Interaction Capture in Immersive Virtual Environments via an Intelligent Floor Surface

Yon Visell, Alvin Law, Jessica Ip, Rishi Rajalingham, Severin Smith, Jeremy R. Cooperstock
McGill University, Montreal, Canada

ABSTRACT

This paper presents a novel, low-cost system and set of techniques designed to enable users to interact naturally on foot in immersive virtual environments. The physical interface consists of an array of networked, rigid floor tiles distributed over an area of several square meters within a CAVE-like virtual environment simulator. Each tile within the array is capable of 3 degree-of-freedom isometric sensing of forces applied by users' feet. The mechanical constraints of this interface are exploited by our data processing algorithms in such a way that it is possible to capture foot-floor interactions with a linear resolution as low as 1-2 cm, up to 30 times better than the nominal resolution of the tile array, over an area of several square meters. The resulting position and force estimates are employed to track contact positions and forces, and thereby to allow users to interact with distributed, deformable ground surfaces, such as soil or ice, presented in a virtual environment, or to operate virtual interface elements such as navigational widgets in a hands-free way. During interaction, users receive visual, auditory, and vibrotactile feedback in the form of physically plausible interactions between the foot and the ground surface or interface. We demonstrate the application of this device to the simulation of deformable ground surfaces in an immersive virtual environment, and to the presentation of virtual floor-based interface elements.

1 INTRODUCTION

Foot operated interfaces have long been instrumental in enabling people to efficiently work in complex real-world environments, ranging from automobiles to dental offices. However, comparatively little research has addressed foot-based interaction with the virtual world, whether via control interfaces or richly textured natural ground surfaces, such as snow or sand. Such interactions could prove instrumental to application domains such as immersive simulation, or to the development of computationally-enhanced control interfaces for real world environments, such as surgical rooms.

Arguably, one obstruction to further development along such lines has been the lack of reproducible human-computer interfaces (analogous to the now-ubiquitous touch screen interfaces for the hands) capable of efficiently capturing foot-floor contact interactions over a distributed area. In this paper, we present an approach to the design of such interfaces, based on a distributed network of low-cost, rigid floor tile components, with integrated sensing and actuation capabilities. In order to make efficient use of this interface, we draw on contact based sensing techniques that are able to capture foot-floor interactions with much finer resolution than is achieved if the tile is regarded as the smallest relevant spatial unit [3].

Most prior literature on interaction with floor surfaces has addressed the development of tactile sensing floor surfaces, based on dense arrays of force sensors (e.g., [9, 12, 8, 14]), with the aim of supporting applications such as person tracking, activity tracking, or musical performance. Such sensing interfaces are now commercially available, but the costs are high and systems complex, due to the large number of electronic elements involved. Further comparison is provided in Section 3.

A second focus of the present contribution is the parsimonious use of contact information captured through the floor interface for interaction with augmented floor surfaces. The basic element of information captured is the contact centroid. It provides an efficient local summary of the foot-floor interaction. As we demonstrate in Section 4, it can conveniently be utilized to channel interaction with a virtual deformable ground surface, providing plausibly realistic visual, auditory, and tactile feedback, or for selection and operation of a floor based multimodal control interface.

2 SYSTEM DESCRIPTION AND COMPONENTS

2.1 Device structure

The distributed floor interface as currently installed in the authors' laboratory is shown in Fig. 1. It consists of a square array of 36 individual tiles, each of which is equipped with force sensing and vibrotactile actuation capabilities (Sec. 2.1.1). In addition, the entire floor is surfaced in gray projection paint. A pair of overhead video projectors is used for visual display, in order to reduce the impact of shadows cast by users; with the addition of a video camera system, more advanced shadow removal techniques could be used [1]. An eight-channel audio display is used to present spatialized auditory stimuli, and in certain applications, a motion capture system (Vicon Mcam2 and M8i) is employed for body position sensing. One goal of the sensing research described in this paper has been to enable multimodal foot-floor interactions without requiring the use of a motion capture system. To this end, the floor is equipped with two-dimensional array of 144 force sensors (Fig. 2), which are integrated in the tiles themselves.

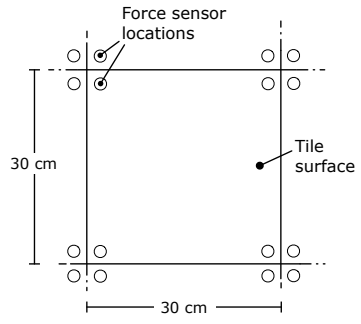

Figure 2: View from above showing the location of the sensors, near tile corners, within the distributed floor surface.

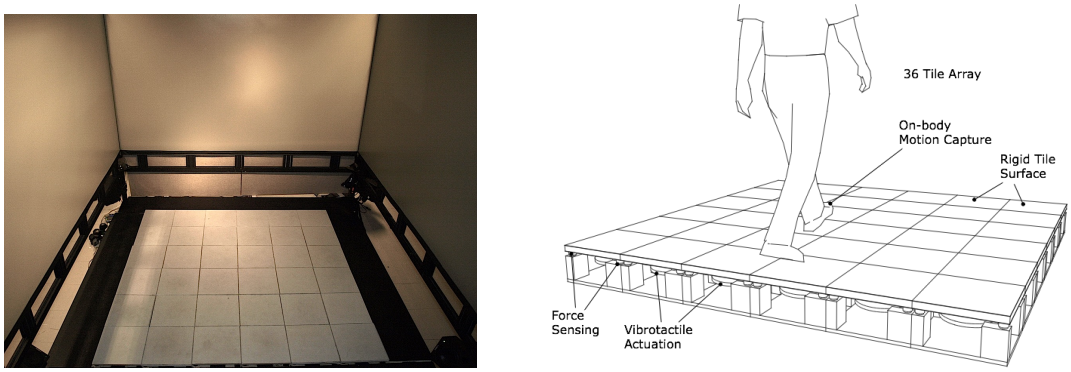

Figure 1: Distributed floor interface as installed in the authors' laboratory. Left: Photo showing the floor interface as situated within an immersive, rear projected virtual environment simulator. Right: illustration showing both sensing and actuating components.

2.1.1 Tile structure

Each tile provides a user interface via a solid plate with dimensions $30.4 \times 30.4 \times 2$ cm. The plate (baltic birch plywood) is supported by vibration mounts (cylindrical SBR rubber elastic supports), which are positioned as shown. A retaining socket surrounding the vibration mount (not present in the single-tile figure) is used to keep the plate in place during interaction. The tiles each incorporate four force sensors located under the vibration mounts, as described below (Sec. 2.2).

A vibrotactile actuator (Clark Synthesis, model TST229) is mounted via an aluminum bracket bonded to the center underside of the plate. The actuator is driven by an audio amplifier providing up to 80 W to an impedance of 4 Ohms (manufacturer's ratings). The actuator signals are generated on a personal computer (see Sec. 2.3, below), and streamed via a 24 bit 96 kHz digital audio interface (Edirol model FA-101). A framed substructure supports the entire floor, incorporating conduits and access plates for the digital and analog data connections.

The engineering of the active response of the vibrotactile floor tile display is described elsewhere [17]. It has a usable vibrotactile bandwidth from about 50 Hz to 1 kHz, and is capable of displaying vibrotactile signals at the highest amplitudes required for the display of virtual ground surface properties via the vibrotactile channel (i.e., more than 40 N across the indicated frequency band).

2.2 Electronic sensing components

In order to capture foot-floor interactions, we sense the normal forces applied to each tile by the user's foot and process these force signals using the methods described in Sec. 3. Positioning the force sensors beneath the plate is feasible, since the bandwidth of the force applied to the plate by the user's body (via the foot) is limited. The sensors are furthermore positioned beneath the elastic suspension of the device so as to better isolate the sensor data from receiving feedback from the actuator signals.

Force sensing is performed via four resistive force sensors (Interlink model 402 FSR) located below the vibration mount located under each corner of each plate. Analog data from the force sensors is conditioned, amplified, and digitized via a custom 32-channel acquisition board based on the Altera FPGA with 16-bit analog-to-digital converters. Data from each sensor is sampled at a rate of 1 kHz and transported to an array of computers over UDP via the board's 10 Mbps ethernet interface, using the Open Sound Control protocol.

A distributed floor array requires a two dimensional $m \times n$ arrays of tiles, requiring a number $N = 4mn$ of sensors. After conditioning, the response of these sensors to an applied force is nonlinear, and varies up to 25% from part to part (according to manufacturer ratings). Consequently, a measurement is performed using a calibrated load cell force sensor, and the measured voltages are linearized to give absolute force values. The calibration function is obtained by fitting the measured force f vs. voltage v values for each sensor to an exponential function $f = a\exp(bv)$ with two free parameters using the least mean squares algorithm. Using this method, a linear response accurate to within about 5% can be obtained using low cost parts. The cost of this sensing system is a small fraction of that of a typical motion capture system.

2.3 System design and network

Figure 4 shows the system architecture. Briefly, the system components all perform digital data exchange over a high speed (Gigabit) ethernet network. An array of 6 small form factor computers is used for interaction capture and audio-tactile synthesis. Each is responsible for one row in the array of tiles. It receives data streamed from one sensor acquisition unit, processes the data using the techniques described below, and broadcasts the interaction data (force and position data from individual contact centroids; see Sec. 3) over the network at a lower sampling frequency (100 Hz). In addition, the computer is responsible for synthesizing vibrotactile feedback in response to the force signals, ensuring a response to force input can be provided at low latencies.

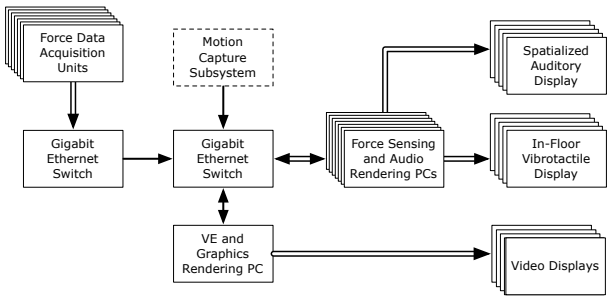

Figure 4: Diagram illustrating the logical connections between system components in the floor interface described above.

3 Contact-based sensing

Intrinsic contact based sensing aims to resolve the locations of contact, forces at the interface, and the moment about

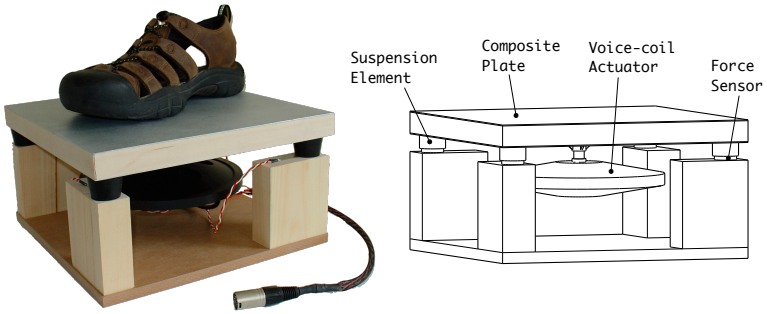

Figure 3: Vibrotactile floor interface hardware for a single tile unit. Left: Photo with large mens' shoe, showing representative size. The force sensor connection cable is visible in the foreground, and the actuator is visible under the tile. Right: View with main components labelled.

the contact normals using internal force and torque measurements [2]. It can be viewed as an alternative to distributed tactile sensing via surface mounted arrays that directly measure the pressure distribution over the area of contact. By comparison, contact based sensing requires a number of sensors that is on the order of the number of kinematic degrees of freedom of the apparatus, far fewer than are required for surface-based techniques. Furthermore, the data one obtains provide can often provide an efficient summary of the information that is most salient to a contact interaction. Although contact based sensing has, to date, been mainly applied to problems in robotic manipulation, here we employ it for foot-ground interaction sensing via a distributed floor interface.

Although these methods can be readily extended to the case of contact with friction, we constrain the discussion in most of what follows to frictionless soft-finger contacts [13]. One reason is that, while friction plays a key role in human movement on foot, peak foot-ground forces in the tangential (shear) direction are, during normal walking, about 80% smaller than peak normal forces [11]. Moreover, our current device possesses only three degrees of force sensing freedom per tile. Six degrees of freedom would be required to resolve contact centroids under the assumption of soft finger contacts with friction (see below and Sec. 5). Finally, the interactive affordances we provide through this interface primarily involve the application of normal forces to areas of the floor by the foot (as in the case of depressing a button, or leaving a footprint in loose soil).

Consider an area R of contact between the foot and floor, over which a frictionless normal force distribution $p_R(xx)$ is distributed. The contact centroid $\mathbf{x}_c$ is a unique point on the floor such that there is a normal force F_c that gives rise to the same intrinsic force measurements as $p_R(\mathbf{x})$ does [2]. A basic result is that the contact centroid is guaranteed to lie within the convex hull of the contact area (dashed line, Fig. 5) [2]. It thus provides a concise summary of the foot-floor contact locus, although it does not provide information about shape or orientation.

Force equilibrium in the normal direction implies that $F_c = \int_R p_R(\mathbf{x}) d\mathbf{x}$. Furthermore, in the frictionless setting, $\mathbf{x}_c$ coincides with the pressure centroid:

$$\mathbf{x}_c = \int_R p_R(\mathbf{x})\,\mathbf{x}\,d\mathbf{x}/F_c. \tag{1}$$

3.1 Problem formulation

The contact sensing problem is simple to state when the contact region of interest, R, lies within a single floor tile (Fig. 5) and the possibility of frictional contact is ignored. The tile has force sensor locations $\mathbf{x}_j$ where internal force measurements f_j are taken. j indexes the tile sensors. The contact

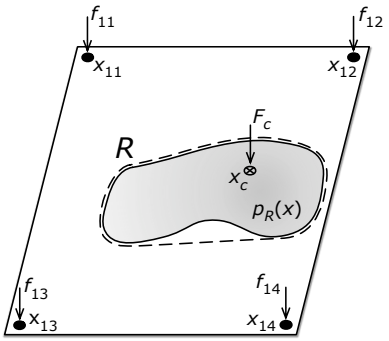

Figure 5: The contact centroid, $\mathbf{x}_c$, and corresponding normal force, f_c, computed from frictionless soft finger contact for a single floor tile. Contact results in a normal force distribution $p_R(\mathbf{x})$ (gray-level shading) with support on R. The contact centroid, computed from the forces captured at the sensor locations $\mathbf{x}_j$, lies within the convex hull of R, indicated by the dotted line.

sensing problem is to recover the contact centroid $\mathbf{x}_c$ and normal force $\mathbf{F}_c = (0,0,F_c)$ from the intrinsic measurements $\mathbf{F}_j = (0,0,f_j)$, which are the scalar (normal) forces captured at the sensor locations $\mathbf{x}_j$. These quantities are related by force and torque equilibrium equations,

$$\begin{aligned} \sum_{j=1}^{4} f_j + F_c + f_p &= 0 \\ \sum_{j=1}^{4} \mathbf{x}_j \times \mathbf{F}_j + \mathbf{x}_c \times \mathbf{F}_c + \mathbf{x}_p \times \mathbf{F}_p &= 0\,. \end{aligned} \tag{2}$$

Here, $\mathbf{F}_p = (0,0,f_p)$ accounts for the (known) force exerted by the plate and actuator, which acts at the tile's center, $\mathbf{x}_p$. The three nontrivial scalar equalities (2) can be readily solved for the contact centroid parameters, yielding:

$$F_c = \sum_{i=1}^{4} f_i - f_p, \quad \mathbf{x}_c = \frac{1}{F_c}\left(\sum_{i=1}^{4} (\mathbf{x}_i - \mathbf{x}_p) f_i + f_c \mathbf{x}_p\right). \tag{3}$$

3.2 Single-foot, multiple tile contact

Although the simple description above breaks down when the foot overlaps two or more tiles, in the case of multi-tile contact, more information is actually available. The contact

area $R = \bigcup_k R_k$, where the contact domain R_k lies within the the kth tile (Fig. 6). The pressure centroid $\mathbf{x}_R$ for the entire contact area can be computed from the contact centroids $\mathbf{x}_{ck}$ for each tile (computed from Eq. (3)) and is given by the weighted average

$$\mathbf{x}_c = w_1\mathbf{x}_{c1} + w_2\mathbf{x}_{c2}\ , \quad w_k = F_i/F \tag{4}$$

Computed in this way, $\mathbf{x}_R$ is manifestly invariant with respect to rigid transformations of the tiling geometry, including those that change the overlap domains. Stated conversely, the position of the weighted contact centroid is preserved as the contact point moves across the tile domain. By abuse of notation, we write it as $\mathbf{x}_c$, using the same denotation as when referring to a single tile. This allows us to track interaction points as they move across tile boundaries, as illustrated in the supplementary video documentation.

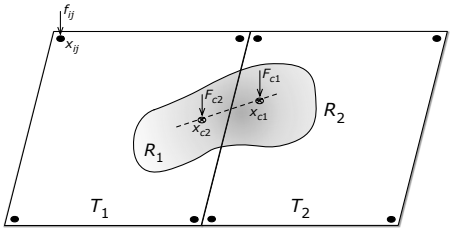

Figure 6: Contact based sensing for a frictionless soft finger contact distribution $p_R(\mathbf{x})$ on a region R spanning two adjacent tiles with domains T_1 and T_2. The contact domains R_1 and R_2 on each tile give rise to centroids x_{c1} and x_{c2}. Their weighted sum is $\mathbf{x}_c$ the contact centroid for the pressure distribution with support $R = R_1 \bigcup R_2$, and lies on the line segment connecting $\mathbf{x}_{c1}$ and x_{c2}. The difference vector $\delta\mathbf{x} = \mathbf{x}_{c1} - \mathbf{x}_{c2}$ provides further information about the contact shape.

Figure 6 illustrates the two tile setting. The contact area $R = R_1 \bigcup R_2$. The pressure centroid $\mathbf{x}_c$ lies on the line segment connecting $\mathbf{x}_{c1}$ and $\mathbf{x}_{c2}$. The difference vector $\delta\mathbf{x} = \mathbf{x}_{c1} - \mathbf{x}_{c2}$ provides additional shape information about the orientation of the contact distribution at the boundary. For a contact distribution that lies on a line segment spanning the two tiles, the direction $\mathbf{n} = \delta\mathbf{x}/|\delta\mathbf{x}|$ of the difference vector coincides with that of the contact shape, and the length $|\delta\mathbf{x}|$ is proportional to the extent of R. For convex contact shapes that are less sharply oriented (such as those of a foot), the range of angles is compressed around the edge normal direction. For simple geometric shapes, such as rectangles or ellipsoids, one can directly compute the expected disparity. Figure 7 provides a qualitative illustration for a family of ellipsoidal contact regions.

3.3 Evaluation of contact position estimates

Figure 8 shows a comparison of measured and estimated contact positions using the contact centroid method. The data acquired from a single calibrated floor tile. Measurements were taken by applying a static (constant) force to measured positions on the surface of the tile, recording the resulting force values, and applying formula (3). Although there is significant distortion near the tile edges, the contact centroid is typically localized with an accuracy < 2 cm, and in the worst case < 5 cm. These numbers compare favorably to the linear dimensions of the tile surface (30 cm) and to the typical dimensions of an adult shoe. Moreover, straight line arrays of centroids remain relatively straight, so that interaction points may be effectively tracked over larger distances, spanning tile boundaries (see also Sec. 4).

Figure 9 shows an instance of an extended object crossing the boundary of a tile at several angles. As can be seen

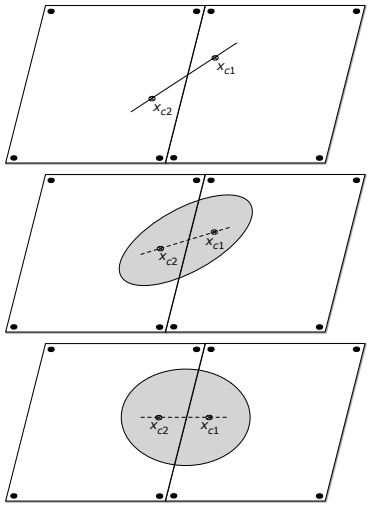

Figure 7: Orientation information in contact distributed between adjacent tiles. Top: For a pressure distributed on a line segment spanning the two tiles, $\delta\mathbf{x} = \mathbf{x}_{c1} - \mathbf{x}_{c2}$ coincides with the direction of the line. Bottom: For a uniform, disk-shaped pressure distribution, δx is always normal to the edge between the tiles. Middle: For an, ellipsoidal distribution, $\delta\mathbf{x}$ points at an angle between the edge normal and the major axis of the ellipsoid. As distribution grows wider (i.e., closer to a disk shape) $\delta\mathbf{x}$ is compressed toward the edge normal.

from the figure, orientation information about the object is preserved, but the range of angles is compressed, due to the extended nature of the object and to distortion of centroid positions near the edges of the tile.

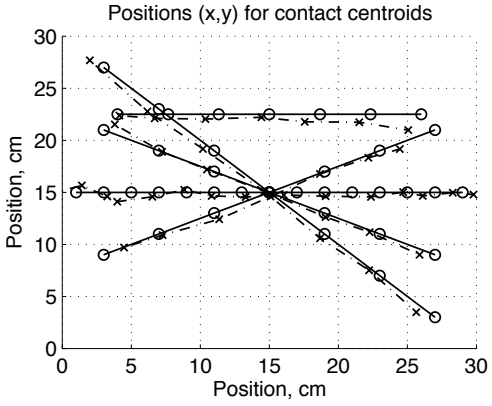

Figure 8: Estimated contact centroids from force measurements due to point contact at various positions along straight lines across a single tile. True positions are shown as circles, estimated positions are shown as Xs.

3.4 Multi-foot contact

Except in cases of users with extraordinarily large shoe sizes, a single foot is able to overlap at most four tiles. The possible configurations up to rigid transformations of the plane are shown in Figure 10. Ignoring kinematic constraints on the body, two feet may be arranged via any combination of these configurations, either overlapping or non-overlapping. As is evident from Figure 11, the contact centroid is less useful as a summary of foot-floor contact when more than one foot lies on a given tile, and is a source of potential artifacts. More generally, optimal determination of

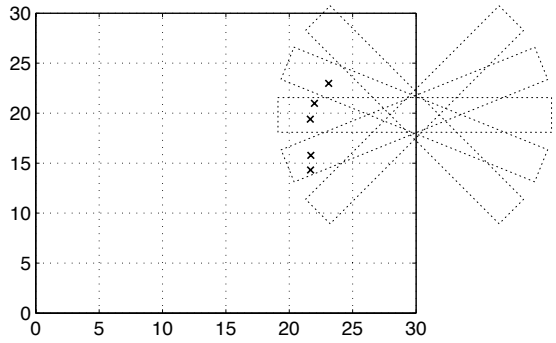

Figure 9: Estimated contact centroids from force measurements from an extended object spanning the border between two adjacent tiles. A uniform force of 51 N was distributed over the area of the object, with the object supported across both tiles. Estimated centroids for the left hand tile are shown as Xs, and the object outline as positioned is shown as a dashed-line rectangle.

foot and body position from coarse in-floor measurements poses a challenging inverse kinematics problem which is being addressed in ongoing work in our lab. In the present contribution, we focus on the simpler problem of resolving well-segmented interaction contact points.

For many applications, including those described in Sec. 4 (below), information about the locations of both feet is not required. Finally, although not a focus of the present contribution, if a motion capture system is employed (as in the scenario described in [5]), it can be used to supply accurate position data, with the force sensors employed to determine contact onsets and forces. As noted above, one goal of our research in this area is, when appropriate, to eliminate the need for such a system.

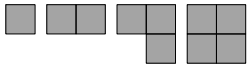

Figure 10: There are four possible tile set shapes overlapped by the contact area of one foot, up to rigid transformations of the plane.

4 Applications to VR and Multimodal HCI

We illustrate two potential applications of these interaction techniques, to the simulation of natural ground surfaces for immersive virtual environments, and the implementation of multimodal human computer interaction via floor surfaces.

4.1 The frozen pond

As an engaging and interactive virtual environment scenario making use of the interaction capture techniques described here, we implemented a demonstration allowing its users to walk over a virtual pond whose surface has frozen over. If a user walks on the pond insufficiently carefully, patterns of cracks form in the surface ice, dependent on the locations and forces of foot-floor contact (Fig. 12). The interactions with this virtual material are rendered using audio, visual, and vibrotactile channels. Video documentation is provided in the supporting material.

The frozen pond scenario provides a useful testing ground, because an interactive response can be rendered without detailed knowledge of the shape of the foot-floor contact region. Unlike simulations of loose ground materials, such as soils, which require that foot posture be tracked via a motion tracking system (e.g., [5]), the simulation described here uses only the in-floor force sensing array.

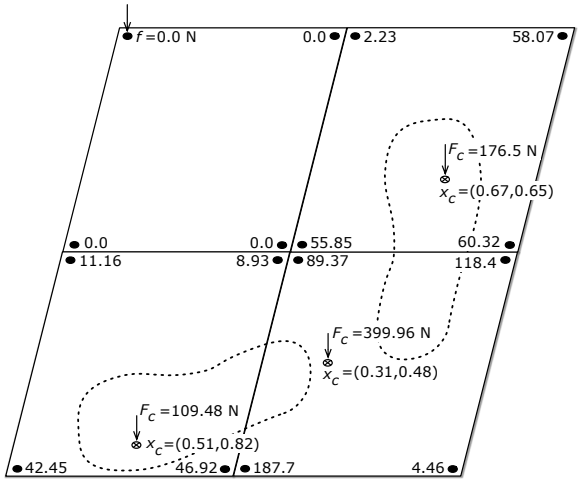

Figure 11: Measured example of a force pattern arising when two feet contact a single tile. Force values (in Newtons) and estimated contact centroids (normalized intra-tile coordinates) for a 2×2 tile floor area are as shown, together with foot locations and orientations.

Figure 12: Still image from video documentation of the frozen pond demonstration. Cracks in the virtual ice sheet dynamically form at loci of contact between feet and floor.

4.1.1 Audio-Tactile rendering

Audio and vibrotactile display channels provide plausibly realistic feedback recapitulating the cracking of the virtual ice sheet underfoot. The two channels are derived from the same signals, which are synthesized by a physically-motivated lumped stochastic process (Fig. 13), similar to those employed in prior work [16, 15]. The process models the distribution of energy loss through each ith inelastic cracking event in the virtual ice sheet, through an event time t_i and energy loss E_i. Separate, independent responses are rendered in parallel for each tile in the array. The elastic part of the stress-strain relationship is taken to be of the form: $f(t) = m\ddot{x} + b\dot{x} + Kx$, where f is the net foot-floor force on the tile, x is a virtual displacement, and m, b, and K are material parameters governing the mass density, viscous damping and stiffness of the virtual ice sheet.

Briefly, the temporal onset of cracking events is modeled as a Poisson process with a rate parameter given by the one-sided "virtual kinetic energy": $\kappa_+(t) = \frac{1}{2}m\dot{x}^2$. We define κ_+

to be zero if $\dot{x} < 0$ (i.e., inelastic slip occurs only when force is increasing). The energy $E(t)$ of a cracking event at time t is sampled from an exponential distribution $p(E) \propto E^{\gamma}$ with free (material-related) parameter γ. Such a choice can be motivated by assuming there is a limited potential energy that can be dissipated during a single footstep [16].

The auditory and tactile signatures of a cracking event are synthesized as discrete transient events, given by an impulsive signal with energy E passed through a representation of the resonant response of the sheet. The latter is modeled as a bank of modal oscillators with impulse response $s(t) = \sum_i a_i e^{-b_i t} \sin(2\pi f_i t)$, determined by amplitudes a_i, decay rates b_i, and resonant frequencies f_i.

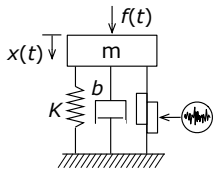

Figure 13: Cracking events at each tile are modeled via a lumped, stochastic process model, depending on parameters governing virtual stiffness, mass density, and damping. Inelastic cracking events are modeled by the plastic slip unit at right, depending on the noise source shown.

4.1.2 Visual rendering

Typical approaches to the animation of surface cracks are based on models of the evolution of the stress distribution within the material as discontinuities are introduced [10, 7, 4, 6]. Here, we adopt a simplified version of this idea, modeling stress release via the temporal stochastic process described above. This preserves good timing correlation between modalities, and ensures that the perceptually harder temporal requirements of the audio-tactile display can always be met. The contact centroid $\mathbf{x}_c$ is treated as an efficient summary of the spatial stress distribution generated by the foot.

A crack pattern consists of a collection of crack fronts, each of which is defined by a linear sequence of nodes, $\mathbf{c}_0, \mathbf{c}_1, \ldots, \mathbf{c}_n$, which are positions on the floor originating with seed locations $\mathbf{p}$, such that $\mathbf{p} = \mathbf{c}_0$. The crack pattern is rendered as a set of line intervals $\ell_k = (\mathbf{c}_k - \mathbf{c}_{k-1})$ on the ice sheet (Fig. 14). The seed locations $\mathbf{p}$ are determined by foot-floor contact at the time they are created. Several fronts are allowed to radiate from each seed. Our method is mesh-free, and the seed locations are unconstrained.

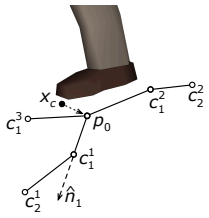

Figure 14: A crack pattern is modeled as a graph consisting of sequences of nodes originating with the crack seed $\mathbf{p}_0$.

A crack event is generated by the audio-tactile process occurs at time t_i with energy $E(t_i)$ as described above. It may either result in the creation of a new seed or the growth of fractures from an existing one. A new seed $\mathbf{p}$ is formed at the location of the dominant contact centroid $\mathbf{x}_c$ (i.e., the one with largest force f_c) if no previously created seed lies within a fixed threshold distance Δp from $\mathbf{x}_c$. This new seed $\mathbf{p}$ is created with a random number N_c of initialized crack fronts, $\mathbf{c}_0^1, \mathbf{c}_0^2, \ldots \mathbf{c}_0^{N_c}$, where we sample N_c uniformly in $2, 3, \ldots 6$. Two of these are then initiated, using the algorithm below.

A crack event propagates a crack front originating with the seed $\mathbf{p}$ nearest to $\mathbf{x}_c$. With probability $1/N_c$ the jth crack front of $\mathbf{p}$ is selected for extension. Crack propagation is determined by a propagation vector $\mathbf{d}_m^j$ such that $\mathbf{c}_m^j = \mathbf{c}_{m-1}^j + \mathbf{d}_m^j$. We take $\mathbf{d}_m^j = \alpha E \hat{\mathbf{n}}_m^j$, where E is the crack energy, α is a global parameter controlling the rate of crack growth, and $\hat{\mathbf{n}}_m^j$ is the direction. Since we do not have information about the principal stress directions at the crack front, we propagate in a direction $\hat{\mathbf{n}}_m^j = \hat{\mathbf{n}}_{m-1}^j + \beta \hat{\mathbf{t}}$, where $\beta \sim N(\beta; 0, \sigma)$ is a Gaussian random variable and $\hat{\mathbf{t}} = \hat{\mathbf{n}}^j \times \hat{\mathbf{u}}$, where $\mathbf{u}$ is the upward surface normal (i.e., $\mathbf{t}$ is a unit vector tangent to $\hat{\mathbf{n}}^j$). The initial directions of crack fronts at $\mathbf{p}_0$ are spaced equally on the circle.

4.2 Floor-based interface elements

The second application we have explored for the interaction capture techniques described above implements an array of floor-based control widgets or gestures (Fig. 15). As in the frozen pond scenario, sensing is performed exclusively using the in-floor force sensing array. The interaction point ("cursor") is taken to be the contact centroid $\mathbf{x}_c$ with largest associated force within the active, control interface region of the floor. In these demonstrations, users are standing on the floor and are able to access nearby control elements (buttons, sliders, or other switches) or regions using their feet, in the same manner as is done with common touch surface interfaces. Normal force thresholds are used to determine when buttons or other controls are being engaged.

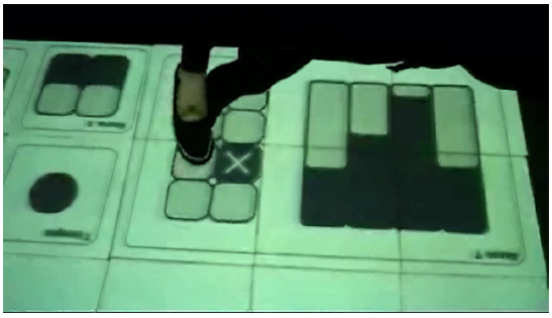

Figure 15: Still image from video documentation of a demonstration application in which users are able to interact with various floor-based interface widgets. The system uses the interaction capture techniques described above. The discrete buttons and switches produce a tactile click when depressed, while the continuous sliders provide a simple, friction-like vibrotactile feedback. The complete video is included in the supplementary materials.

Concurrent audio and vibrotactile feedback, in the form of clicks, taps, or rubbing sounds or vibrations, is supplied accompanying the discrete or continuous response of a control. Video documentation of one of the demonstration applications is provided in the supplementary material. As can be seen in the video, the appropriate size of a control can easily be limited as much by the dimension of the foot as by the resolution of the contact-based position sensing, as we intend to investigate in future work.

Interactions with floor controls as these are common in many fields, such as manufacturing, surgery, or dentistry,

where ergonomic problems with embodied control interfaces have been documented [18]. Potential advantages of virtualized floor control interface are readily identified, such as their software reconfigurability and lack of obtrusiveness when not in use.

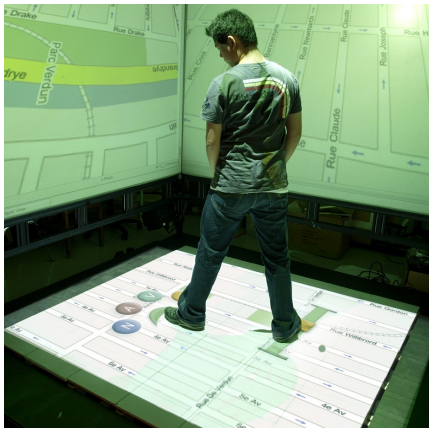

Figure 16: Mockup of a proposed application in which users interact with floor based interface widgets in order to control an immersive map application.

5 Conclusion

This paper described interaction techniques based on intrinsic contact based force sensing in a distributed floor interface. Such an approach may well-suited to situations in which foot-floor contact interactions are of particular interest, because such information may not be available through other commonly used sensing channels, such as motion capture. The methods presented here are low in cost and system complexity, and can be deployed for multiple simultaneous users, without requiring them to wear any specialized equipment. In addition, this paper demonstrates the integration of these interaction techniques within multimodal displays implementing virtual ground surface simulations or floor-based control interfaces.

Despite the promising nature of these results, there are several aspects in which the present system might be improved or extended:

- The current tile interface is capable of only three degree-of-freedom (DOF) sensing. In this setting, the contact sensing problem can be solved only under the assumption, made in Sec. 3, of frictionless soft contact. A future version of this interface, capable of sensing the full six rigid DOF of the tile, would capture contact interactions more accurately in uncontrolled settings by accounting for friction effects.

- A floor interface with a denser array of tiles would be capable of capturing more information about footground contact shape.

- During multi-tile foot-floor contact, a contact-based sensing approach results in clusters of contact centroids. New techniques are needed in order to acquire the information arising from such features.

- Improvements would be expected from the parsimonious use of available (prior) information about kinematic (skeletal) constraints on users, which make certain force patterns significantly more or less likely. Similarly, one anticipates that good use can be made of temporal constraints, such as continuity. We are addressing these issues within a Bayesian filter framework, that aims to optimally track foot and body posture during the course of interaction via simplified skeletal models

- Further work is needed in order to develop usability guidelines for floor-based interfaces, including basic guidelines related to control element size, display scale, and other aspects salient to implementing such a UI, such as we have only begun to observe here.

It is hoped that the present contribution convinces the reader of the potential of such floor-based interaction methods, and that other researchers are inspired to contribute in areas such as those noted above, or in others that have not yet been anticipated.

Supplementary Material

The supplementary materials consist of video documentation of the interactive scenarios described in Section 4 of the paper:

1. **Frozen pond:** A floor-based virtual environment in which users are able to walk across a frozen pond. Their footsteps are sensed using the techniques described in this paper, and users are able to see, hear, and feel cracks forming beneath their feet as they walk.

2. **Floor-based control interfaces:** Demonstration of a number of user interface elements controlled via foot-floor contacts. The video also illustrates the system's ability to track foot interaction locations across tile boundaries. (The authors are re-shooting this documentation, both to improve the video quality and to eliminate occasionally visible lag that arose during recording for extraneous reasons.)

Acknowledgements

The authors gratefully acknowledge support from the MDEIE of Quebec for the EU FP7 project NIW (no. 222107). The first author acknowledges support from the ESF COST Action on Sonic Interaction Design (no. IC0601).

References

[1] S. Audet and J. Cooperstock. Shadow removal in front projection environments using object tracking. In *IEEE Conference on Computer Vision and Pattern Recognition, 2007. CVPR'07*, pages 1–8, 2007.

[2] A. Bicchi, J. Salisbury, and D. Brock. Contact sensing from force measurements. *The International Journal of Robotics Research*, 12(3):249, 1993.

[3] T. Delbrueck, A. Whatley, R. Douglas, K. Eng, K. Hepp, and P. Verschure. A tactile luminous floor for an interactive autonomous space. *Robotics and Autonomous Systems*, 55(6):433–443, 2007.

[4] K. Hirota, Y. Tanoue, and T. Kaneko. Generation of crack patterns with a physical model. *The Visual Computer*, 14(3):126–137, 1998.

[5] A. Law, B. Peck, Y. Visell, P. Kry, and J. Cooperstock. A multi-modal floor-space for displaying material deformation underfoot in virtual reality. In *Proc. of the IEEE Intl. Workshop on Haptic Audio Visual Environments and Their Applications*, 2008.

[6] A. Norton, G. Turk, B. Bacon, J. Gerth, and P. Sweeney. Animation of fracture by physical modeling. *The Visual Computer*, 7(4):210–219, 1991.

[7] J. O'Brien and J. Hodgins. Graphical modeling and animation of brittle fracture. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 137–146. ACM Press/Addison-Wesley Publishing Co. New York, NY, USA, 1999.

[8] R. J. Orr and G. D. Abowd. The smart floor: a mechanism for natural user identification and tracking. In *CHI '00: CHI '00 extended abstracts on Human factors in computing systems*, pages 275–276, New York, NY, USA, 2000. ACM.

[9] J. Paradiso, C. Abler, K.-y. Hsiao, and M. Reynolds. The magic carpet: physical sensing for immersive environments. In *CHI '97: CHI '97 extended abstracts on Human factors in computing systems*, pages 277–278, New York, NY, USA, 1997. ACM.

[10] M. Pauly, R. Keiser, B. Adams, P. Dutré, M. Gross, and L. Guibas. Meshless animation of fracturing solids. *Proceedings of ACM Siggraph 2005*, 24(3):957–964, 2005.

[11] J. Perry. *Gait analysis: normal and pathological function*. SLACK incorporated, 1992.

[12] R. F. Pinkston. A touch sensitive dance floor/midi controller. *The Journal of the Acoustical Society of America*, 96(5):3302–3302, 1994.

[13] J. Salisbury Jr. Interpretation of contact geometries from force measurements. In *1984 IEEE International Conference on Robotics and Automation. Proceedings*, volume 1, 1984.

[14] A. Schmidt, M. Strohbach, K. v. Laerhoven, A. Friday, and H.-W. Gellersen. Context acquisition based on load sensing. In *UbiComp '02: Proceedings of the 4th international conference on Ubiquitous Computing*, pages 333–350, London, UK, 2002. Springer-Verlag.

[15] Y. Visell, J. Cooperstock, B. Giordano, K. Franinovic, A. Law, S. McAdams, K. Jathal, and F. Fontana. A Vibrotactile Device for Display of Virtual Ground Materials in Walking. *Lecture Notes in Computer Science*, 5024:420, 2008.

[16] Y. Visell, F. Fontana, B. Giordano, R. Nordahl, S. Serafin, and R. Bresin. Sound design and perception in walking interactions. *International Journal of Human-Computer Studies*, 2009.

[17] Y. Visell, A. Law, and J. Cooperstock. Touch Is Everywhere: Floor Surfaces as Ambient Haptic Interfaces. *IEEE Transactions on Haptics*, 2009.

[18] L. Wauben, M. van Veelen, D. Gossot, and R. Goossens. Application of ergonomic guidelines during minimally invasive surgery: a questionnaire survey of 284 surgeons. *Surgical endoscopy*, 20(8):1268–1274, 2006.