TeslaTouch: Electrovibration for Touch Surfaces

Olivier Bau[1,2], Ivan Poupyrev[1], Ali Israr[1], Chris Harrison[1,3]

[1] Disney Research, Pittsburgh
4615 Forbes Avenue
Pittsburgh, PA 15213
{bau, ivan.poupyrev, israr}
@disneyresearch.com

[2] In|Situ|, INRIA Saclay,
Building 490,
Université Paris-Sud,
91405 Orsay, France

[3] Human-Computer Interaction Institute
Carnegie Mellon University
5000 Forbes Avenue,
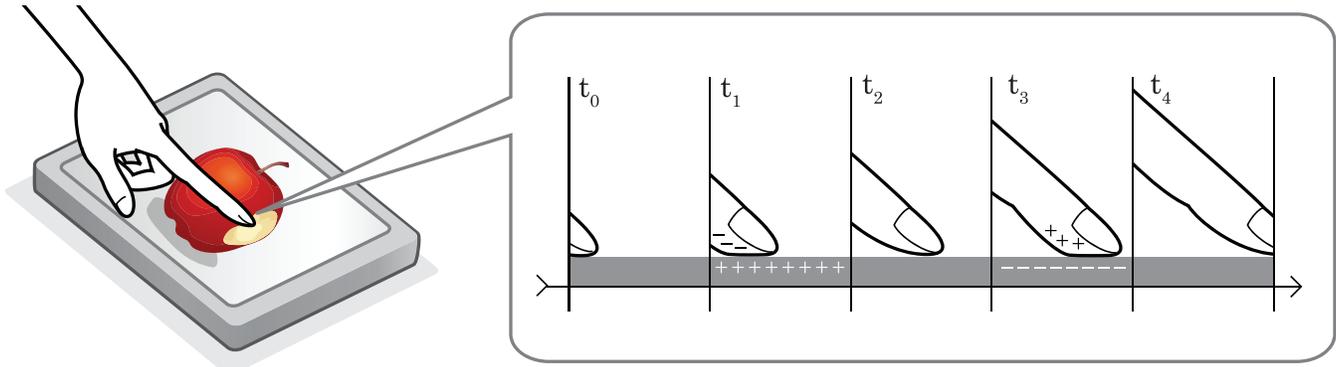Pittsburgh, PA 15213
chris.harrison@cs.cmu.edu

Figure 1: TeslaTouch uses electrovibration to control electrostatic friction between a touch surface and the user's finger.

ABSTRACT

We present a new technology for enhancing touch interfaces with tactile feedback. The proposed technology is based on the electrovibration principle, does not use any moving parts and provides a wide range of tactile feedback sensations to fingers moving across a touch surface. When combined with an interactive display and touch input, it enables the design of a wide variety of interfaces that allow the user to feel virtual elements through touch. We present the principles of operation and an implementation of the technology. We also report the results of three controlled psychophysical experiments and a subjective user evaluation that describe and characterize users' perception of this technology. We conclude with an exploration of the design space of tactile touch screens using two comparable setups, one based on electrovibration and another on mechanical vibrotactile actuation.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces - Graphical user interfaces, Input devices and strategies, Haptic I/O.

General terms: Design, Measurement, Human Factors.

Keywords: Tactile feedback, touch screens, multitouch.

INTRODUCTION

Interest in designing and investigating haptic interfaces for touch-based interactive systems has been rapidly growing. This interest is partially fueled by the popularity of touch-based interfaces, both in research and end-user communities. Despite their popularity, a major problem with touch interfaces is the lack of dynamic tactile feedback. Indeed, as observed by Buxton as early as 1985 [6], a lack of haptic feedback 1) decreases the realism of visual environments, 2) breaks the metaphor of direct interaction, and 3) reduces interface efficiency, because the user can not rely on familiar haptic cues for accomplishing even the most basic interaction tasks.

Most previous work on designing tactile interfaces for interactive touch surfaces falls into two categories. First, the touch surface itself can be actuated with various electromechanical actuators such as piezoelectric bending motors, voice coils, and solenoids [10, 27]. The actuation can be designed to create surface motion either in the normal [27] or lateral directions [4]. Second, the tools used to interact with a surface, such as pens, can be enhanced with mechanical actuation [9, 19].

In this paper, we present an alternative approach for creating tactile interfaces for touch surfaces that does not use any form of mechanical actuation. Instead, the proposed technique exploits the principle of *electrovibration*, which allows us to create a broad range of tactile sensations by controlling *electrostatic friction* between an instrumented touch surface and the user's fingers. When combined with an input-capable interactive display, it enables a wide variety of interactions augmented with tactile feedback.

Tactile feedback based on electrovibration has several compelling properties. It is fast, low-powered, dynamic, and can

be used in a wide range of interaction scenarios and applications, including multitouch interfaces. Our system demonstrates an exceptionally broad bandwidth and uniformity of response across a wide range of frequencies and amplitudes. Furthermore, the technology is highly scalable and can be used efficiently on touch surfaces of any size, shape and configuration, including large interactive tables, hand-held mobile devices, as well as curved, flexible and irregular touch surfaces (e.g. [3, 29]). Lastly, because our design does not have any moving parts, it can be easily added to existing devices with minimal physical modification.

The contributions of this paper are four-fold. 1) We present the principles and implementation of electrovibration-based tactile feedback for touch surfaces. 2) We report the results of three controlled psychophysical experiments and a subjective user evaluation, which describe and characterize users' perception of this technology. 3) We analyze and compare our design to traditional mechanical vibrotactile displays and highlight their relative advantages and disadvantages. 4) We explore the interaction design space.

BACKGROUND AND RELATED WORK
The effect of electrovibration was discovered in 1954 by accident. Mallinckrodt et al. [23] reported that dragging a dry finger over a conductive surface covered with a thin insulating layer and excited with a 110 V signal, created a characteristic "rubbery" feeling. They explained this effect by suggesting that the insulating layer of dry outer skin formed the dielectric layer of a capacitor, in which conductive surfaces and fluids in the finger's tissue are the two opposing plates. When alternating voltage is applied to the conductive surface, an intermittent attraction force develops between the finger and conductive surface. While this force is too weak to be perceived when the finger is static, it modulates *friction* between the surface and skin of the moving hand, creating the rubbery sensation. This effect was named "electrovibration" [32].

It is important to highlight the differences between electrocutaneous, electrostatic, and electrovibration tactile actuation. *Electrocutaneous* displays stimulate tactile receptors in human fingers with electric charge passing through the skin [18]. In contrast, there is no passing charge in electrovibration: the charge in the finger is induced by a charge moving on a conductive surface (Figure 1). Furthermore, unlike electrocutaneous tactile feedback, where current is directly stimulating the nerve endings, stimulation with electrovibration is *mechanical*, created by a periodic electrostatic force deforming the skin of the sliding finger.

In the *electrostatic* approach, a user is manipulating an intermediate object, such as a piece of aluminum foil [37], over an electrode pattern. A periodic signal applied to this pattern creates weak electrostatic attraction between an object and an electrode, which is perceived as vibration when the object is moved by the user's finger. The tactile sensation, therefore, is created indirectly: the vibration induced by electrostatic force on an object is transferred to the touching human finger. In case of electrovibration, no in-termediate elements are required; the tactile sensation is created by directly actuating the fingers.

Although electrovibration was discovered in 1954, there was no attempt to use it for haptic applications until 1970, when Strong [32] proposed a tactile display consisting of an array of pins insulated with a thin layer of dielectric. Different voltages were applied to different pins so that users could feel various tactile shapes. A similar configuration was reported by Tang and Beebe [33], where the pin array was created using lithographic microfabrication, resulting in a thin and durable tactile display.

Similar to mechanical vibration, electrovibration is not a technology per se, but a category of tactile sensation that can be generated in many different ways. In all previous work (e.g. [7, 32, 33]), electrovibration was delivered using opaque patterns of electrodes, such as the dense arrays of metal pins described earlier, which makes combination with tracking and display technologies challenging. Furthermore, the technique does not scale to large surfaces. In Tesla-Touch, we deliver electrovibration via a transparent electrode on a clear substrate. This allows electrovibration to be used with a wide variety of display and input technologies.

E-Sense technology from Senseg corporation [31] produces tactile feedback by charging a conductive film attached to the touch panel. It has been developed in parallel to Tesla-Touch[1] and is based on the same physical principles. However, the technology has not been released on the market, nor have implementation details been disclosed. Therefore, it cannot be reproduced and compared to TeslaTouch.

In general, adding tactile feedback to touch interfaces has been challenging. One research direction has been the design of tactile feedback for touch interfaces on small hand-held devices by mechanically vibrating the entire touch surface with piezoelectric actuators, voice coils and other actuators [4, 10, 27]. With low frequency vibrations, a simple "click" sensation can be simulated [27]. When ultrasonic frequencies are used [4, 35], a sensation of variable friction between the finger and surface can be created.

A major challenge in using mechanical actuation with mobile touch surfaces is the difficulty of creating actuators that fit into mobile devices and produce sufficient force to displace the touch surface. Creating tactile interfaces for large touch screens [30] such as interactive kiosks and desktop computers allows for larger actuators. Larger actuated surfaces, however, begin to behave as a flexible membrane instead of a rigid plate. Complex mechanical deformations occur when larger plates are actuated, making it difficult to predictably control tactile sensation or even provide enough power for actuation.

An alternative approach to actuation of the touch surface is to decouple the tactile and visual displays. In the case of mobile devices, tactile feedback can be provided by vibrating the backside of the device, stimulating the holding hand

---

[1] Preliminary explorations of basic TeslaTouch technology started when the second author was at Sony CSL Inc. [28]
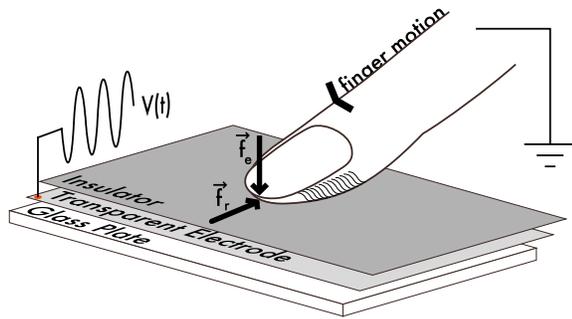
Figure 2: TeslaTouch operating principle.

[5]. Alternatively, it is possible to embed localized tactile actuators into the body of a mobile device [22] or into tools used in conjunction with touch interfaces [9, 19]. This approach, however, breaks the metaphor of direct interaction, requires external devices and still does not solve the problem of developing tactile feedback for large surfaces.

TESLATOUCH

To investigate the tactile properties of our approach, we combined it with a specific input-tracking technique: a diffuse illumination-based multitouch setup [24]. However, the fundamental technology is generic and can be easily extended to many input and display technologies.

TeslaTouch Tactile Feedback Apparatus

We used a 3M Microtouch panel [1] originally designed for capacitive-based touch sensing. It is composed of a transparent electrode sheet applied onto a glass plate coated with an insulator layer (Figure 2). We then excite the transparent electrode with a periodic electrical signal $V(t)$ applied to connectors normally used by the position-sensing driver. When an input signal of sufficient amplitude is provided, an electrically induced attractive force $\vec{f_e}$ develops between a sliding finger and the underlying electrode, increasing the dynamic friction $\vec{f_r}$ between the finger and the panel surface (Figure 2). Because the amplitude of $\vec{f_e}$ varies with the signal amplitude, changes in friction $\vec{f_r}$ will also be periodic, resulting in periodic skin deformations as the finger slides on the panel. These deformations are perceived as vibration or friction and can be controlled by modulating the amplitude and frequency of the applied signal. Note that only digits in motion perceive this effect.

The tactile signal in our current implementation is generated by the Pure Data sound programming environment, outputted by a standard sound card and amplified from ~1.5 Volts peak-to-peak (Vpp) to 5 Vpp using an operational amplifier. It is then further amplified up to a maximum of 120 Vpp signal with a power transformer (Figure 3). In our current implementation, we use pure sinusoidal waveforms. However, other waveforms can be used, e.g. square or triangular [17]. Importantly, the input signal is uniformly propagated across the conductive layer of the plate; therefore, the resulting tactile sensation is *spatially uniform*.

Grounding Strategies

We instrumented users with a *return ground path* for the signal [32]. We found that, although our bodies provide a natural link to the ground, creating a direct ground connec-
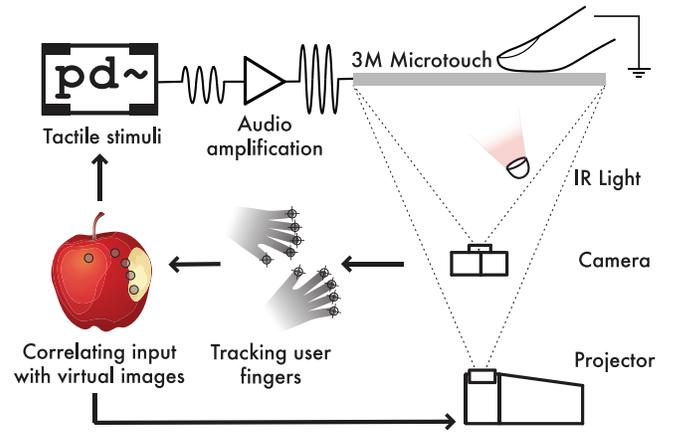

Figure 3: Integrating TeslaTouch into multitouch tabletop interactive surface.

tion significantly increased the intensity of the tactile sensation. Without such grounding, the voltage must be increased to provide the same intensity of sensation. This grounding can be achieved by wearing a simple ground electrode, e.g. an antistatic wristband. Users can also sit or stand on a grounded pad [8]. In the case of mobile devices, the backside of the enclosure, which contacts the user when grasped, could be used as the ground.

TeslaTouch Safety

The critical factor for safe operation is current, rather than voltage. We emphasize that there is no actual *charge* passing through the skin and the amount of induced current flowing through the user's hand is negligible. The current supplied to the TelsaTouch panel is limited to 0.5 mA, which is considered safe for humans [36]. Current limitation is defined by the power rating of the operational amplifier used in the driving circuit. In fact, users experience the same amount of current while using conventional capacitive touch panels [1]. To further protect the user, we use a simple current limiting circuit.

Instrumenting Touch Surfaces With TeslaTouch

For our initial implementation, we chose to implement a TeslaTouch tactile display for multitouch interactive tabletop surfaces [24] (Figure 3). The capacitive touch panel was used as a projection and input surface. An additional diffuser plane was installed behind the panel; a projector was used to render graphical content. To capture the user input, the panel was illuminated from behind with infrared illuminators. An infrared camera captured reflections of user fingers touching the surface. We used the open source CCV project (http://ccv.nuigroup.com) for multitouch tracking at 60 frames per second. Finger positions were sent using the TUIO protocol (http://tuio.org) to the main application responsible for controlling interactive features, visual display, and tactile output. The latter was achieved by sending frequency and amplitude data over UDP to the Pure Data sound-programming environment. All software runs on a single iMac computer in real-time.

The implementation described above is scalable and can be adapted to other input techniques, including frustrated in-

ternal reflection [12] and surface acoustic tracking, among others. It can be easily extended, modified and applied to any surface or device. Indeed, since there is no mechanical motion, almost any object can be instrumented with elec-trovibration-based tactile feedback. The electrodes can be transparent or opaque, be painted on curved and irregular surfaces, and added to any display, hand tool, or appliance.

HUMAN FACTORS OF TESLATOUCH

Designing applications for TeslaTouch requires understand-ing the basic human factors and usability characteristics of electrovibration. Will users like it? Will it feel more like friction or vibration? What are the lowest signal levels that users can feel? How well can users differentiate changes in signal frequencies and amplitudes? Answering these ques-tions is important for designing effective tactile interfaces based on electrovibration. We use this section to discuss some of these issues.

Subjective Evaluation of TeslaTouch

Unlike mechanical vibrations, electrovibrations are not normally experienced in everyday life. Therefore, we con-ducted subjective evaluations to better understand how us-ers interpret the tactile sensations produced by TeslaTouch.

Procedure and Participants

Ten participants (min. age 25, max. 40) felt four Tesla-Touch textures produced by four frequency-amplitude com-binations: 80 Hz and 400 Hz each at 80 and 115 Vpp. These frequencies were perceptually distinct as they represent two ends of our test frequency range.

For each texture, participants filled out a three-section ques-tionnaire. The first section asked participants to describe each sensation in their own words. The second section in-troduced 11 nouns (e.g., fur, silk, jeans, sand paper, skin) and asked participants to select nouns that described the tactile sensations as closely as possible. In the final section, participants rated different dimensions of sensations on a five-point Likert scale (e.g. from *smooth* to *sticky*). Partici-pants could experience each texture for as long as they wished; sessions took between 20 and 35 minutes.

Results

As we expected, low frequency stimuli were perceived as rougher compared to high frequencies. They were often likened to "wood" and "bumpy leather", versus "paper" and "a painted wall" for higher frequency stimuli.

The effect of *amplitude* depended on stimuli frequency. For high frequency textures (e.g. 400 Hz) an increase of ampli-tude increased perceived *smoothness* of tactile sensations. Indeed, while at 80 Vpp textures were mostly compared to "cement surface" and "cheap paper"; at 115 Vpp they were compared to "paper" or "a painted wall". Some participants explicitly pointed out this increase in perceived smoothness. For example, one participant described the 80 Vpp stimuli as feeling like a "painted wall", and later referred to the 115 Vpp texture as a "smoother painted wall".

At low frequencies (e.g. 80 Hz), an increase in stimuli am-plitude heightened the perception of *stickiness*. While some participants referred explicitly to a "sticky" sensation, oth-
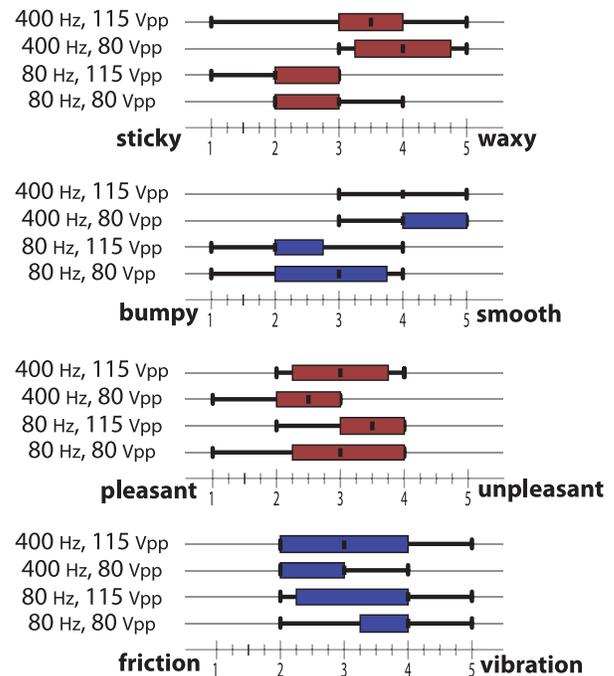

Figure 4: Ratings of stickiness, smoothness, pleasure and level of friction vs. vibration.

ers compared the sensation to that of touching a "motorcy-cle handle" or "rubber". Other participants associated vis-cosity with this type of texture; one participant compared his experience to "running fingers through viscous liquid".

In the final section of the study, subjects rated tactile sensa-tions with a set of five-point Likert scales on such proper-ties as *stickiness* (sticky to waxy), *fineness* (coarse to fine), *smoothness* (bumpy to smooth), *pleasantness* (pleasant to unpleasant), *friction versus vibration* and *gentleness* (gentle to brisk).

These results are presented in Figure 4 and generally agree with our earlier observations. Indeed, frequency had a sig-nificant effect on perception of *stickiness* ($F(1,9)=18.45$; $p<0.01$): low frequencies were associated with *sticky* sensa-tions, while high frequencies felt *waxy*. Similarly, high fre-quency stimuli were rated significantly *smoother* than low frequency stimuli (mean ratings of 4.1 and 2.55 respec-tively, $F(1,9) = 67.04$; $p<0.01$). Stimuli with high amplitude were rated less pleasant than lower amplitude stimuli, with a mean rating of 3.2 versus 2.7 ($F(1,9) = 5.16$; $p<0.05$).

When describing tactile sensations produced by Tesla-Touch, participants often described them as a combination of *vibration* and *friction* sensations. High frequency stimuli were rated as more related to friction than low frequency stimuli, which were related more to vibration (mean ratings of 2.9 and 3.6 respectively). However, this effect was not statistically significant. This seeming duality of tactile sen-sation elicited by TeslaTouch is an interesting direction for future experimentation.

Psychophysics of TeslaTouch

In this section, we investigate perception-based characteris-tics of electrovibration. These include absolute detection
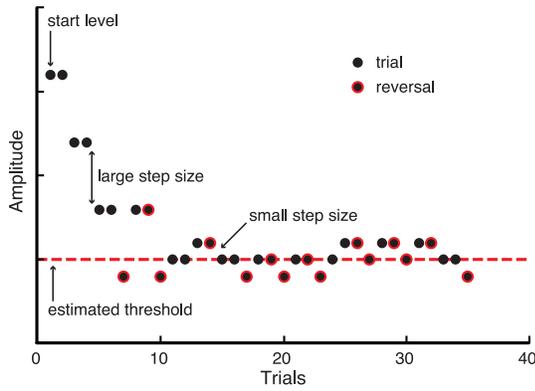
Figure 5: Example of a trial session that follows our adaptive procedure.

Figure 6: Experimental set up used to test absolute detection threshold (left) and JND thresholds (right).

thresholds and frequency and amplitude discrimination thresholds. The *absolute detection threshold* is an important psychophysical measure that defines the baseline of human sensitivity. In the case of electrovibration it is the minimum voltage amplitude that creates a barely detectable sensation at a specific frequency. Voltages below the detection threshold are not usable in creating haptic sensations.

The *amplitude* and *frequency discrimination thresholds* are typically referred to as just-noticeable-differences (JNDs), which are the smallest detectable differences between two stimuli. The detection and discrimination thresholds together form a set of fundamental measures that describe the dynamic range and processing capabilities of electrovibration sensations. These measures can be used to design interfaces and applications using TeslaTouch.

Methods

Detection and *discrimination thresholds* were estimated for five frequencies equally spaced on a logarithmic scale: 80, 120, 180, 270, and 400 Hz. The order of frequencies was randomized to control for order effects. We employed a widely used one-up/two-down adaptive staircase procedure [21]. The advantage of this procedure is that it allows accurate estimation of detection and discrimination thresholds with relatively small number of trials [20].

The *absolute detection thresholds* were estimated by using a two-alternative forced-choice paradigm [20]. A touch screen was split into two areas marked as A and B (Figure 6); one of the areas had a tactile stimulus, while the other had none. In each trial, the stimulus was randomly assigned to one of the two areas. The participant's task was to determine which area provided a tactile sensation.

Each session started with the stimulus amplitude sufficiently higher than the anticipated detection threshold at the corresponding frequency. The voltage amplitude was then reduced by 1 dB if the participant had made two consecutively correct responses. When the user made an incorrect response, the voltage was *increased* by 1 dB, making tactile stimulus more prominent. A change from decreasing to increasing voltage amplitudes, and vice versa, is referred to as *reversal*. After the first three reversals the step size of the
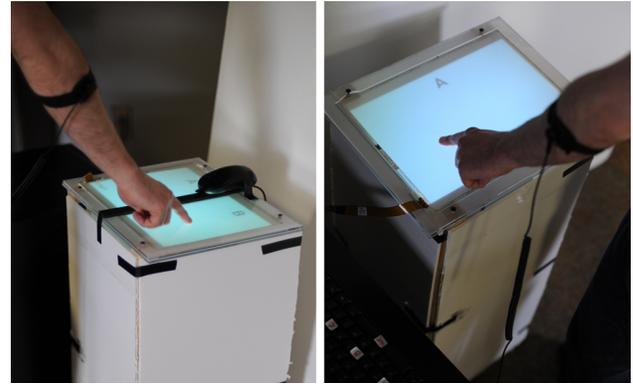
voltage change was reduced to 0.25 dB. The initially large 1 dB step size ensured faster convergence of amplitude towards the threshold level and the following smaller 0.25 dB step size guaranteed fine resolution of the threshold estimation. The session was terminated after 12 reversals at the 0.25 dB step size and the average amplitude from the last 12 reversals was taken as an estimate of the threshold level. Figure 5 illustrates this procedure.

The *frequency* and *amplitude discrimination* thresholds (JNDs) were determined in a similar manner. The stimulus set consisted of the same five reference frequencies at amplitude levels 15 dB above the detection threshold. JND values were estimated using a three-alternative forced-choice paradigm [20]. In each trial, three tactile stimuli were presented one after another. Participants were requested to identify *test stimuli*, which was different from the two identical *reference stimuli*. In amplitude discrimination experiments, the amplitude of test stimuli A_{test} differed from the amplitude of reference stimuli A_{ref} by a variable increment: $A_{\text{test}} = A_{ref} + \Delta A$, and in frequency discrimination experiments, $F_{test} = F_{ref} + \Delta F$. The order of test and reference stimuli was randomized in each trial.

At the start of the experimental session the test and reference stimuli were selected to be easily discernable, i.e. ΔA and ΔF were well above the anticipated JNDs. Two consecutively correct responses decreased and one incorrect response increased ΔF by factor of 1.58 and ΔA by 1 dB for the first three reversals and by 1.12 and 0.25 dB for the rest of the session. The session was terminated after 12 reversals at the smaller step size. The average of ΔA and ΔF from the last 12 reversals were then taken as JND estimates.

Experimental apparatus and procedure

The experimental apparatus is shown in Figure 6. Participants stood in front of the interactive touch table instrumented with TeslaTouch tactile feedback (touch panel size 316×254 mm). They were requested to wear an electrostatic ground strap on their dominant forearm and slide the pad of their index finger on the interactive surface.

All participants completed detection threshold experiments before discrimination threshold experiments. In the *absolute*

detection threshold experiments, participants were presented with two equally sized areas marked with letters A and B separated by a cardboard piece (Figure 6). Participants had eight seconds to compare areas A and B and respond by clicking a mouse button. In *discrimination thresholds* experiments, three screens were presented one after another marked with letters A, B and C. Participants had as much time as needed to feel tactile sensations on each screen. They progressed to the next screen by pressing the spacebar and were not allowed to return to the previous screen. After finishing all three screens, participants were prompted to select one that was different from the other two by pressing marked keys on the keyboard.

Participants
Ten right-handed participants (9 male, mean age 30 years old) took part in the detection threshold experiments. They conducted between 50 and 100 trials for each of the five reference frequencies. Each session lasted no more than 15 minutes; the total experiment time for each subject was 45-60 minutes. Seven right-handed participants (all male, mean age 30 years old) were tested in frequency and amplitude discrimination experiments. Each session lasted between 6 and 10 minutes and consisted of approximately 35-70 trials. The total time to complete both frequency and amplitude discrimination experiments was between 60 and 100 minutes for each participant.

Results
The detection and discrimination thresholds were analyzed across frequencies using repeated measures ANOVA with Greenhouse-Geisser correction for univariate analysis. A null hypothesis of significant effect was rejected if the resulting p-value was less than $\alpha = 0.05$.

Absolute Detection Thresholds
The absolute detection thresholds for five reference frequencies are shown in Figure 7. The thresholds are defined in "dB re 1 V peak" units computed as $20 \times \log_{10}(A)$ where A is the signal amplitude in Volts. Using this unit is a standard practice in psychophysical experiments due to linearity of human perception in logarithmic scale [15, 34]. For comparison, a force detection threshold curve [15] is also plotted in Figure 7. There was a statistically significant effect of frequency on the threshold levels ($F(4,36)=12.8$;

p<0.001), indicating that the threshold levels depend on the stimulus frequency.

Frequency Discrimination Thresholds
Frequency JNDs for each of the reference frequency are presented in Figure 8. The JNDs are defined as percentage of reference frequencies, i.e. for 80 Hz reference, the JND was found to be 25% or 20 Hz. Relative units are commonly used in JND experiments due to Weber's law that states that the ratio of the discrimination threshold to stimulus intensity is constant [11].

The effect of frequency on JND was statistically significant ($F(4,24)=6.46$; p<0.01). Post-hoc comparison divided the frequency range in to two groups: 80 Hz, 120 Hz, 180 Hz and 270 Hz, 400 Hz. The average JNDs were 25% and 12% for lower and higher frequency groups, respectively.

Amplitude Discrimination Thresholds
The amplitude JNDs are presented in Figure 9 as a function of reference frequency. The amplitude JNDs are also defined in dB units relative to the reference voltage. The ANOVA analysis failed to show significant effect of frequency on the amplitude JND ($F(4,24)=0.43$; p=0.79), indicating that the JND of 1.16 dB remains constant across all tested frequencies, thus obeying Weber's law.

Discussion
The detection threshold levels for electrovibrations closely coincide with the force detection threshold levels for sinusoidal stimulus determined earlier by Israr et al. [15] (see Figure 7). The plot suggests that sensations created with TeslaTouch are closely related to perception of forces lateral to the skin. The relation between electrovibration voltages and perceived forces might not be linear; the exact relation will be determined in future studies.

The detection threshold levels shown in Figure 7 provide important guidelines for designing tactile interfaces using electrovibration. For example, they inform the designer that at each frequency the applied voltage must be above the corresponding detection threshold level in order to provide a tactile sensation that a user can perceive. They also allow optimizing power requirements. For example, at 400 Hz the tactile signal would create an easily discernable tactile sensation at 18 dB re 1 V level or 16 Vpp. On the other hand, at 180 Hz the voltage threshold level is half of that, requir-
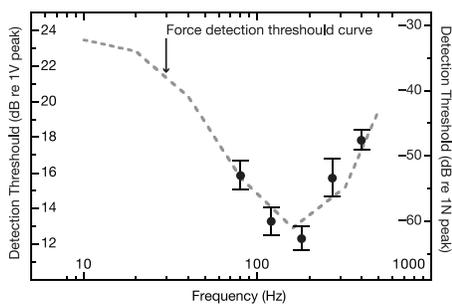

Figure 7: Mean detection threshold of electrovibrations with standard error bars (left axis) and force detection threshold curve from [15] (right axis).
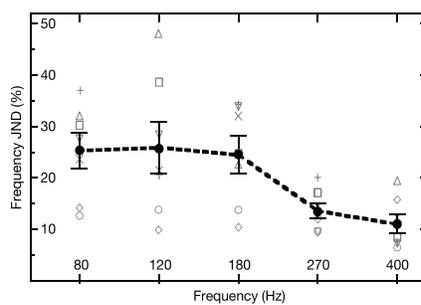
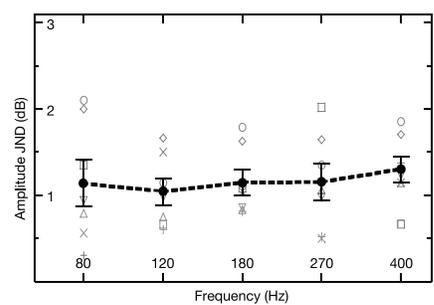Figure 8: Frequency JNDs for each participant and average values with standard error bars.

Figure 9: Amplitude JNDs for each participant and average values with standard error bars.

ing significantly less power (12 dB re 1 V peak or 8 Vpp). Therefore, tactile feedback can be optimized to require less power, which is especially important for mobile devices.

The average *frequency* JND varied from 11% at 400 Hz to 25% at 120 Hz which was similar to 13% – 38% thresholds determined for pure mechanical vibrations [16]. The average *amplitude* JND was 1.16 dB and constant across all frequencies thus following Weber's law. It was slightly lower than previously reported JNDs, which were in a 1.5 – 2.5 dB range [16, 34].

The frequency and amplitude discrimination thresholds describe the resolution of human perception: they determine the *granularity* of tactile sensations that can be used in designing interfaces. For example, if designers want to create two distinct tactile sensations, they must make sure that the amplitude of voltages for each sensation are at least 1.16 dB apart for the user to be able to differentiate them. Similar considerations also apply for frequency of stimuli.

ELECTROVIBRATION VS. MECHANICAL STIMULATION

TeslaTouch offers several significant advantages over conventional mechanical vibrotactile actuation on touch screens. In addition to motivating our approach, this section also underscores the unique interaction opportunities that TeslaTouch enables. We discuss how these properties can be employed for interactive purposes in the next section.

To inform this comparison, we built two identical and fully functional interactive surfaces, one equipped with Tesla-Touch and another with mechanical actuation provided by a pair of voice coil actuators. To ensure the highest degree of comparability, both were built on the same chassis with the same sized screen. They both employ the same diffused illumination multitouch optical tracking [24] and can run applications interchangeably.

The Effect of Mechanical Motion

The *absence of mechanical motion* is the most immediate difference between TeslaTouch and conventional mechanical actuation. This feature has several notable implications.

Spatial Uniformity of Tactile Feedback

Regardless of the type of material, any plane of material will flex when actuated. This problem is exacerbated when the plate is large and actuation forces are applied on its periphery, which is common when designing tactile feedback for touch surfaces. Consequently, not only are vibrotactile solutions not feasible for large interactive surfaces, but even for small screen sizes, different parts of the screens would have different magnitudes of physical displacement and, therefore, different tactile sensations (see Figure 10). Electrovibration, on the other hand, does not suffer from this effect as electrical potential is evenly and instantaneously distributed over the entire plate. The tactile feedback in TeslaTouch is uniform across surfaces of *any size*.

Attenuation of Tactile Sensations at Different Frequencies

When a periodic force vibrates a plate of material, as is the case for vibrotactile displays, the plate spring properties are combined with dampening, which is inherent due to attachment of the plate to an enclosure or chassis, and to-
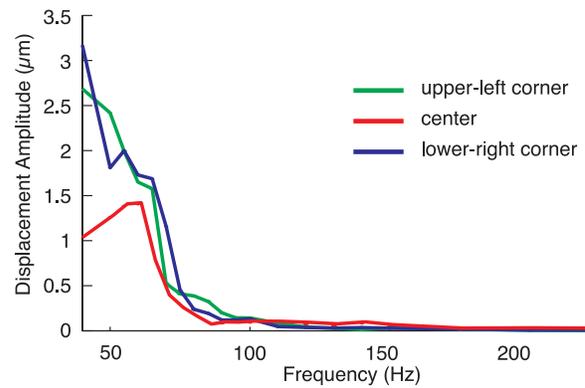

Figure 10; Acceleration data collected at three locations spaced along the diagonal of our mechanically actuated tactile screen.

gether they result in a *highly attenuated frequency response* of the plate. As a result, for a signal of the same amplitude, the mechanical displacement of the plate will be different at different frequencies, peaking close to its resonant mechanical frequency and then dramatically decreasing (Figure 10.) These complex signal attenuations make it essentially impossible to engineer a flat response – even software amplitude correction cannot hope to counter these laws of physics. Because electrovibration requires no moving parts, it suffers from neither of these effects.

Magnitude of Tactile Sensation

Traditional vibrotactile feedback can deliver very strong forces to the user's fingers with current electromagnetic or piezoelectric actuators. TeslaTouch delivers more subtle tactile experiences. However, the tactile effect is readily and immediately apparent, similar to how sand paper is clearly different from glass. Moreover, our psychophysics experiments show that participants are able to detect the tactile sensation with as little as 8 Vpp, which is less than 7% of our total voltage output. Additionally, although our current implementation tops out at ~120 Vpp, higher voltages could be used, both safely and comfortably, as demonstrated in famous electrostatic experiments by Nikola Tesla. For comparison, an electrostatic "spark" received from a household doorknob is in the order of thousands of volts and yet remains fairly harmless.

Noise

Byproduct noise is a serious consideration when designing end-user interactive systems. We accept that our kitchen blenders are noisy, but we use them rarely, and then only briefly. This level of noise would not be acceptable in a computing device we hope to use for extended period of time. Unfortunately, physical vibrations are often noisy, e.g. consider a mobile phone vibrating on a table. Compounding this problem is the fact that interactive surfaces tend to have large surface areas, which displace a considerable volume of air, essentially turning their screens into speakers. Because there is no physical motion in our TeslaTouch apparatus, it is entirely silent.
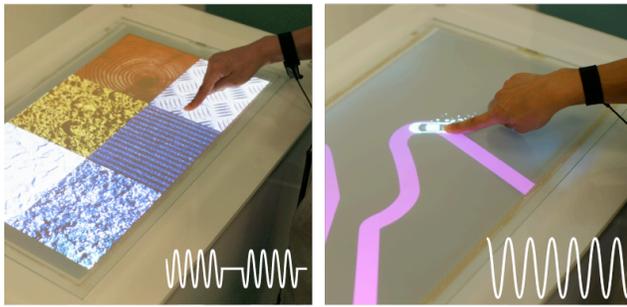
Figure 11: Left: different textures produce different sensations, e.g. simulated corduroy. Right: a racing track where friction increases as the car "squeaks" around corners.

Figure 12: A visual star field in concert with a tactile layer conveying radiation intensity.

Reliability

Moving parts naturally wear over time, which alters their performance characteristics and may eventually lead to failure. In addition, the vibrating screen must be separated from the enclosure with a small seam to accommodate movement, which allows dust, liquid and other debris inside the device. Sealing this seam, however, dampens vibrations, which decreases the intensity of tactile feedback. None of these issues are relevant in the case of TeslaTouch.

Feedback Localization and Multitouch

Vibrotactile actuation delivers tactile feedback by displacing the entire surface. As a result, all fingers resting on the surface will be stimulated and any physical object located on the surface is likely to chatter and move around, which is less favorable. In general, there is no way to localize tactile feedback to particular digits when vibrotactile feedback is used with interactive surfaces.

Although TeslaTouch can only provide one tactile signal to the entire surface, *only moving fingers* feel the tactile feedback. By carefully designing interactive sequences so that out of all fingers touching the surface only one moves at a time, we can create an illusion of localized tactile feedback. This approach enables the design of tactile feedback solutions for multitouch scenarios, leading to new interaction opportunities, which we discuss next.

DESIGNING INTERACTIONS WITH TESLATOUCH

The previous section discussed key operational differences between vibrotactile screens and TeslaTouch. We now draw upon these attributes to explore the design space of tactile interactions on touch screens. While many are supported by both approaches, we also identify interactions unique to each method. We include figures from many demo applications that we created as part of this design exercise.

General Applications

There are many applications that can be implemented by both actuation approaches. However, the superior frequency range and uniform response characteristics of TeslaTouch yield more accurate tactile representations and allow for richer user experiences. Regardless, we treat both techniques as interchangeable in this subsection.

Simulations

This class of applications includes such tactile effects as textures for virtual objects, simulation of friction between objects or objects and a virtual surface, and activities like painting and drawing, where tools are manipulated on top of a canvas (Figure 11). There has been considerable work in this domain, e.g. [30].

Tactile Information Layers

Tactile feedback on touch screens allows for non-visual information layers. For example, a visual image of a star field could be supplemented with a "tactile image" of radiation intensity, felt by fingers running over the areas of interest (Figure 12). The tactile channel can be dynamic in both amplitude and frequency, potentially offering two additional channels of information.

GUI Widgets with Tactile Feedback

There are many interesting avenues for infusing conventional GUI elements with tactile feedback [2, 13, 25, 30]. For example, sliders can report their drag extent by changing the tactile feedback frequency. Similarly, a user could run their fingers over a list of emails to sense those that are new or with the highest priority. There are numerous other interaction design ideas that can be explored in this area.

Supporting Direct Manipulation

Direct manipulation is ripe for tactile augmentation, especially in touch interfaces where occlusion can be problematic. Files, icons and other dragable items could be augmented with variable levels of friction to not only confirm that the target was successfully captured, but also convey properties like file size and drag-and-drop applicability (Figure 13). Object alignment, snapping and grid-based layouts could be also supplemented with tactile feedback (Figure 13). Such tactile augmentation could enable eyes-free interaction with sufficient practice.

Rubbing interactions

Repeated cursor motion over a region, i.e. *rubbing*, has been used in image editing applications for erasing, smoothing, desaturating and other procedures that incrementally increase or decrease some attribute of the image (Figure 14). Olwal et al. [26] investigated rubbing as a general technique for selection, targeting and zooming. Rubbing interaction offers an interesting application of dynamic tac-
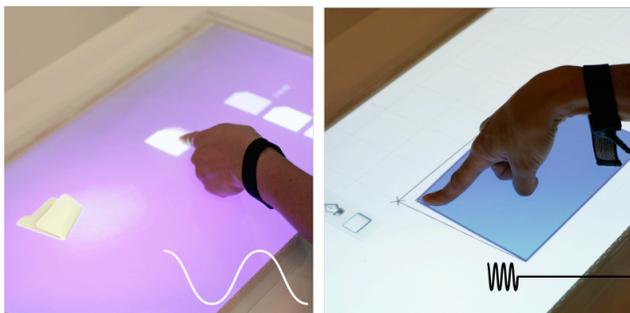
Figure 13: Left: Files being dragged to a folder have variable levels of friction based on their size. Right: Vibration diminishes as an object is dragged into alignment with neighboring items.

Figure 14: Friction between a user's finger and the touch surface decreases as the user increasingly erases a projected image by rubbing it.

tile feedback, e.g. as we progressively wipe out pixels in an area of an image, tactile sensation would decrease.

Applications Unique to Electrovibration

A unique quality of TeslaTouch is that only fingers in motions are stimulated. Therefore, it allows for *multitouch tactile feedback* so long as at each moment only one finger is moving on the surface. There are at least two examples where this can be employed in a unique and useful manner.

Anchored Gestures

These are gestures where one finger defines a reference point, while another finger is used for manipulation. A selection from a pie menu is one example, where one finger is static while another moves rotationally to select an item (Figure 15). Similarly shape transformations can be implemented, where one finger defines a static reference point while a moving finger specifies the amount of transformation, e.g. stretching, rotation or zooming. In all such operations, a moving finger can be easily supplemented with tactile feedback using TeslaTouch.

Two-Handed Asynchronous Manipulation

These are gestures that employ asymmetric separation of labor between the two hands [14]. For example, a non-dominant hand could perform a gross manipulation, such as orienting a sheet of paper, while the dominant hand performs a fine-grained interaction, such as writing. Another setup could use one or more modal buttons to define operation of a common slider (Figure 15). As in the previous example, one or more fingers are static, while one or more are
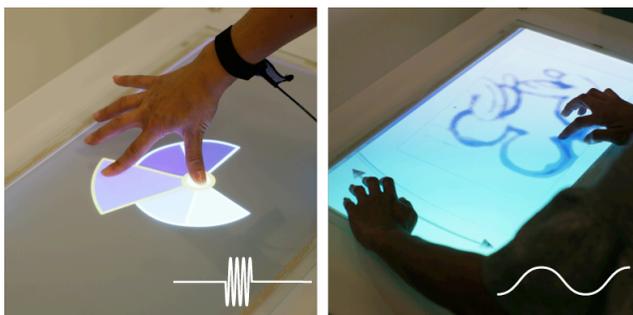

Figure 15: Left: A pie menu anchored by the middle finger and traversed by the forefinger. Right: The dominant hand is used to sketch; the non-dominant hand controls orientation.

engaged in movement and provided with tactile feedback using TeslaTouch.

Applications Unique to Mechanical Actuation

The fact the entire screen surface moves when mechanically actuated mean that vibrotactile actuation can provide effective tactile feedback to static, non-moving finger pressing against the surface of the screen. Therefore, such operations as feedback on button presses or a target acquisition are only supported by vibrotactile surfaces and cannot be implemented with TeslaTouch. In a similar manner, TeslaTouch is not able to provide tactile feedback for "press and hold" interactions, e.g. when finger touch or displacement is mapped to speed of movement of a virtual character.

CONCLUSION

This paper introduced TeslaTouch: a new technology for tactile display based on electrovibration. This technology can be adapted to a wide range of input tracking strategies, and can be used in many applications. Four experiments were conducted to characterize users' perception of TeslaTouch, providing a foundation for designing effective tactile sensations. A comparison between mechanical actuation and electrovibration led to an overview of the TeslaTouch applications design space. This first investigation of electrovibration opens up a wide range of possibilities for further research, both in human perception and interaction.

REFERENCES

1. 3M. *Microtouch technology brief*. Available from: http://solutions.3m.com.
2. Bau, O., U. Petrevski, and W. Mackay. *BubbleWrap: a textile-based electromagnetic haptic display*. in CHI'2009, Ext. Abstracts. 2009: ACM. pp. 3607-3612

3. Benko, H., A. Wilson, and R. Balakrishnan. *Sphere: multi-touch interactions on a spherical display*. in UIST'2008. 2008: ACM. pp. 77-86

4. Biet, M., F. Giraud, and B. Lemaire-Semail, *Implementation of tactile feedback by modifying the percieved friction*. The European Physical Journal: Applied Physics, 2008. **43**(1): pp. 123-135.

5. Brewster, S., F. Chohan, and L. Brown. *Tactile feedback for mobile interactions*. in CHI'2007. 2007: ACM. pp. 159-162

6. Buxton, W., R. Hill, and P. Rowley. *Issues and Techniques in Touch-Sensitive Tablet Input*. in SIGGRAPH'85. 1985: ACM. pp. 215-224

7. Collins, A., *Tactile television - mechanical and electrical image projection*. IEEE Transactions on Man-Machine Systems, 1970: pp. 65-71.

8. Dietz, P. and D. Leigh. *DiamondTouch: A Multi-User Touch Technology*. in UIST. 2001: ACM. pp. 219–226

9. Forlines, C. and R. Balakrishnan. *Evaluating tactile feedback and direct vs. indirect stylus input in pointing and crossing selection tasks*. in CHI'08. 2008: ACM. pp. 1563–1572

10. Fukumoto, M. and S. Toshiaki. *ActiveClick: Tactile Feedback for Touch Panels*. in CHI'2001, Extended Abstracts. 2001: ACM. pp. 121-122

11. Gescheider, G.A., *Psychophysics: The Fundamentals*. 3rd edition. 1997: Psychology Press.

12. Han, J. *Low-cost multi-touch sensing through frustrated total internal reflection*. in UIST'05. 2005: ACM. pp. 115-118

13. Harrison, C. and S. Hudson. *Providing dynamically changeable physical buttons on a visual display*. in CHI'2009. 2009: ACM. pp. 299-308

14. Hinckley, K., R. Pausch, D. Proffitt, J. Patten, and N. Kassell. *Cooperative bimanual action*. in CHI'97. 1997: ACM. pp. 27-34

15. Israr, A., S. Choi, and H.Z. Tan. *Detection Threshold and Mechanical Impedance of the Hand in a Pen-Hold Posture*. in IEEE/RSJ Intern. Conference on Intelligent Robots and Systems. 2006: IEEE. pp. 472-477

16. Israr, A., H.Z. Tan, and C.M. Reed, *Frequency and amplitude discrimination along the kinesthetic-cutaneous continuum in the presence of masking stimuli*. The Journal of the Acoustical society of America, 2006. **120**(5): pp. 2789-2800.

17. Kaczmarek, K., K. Nammi, A. Agarwal, M. Tyler, S. Haase, and D. Beebe, *Polarity effect in electrovibration for tactile display*. IEEE Transactions on Biomedical Engineering, 2006. **10**(53): pp. 2047-2054.

18. Kaczmarek, K., J. Webster, P. Pach-y-Rita, and W. Tompkins, *Electrotactile and vibrotactile displays for sensory substitution systems*. IEEE Transactions on Biomedical Engineering, 1991. **38**(1): pp. 1-16.

19. Lee, J., P. Dietz, D. Leigh, W. Yerazunis, and S. Hudson. *Haptic Pen: A Tactile feedback stylus for touch screens*. in UIST'2004. 2004: ACM. pp. 291-294

20. Leek, M.R., *Adaptive procedures in psychophysical research*. Perception and Psychophysics, 2001. **63**(8): pp. 1279-1292.

21. Levitt, H., *Transformed up-down methods in psychoacoustics*. The Journal of the Acoustical society of America, 1971. **49**(2): pp. 467-477.

22. Luk, J., J. Pasquero, S. Little, K. MacLean, V. Levesque, V. Hayward. *A role for haptics in mobile interaction: Initial design using a handheld tactile display prototype*. in CHI'06. 2006: ACM. pp. 171-180

23. Mallinckrodt, E., A. Hughes, and W. Sleator, *Perception by the Skin of Electrically Induced Vibrations*. Science, 1953. **118**(3062): pp. 277-278.

24. Matsushita, N. and J. Rekimoto. *HoloWall: designing a finger, hand, body and object sensitive wall*. in UIST'97. 1997: ACM. pp. 209-210

25. Miller, T. and R. Zeleznik. *The design of 3D haptic widgets*. in Symposium on Interactive 3D Graphics. 1999: ACM. pp. 97-102

26. Olwal, A., S. Feiner, and S. Heyman. *Rubbing and tapping for precise and rapid selection on touch-screen displays*. in CHI'08. 2008: ACM. pp. 295-304

27. Poupyrev, I. and S. Maruyama. *Tactile interfaces for small touch screens*. in UIST. 2003: ACM. pp. 217-220

28. Poupyrev, I. and S. Maruyama, *Method and programming for haptics display and feedback*, Japan Patent Office, P2008-257629. 2008.

29. Poupyrev, I., H. Newton-Dunn, and O. Bau. *D20: Interaction with Multifaceted Display Devices*. in CHI'2006, Extended Abstract. 2006. pp. 1241-1246

30. Poupyrev, I., M. Okabe, and S. Maruyama. *Haptic Feedback for Pen Computing: Directions and Strategies*. in CHI'2004. 2004: ACM. pp. 1309-1310

31. Senseg. Available from: http://www.senseg.com/.

32. Strong, R.M. and D.E. Troxel, *An electrotactile display*. IEEE Transactions on Man-Machine Systems, 1970. **11**(1): pp. 72-79.

33. Tang, H. and D. Beebe, *A microfabricated electrostatic haptic display for persons with visual imairments*. IEEE Transactions on Rehabilitation Engineering, 1998. **6**(3): pp. 241-248.

34. Verrillo, R.T. and G.A. Gescheider, Perception via the sense of touch, in *Tactile Aids for the Hearing Impaired*, I.R. Summers, Editor. 1992, Whurr Publishers: London. pp. 1-36.

35. Watanabe, T. and S. Fukui. *A method of contolling tactile sensation of surface roughness using ultrasonic vibrations*. in International Conference on Robotics and Automation. 1995: IEEE. pp. 1134-1139

36. Webster, J., Editor. *Medical instrumentation: Application and design*. 3rd ed. 1998, Wiley. p. 691.

37. Yamamoto, A., S. Nagasawa, H. Yamamoto, and T. Higuchi, *Electrostatic tactile display with thin film slider and its application to tactile telepresentation systems* IEEE Transactions on Visualization and Computer Graphics, 2006. **12**(2): pp. 168-177.