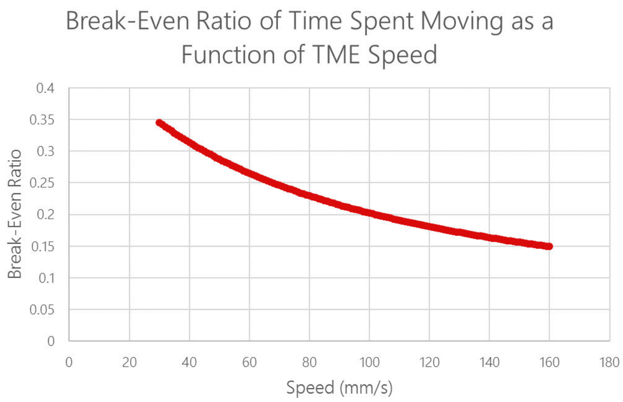

Figure 13: The fraction of total operation spent moving a RTE below which the switchable permanent magnet system is more power efficient than the micro-coil version. This break-even point is lower when moving the magnet at greater speeds on the switchable magnet system.

the total power consumption. Then:

$$P = \frac{P_s * T_s + P_m * T_m}{T_s + T_m}$$

For the PCB-coil version, $P_s = P_m = 0.55W$. For the switchable permanent magnet variant, $P_s = 0$, and P_m is a function of v, the traveling speed of the magnet in mm/s: $P_{m[switch]} = 0.0161v + 1.109$ Setting these equations equal for the two prototypes, and substituting in the empirically measured power consumptions, we can determine the point at which the switchable magnet variation becomes power efficient:

$$\frac{T_m}{T_s + T_m} = \frac{0.55}{P_{m[switch]}} = \frac{0.55}{0.0161v + 1.109}$$

This equation shows the break-even percentage of time spent moving for power consumption between the two systems, as a function of the speed of the switchable magnet variation. That is, if the RTE is moving for more than $\frac{0.55}{P_{m[switch]}}$ of the total time of operation, then the switchable permanent magnet version is more power efficient. The break-even ratio is plotted as a function of speed in Figure 13. For example, at a speed of 50 mm/s, the switchable magnet variation is more efficient if the RTE is moving less than 28.7% of the time. This implies that the ideal drive mechanism (from a power perspective) is dependent upon the characteristics of the target application. A technical comparison of the two prototypes is presented in Table 1.

APPLICATIONS

shiftIO enables a wide range of mobile interactions that can leverage dynamic interface elements, tactile notifications, and

	PCB Micro-Coil	Switchable Magnet
Depth*	0.125 mm	9.9 mm
Flexibility	Yes	No
Power (stationary)	0.55 W	0 W
Power (moving)	0.55 W	0.0161v + 1.109**
Maximum speed	80 mm/s	160 mm/s
Lateral force	.01 N	.02 N

*

Protrusion into device, RTE not included

** v is the speed of the RTE in mm/s

Table 1: Comparison between the two implemented shiftIO prototypes. Power consumption is per RTE.

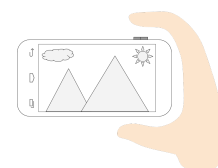

(a) A physical camera button which adjusts according to orientation

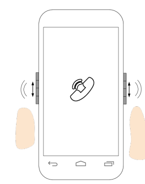

(b) Haptic notifications which tap the user's finger during an alert

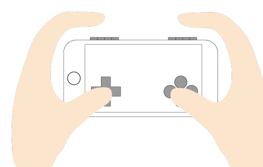

(c) Physical game controls

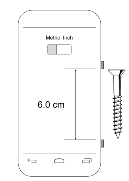

(d) Using a dynamic physical tool in the form of interactive calipers to measure the external world.

Figure 14: Example applications leveraging shiftIO.

rich haptic feedback. Here we describe a number of applications, see examples in Figure 14, primarily for the mobile device with integrated Reconfigurable RTEs, though many could be extended to a wearable band style device as well.

Context-Specific Controls

Physical controls have a number of advantages over touchscreen interactions in many applications. As an example interaction, when the user starts a game on the mobile device, two shoulder buttons emerge to become the interface elements (See Figure 15). The elements could also dynamically move in correspondence with game elements, such as acting as physical paddles in a game of Pong.

Self-Adjusting Interface Elements

Interface elements can adjust to accommodate new modes or states. For example, when the user opens a camera application, a physical shutter button appears, to allow the taking of photos without the need to locate a graphical button on the screen. Further, when the device is rotated, the button can adjust to ensure that the control always remains in the user's preferred position (e.g. top right). We can similarly have physical

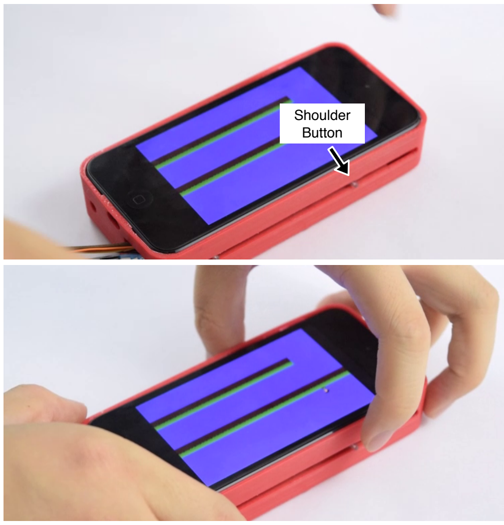

Figure 15: Two tactile elements appear as dynamic, physical shoulder buttons with touch input for a mobile video game.

scrollbars or other indicators which adjust to render the state of the application.

Increasing Effective Screen Space

Because of their small size, mobile devices have limited space for interaction. In particular, interface elements rendered on the screen can subtract from real estate otherwise used for display, since the display is necessarily occluded by the user's hands when they interact with on-screen controls. By creating controls on the sides of the device, we can free up the entire screen for display, and enable interactions without occluding the screen. For example, we can provide a video player which has a scrubber/playback slider on one side of the device, and volume controls on another side. This allows the video to play in full-screen without GUI elements. Figure 16 shows a scrollbar RTE that moves with the current document position.

Rich Haptic Notifications

There are numerous situations in which users cannot necessarily devote their visual attention to their mobile device, e.g., due to safety (e.g. while riding a bike), inappropriate for social reasons (such as at a dinner or in the cinema), or due to lighting conditions (strong sunlight). In such cases, we can convey notifications and other information haptically, such as representing unread notifications as the number of tactile elements lining the side of the device, or by altering their position, as shown in Figure 17.

LIMITATIONS AND FUTURE WORK

Though the use of magnetic actuation has numerous advantages, our current designs do have some limitations.

While shiftIO is significantly more power efficient than most traditional shape displays, power consumption could be further optimized to increase suitability for mobile and wearable devices. The thin, flexible micro-coil version uses power not only to move the RTEs, but also to keep them in a steady state. We envision adding a mechanical clutch to lock all RTEs in their current location at once, allowing us to effectively unpower all but the sensing subsystems while the RTEs are stationary. Rather than completely locking the elements, the clutch could provide a high friction force, holding the RTEs in place while still allowing users to reposition them for input. Another approach would be to add ferromagnetic material beneath the PCB, such that the RTE is passively attracted with a normal force, even when power is not supplied.

shiftIO is currently limited to low forces, which precludes more advanced haptic interactions with the ability to apply larger force to the user's fingers. The use of electro-permanent magnets [19], which are similar to our bistable magnets but can generate higher magnetic fields (at the cost of the ability of changing polarity), could allow us to generate stronger forces. In addition, with our current implementations, when a user places the device into a pocket, the RTEs would likely be dislodged from their target positions from strong contact forces.

Figure 16: A tactile element represents the current position while scrolling through a document. Touching the element bookmarks the position

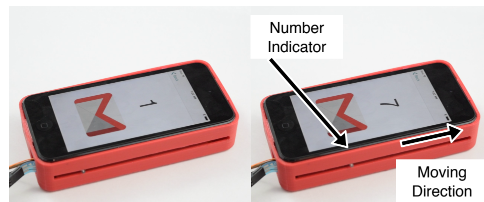

Figure 17: A tactile element changes position based on the number of e-mails available for subtle, tactile notifications.

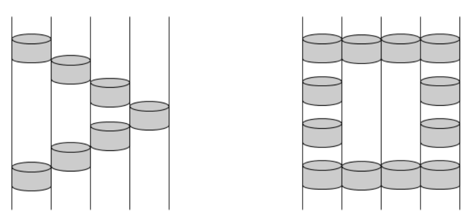

Figure 18: Future iterations could combine multiple rows of RTE tracks to allow for the display of more complex shapes, such as these playback controls.

By implementing closed-loop control, we could potentially detect these unintentional movements and return the RTEs to their previous positions when the resistance is removed.

As previously described, the neodymium magnets require a minimum separation to maintain stability, which prevents current applications from having several RTEs in close proximity. When RTEs approach a certain distance, they snap together and need to be mechanically separated as the system's dynamic magnetic fields are not strong enough to pull apart the strong N52 magnets. We are currently exploring magnetic shielding materials on the exterior of the RTEs. While our early efforts significantly decreased the minimum separation, shielding adds additional weight to the RTE, and thus additional tuning of the dimensions and materials is required to maintain the performance of the system. Future versions could also introduce a small linear actuator with a simple wedge to separate RTEs, thus allowing for different lengths of RTEs to be ejected. Currently, there is also a limited number of RTEs that can be stored in the "hidden" state, due to the minimum spacing required between RTEs. Further sophistication of the reservoir where the RTEs are stored is needed.

Magnetic fields external to shiftIO could interfere with normal operation, and the magnetic activity from shiftIO could also cause problems for other magnetically sensitive devices, such as a magnetometer or credit card. In simulations using our neodymium magnets, we find that the problematic distance at which demagnetization of a standard credit card becomes likely is far less than the minimum possible separation (due to the casing of the device). However, since we have not tested these effects directly, they remain an open concern.

An additional technical issue arises due to eddy currents generated within the neodymium magnet. As the element moves through a magnetic field, the induced current causes heating, which can become problematically hot over long periods of operation in our current prototype. In addition to damaging other internals within the device and possibly causing harm to the user, a RTE left in a high-temperature state for too long could lose its magnetization. This issue can be addressed by refining our choice of magnet. An RTE consisting of a laminated magnet or a magnet with a high electrical resistance would have significantly reduced eddy currents.

Furthermore, our use of a single soft potentiometer for registering input precludes us from recognizing multi-touch input. This means that multi-touch techniques (such as pinch input) require additional sensing. While we are currently using the system in an open loop configuration, we envision that a hall effect sensor array would improve performance and also make the system more robust to disturbances. Using a pressure-sensitive soft potentiometer, we hope to enable multi-stage touch interaction. For example, in a camera shutter application, a light touch on the RTE could initiate auto-focusing and illumination, while a solid press would then take the photo.

There are also limitations inherent to RTE interfaces in general, one of which is the inability to instantaneously "display" a RTE. Unlike the rendering of a pixel, the RTE takes an appreciable amount of time to travel to its intended location. In a 1D implementation this poses a critical path planning problem if there are heterogeneous RTEs. In addition, it could be hard for a user to discriminate between preparatory motion of an RTE into position and intentional motion for display. By better utilizing (or increasing) the thickness of the device, we can implement parallel rows for RTE display. With an added method of changing tracks, we could move RTEs into position along the hidden tracks, then bring them to the surface for interaction. In addition, this would create a 2D tactile display along the edges of the device, enabling the rendering of more complex shapes. For example, in a music player application, the play, fast-forward, and rewind functions could be rendered as tactile shapes, that could be recognized via the sense of touch, (see Figure 18).

CONCLUSION

We have presented a novel approach to mobile and wearable haptics in the introduction of the Reconfigurable Tactile Element Interface, which seeks to augment existing devices with dynamic physical controls and feedback, without altering the existing form factor. We have shown two possible implementations of such a system in the form of shiftIO, and discussed their technical tradeoffs and limitations. Finally, we have explored the design considerations for interactions of a RTE Interface, and presented a number of example applications which leverage the advantages of dynamically reconfigurable tactile elements on a mobile device.

ACKNOWLEDGMENTS

This work was funded in part by a Google Faculty Research Award. We thank Sungjune Jang for his large contribution to the engineering of the flexible PCB system. We thank David Christensen for his advice about electromagnetic drive systems and simulation.

REFERENCES

1. Kota Amano and Akio Yamamoto. 2012. Tangible interactions on a flat panel display using actuated paper sheets. In *Proceedings of the 2012 ACM international conference on Interactive tabletops and surfaces*. ACM, 351–354.

2. Olivier Bau, Ivan Poupyrev, Ali Israr, and Chris Harrison. 2010. TeslaTouch: electrovibration for touch surfaces. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*. ACM, 283–292.

3. Patrick Baudisch and Gerry Chu. 2009. Back-of-device Interaction Allows Creating Very Small Touch Devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 1923–1932. DOI: http://dx.doi.org/10.1145/1518701.1518995

4. Gábor Blaskó and Steven Feiner. 2004. Single-handed Interaction Techniques for Multiple Pressure-sensitive Strips. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems (CHI EA '04)*. ACM, New York, NY, USA, 1461–1464. DOI: http://dx.doi.org/10.1145/985921.986090

5. Stephen Brewster, Faraz Chohan, and Lorna Brown. 2007. Tactile Feedback for Mobile Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 159–162. DOI: http://dx.doi.org/10.1145/1240624.1240649
6. Alex Butler, Shahram Izadi, and Steve Hodges. 2008. SideSight: Multi-"Touch" Interaction Around Small Devices. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology (UIST '08)*. ACM, New York, NY, USA, 201–204. DOI: http://dx.doi.org/10.1145/1449715.1449746
7. David Cappelleri, Dimitrios Efthymiou, Ashesh Goswami, Nikolaos Vitoroulis, and Michael Zavlanos. 2014. Towards mobile microrobot swarms for additive micromanufacturing. *International Journal of Advanced Robotic Systems* 11 (2014).
8. Craig Michael Ciesla and Micah B Yairi. 2012. Tactus User interface system. (April 10 2012). US Patent 8,154,527.
9. Ronald S Fearing. 1996. A planar milli-robot system on an air bearing. In *Robotics Research*. Springer, 570–581.
10. John Hardy, Christian Weichel, Faisal Taher, John Vidler, and Jason Alexander. 2015. ShapeClip: Towards Rapid Prototyping with Shape-Changing Displays for Designers. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 19–28. DOI: http://dx.doi.org/10.1145/2702123.2702599
11. Chris Harrison and Scott E Hudson. 2009. Providing dynamically changeable physical buttons on a visual display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 299–308.
12. Fabian Hemmert, Gesche Joost, André Knörig, and Reto Wettach. 2008. Dynamic knobs: shape change as a means of interaction on a mobile phone. In *CHI'08 Extended Abstracts on Human Factors in Computing Systems*. ACM, 2309–2314.
13. David Holman, Andreas Hollatz, Amartya Banerjee, and Roel Vertegaal. 2013. Unifone: Designing for Auxiliary Finger Input in One-handed Mobile Interactions. In *Proceedings of the 7th International Conference on Tangible, Embedded and Embodied Interaction (TEI '13)*. ACM, New York, NY, USA, 177–184. DOI: http://dx.doi.org/10.1145/2460625.2460653
14. Sungjae Hwang, Myungwook Ahn, and Kwang-yun Wohn. 2013. MagGetz: customizable passive tangible controllers on and around conventional mobile devices. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*. ACM, 411–416.
15. Hiroshi Ishii and Brygg Ullmer. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. ACM, 234–241.
16. Sungjune Jang, Lawrence H. Kim, Kesler Tanner, Hiroshi Ishii, and Sean Follmer. 2016. Haptic Edge Display for Mobile Tactile Interaction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3706–3716. DOI: http://dx.doi.org/10.1145/2858036.2858264
17. Yvonne Jansen, Pierre Dragicevic, and Jean-Daniel Fekete. 2012. Tangible Remote Controllers for Wall-size Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 2865–2874. DOI: http://dx.doi.org/10.1145/2207676.2208691
18. Wuming Jing, Nicholas Pagano, and David J Cappelleri. 2013. A novel micro-scale magnetic tumbling microrobot. *Journal of Micro-Bio Robotics* 8, 1 (2013), 1–12.
19. Ara Nerses Knaian. 2010. *Electropermanent magnetic connectors and actuators: devices and their application in programmable matter*. Ph.D. Dissertation. Massachusetts Institute of Technology.
20. Mathieu Le Goc, Lawrence H. Kim, Ali Parsaei, Jean-Daniel Fekete, Pierre Dragicevic, and Sean Follmer. 2016. Zooids: Building Blocks for Swarm User Interfaces. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST '16)*. ACM, New York, NY, USA, 97–109. DOI: http://dx.doi.org/10.1145/2984511.2984547
21. Joseph Luk, Jerome Pasquero, Shannon Little, Karon MacLean, Vincent Levesque, and Vincent Hayward. 2006. A Role for Haptics in Mobile Interaction: Initial Design Using a Handheld Tactile Display Prototype. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*. ACM, New York, NY, USA, 171–180. DOI: http://dx.doi.org/10.1145/1124772.1124800
22. David Meeker. 2014. Finite Element Method Magnets. (2014). http://www.femm.info/wiki Accessed Aug. 2016.
23. Gian Pangaro, Dan Maynes-Aminzade, and Hiroshi Ishii. 2002. The Actuated Workbench: Computer-controlled Actuation in Tabletop Tangible Interfaces. In *Proceedings of the 15th Annual ACM Symposium on User Interface Software and Technology (UIST '02)*. ACM, New York, NY, USA, 181–190. DOI: http://dx.doi.org/10.1145/571985.572011
24. J. Pasquero, J. Luk, V. Levesque, Qi Wang, V. Hayward, and K.E. MacLean. 2007. Haptically Enabled Handheld Information Display With Distributed Tactile Transducer. *Multimedia, IEEE Transactions on* 9, 4 (June 2007), 746–753. DOI: http://dx.doi.org/10.1109/TMM.2007.895672
25. James Patten and Hiroshi Ishii. 2007. Mechanical Constraints As Computational Constraints in Tabletop Tangible Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 809–818. DOI: http://dx.doi.org/10.1145/1240624.1240746

26. Esben Warming Pedersen and Kasper Hornbæk. 2011. Tangible Bots: Interaction with Active Tangibles in Tabletop Interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '11)*. ACM, New York, NY, USA, 2975–2984. DOI: http://dx.doi.org/10.1145/1978942.1979384

27. Ron Pelrine. 1992. Magnetically levitated apparatus. (March 24 1992). US Patent 5,099,216.

28. Ron Pelrine, Annjoe Wong-Foy, Brian McCoy, Dennis Holeman, Rich Mahoney, Greg Myers, Jim Herson, and Tony Low. 2012. Diamagnetically levitated robots: An approach to massively parallel robotic systems with unusual motion properties. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*. IEEE, 739–744.

29. Ivan Poupyrev and Shigeaki Maruyama. 2003. Tactile interfaces for small touch screens. In *Proceedings of the 16th annual ACM symposium on User interface software and technology*. ACM, 217–220.

30. Simon Robinson, Céline Coutrix, Jennifer Pearson, Juan Rosso, Matheus Fernandes Torquato, Laurence Nigay, and Matt Jones. 2016. Emergeables: Deformable Displays for Continuous Eyes-Free Mobile Interaction. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 3793–3805. DOI: http://dx.doi.org/10.1145/2858036.2858097

31. Dan Rosenfeld, Michael Zawadzki, Jeremi Sudol, and Ken Perlin. 2004. Physical objects as bidirectional user interface elements. *Computer Graphics and Applications, IEEE* 24, 1 (2004), 44–49.

32. Evan Strasnick and Sean Follmer. 2016. Applications of Switchable Permanent Magnetic Actuators in Shape Change and Tactile Display. In *Proceedings of the 29th Annual ACM Symposium on User Interface Software and Technology (UIST '16)*. ACM, Tokyo, Japan, 4. DOI: http://dx.doi.org/10.1145/2984751.2985728

33. John Tiab and Kasper Hornbæk. 2016. Understanding Affordance, System State, and Feedback in Shape-Changing Buttons. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM, 2752–2763.

34. Eone Time. 2016. Eone Bradley tactile watch. (2016). https://www.eone-time.com

35. Nicolas Villar and Hans Gellersen. 2007. A malleable control structure for softwired user interfaces. In *Proceedings of the 1st international conference on Tangible and embedded interaction*. ACM, 49–56.

36. Akira Wakita, Akito Nakano, and Nobuhiro Kobayashi. 2011. Programmable blobs: a rheologic interface for organic shape design. In *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction*. ACM, 273–276.

37. Malte Weiss, Florian Schwarz, Simon Jakubowski, and Jan Borchers. 2010. Madgets: Actuating Widgets on Interactive Tabletops. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, New York, NY, USA, 293–302. DOI: http://dx.doi.org/10.1145/1866029.1866075

38. Malte Weiss, Chat Wacharamanotham, Simon Voelker, and Jan Borchers. 2011. FingerFlux: near-surface haptic feedback on tabletops. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. ACM, 615–620.