

Electric Field Sensing For Graphical Interfaces



Joshua Smith, Tom White, Christopher Dodge,
Joseph Paradiso, and Neil Gershenfeld
MIT Media Laboratory

David Allport
VSIS

We outline the theory and implementation of low-frequency electric fields as sensing techniques, then present a range of applications developed for interacting with computer graphics.

The earliest bit-mapped graphical computers have progressed to bring real-time 3D rendering and digital video to the desktop, but the common physical interface remains unchanged from the first workstations—keyboards and mice. As a result, many applications, such as modeling or navigating in virtual worlds, are often limited not by processing speed but by the users' difficulty in conveying desired actions to the computer.

The range of alternative controllers tried is notable not just for diversity, but for their compromising constraints or relatively disappointing overall performance. They divide broadly into noncontact approaches and those that require contact with the user. Video—the most familiar noncontact interface—uses one or more cameras to determine the users' actions.¹ The canonical problems with such computer vision approaches include

- difficulty in obtaining estimates faster than standard video frame rates,
- the need for enormous input bandwidth and computational power to process the images, and
- the requirement to constrain the scene background, activity, and illumination to make the recognition task feasible.

Given people's limiting performance on the order of millimeter displacements in milliseconds, a large gulf exists between what people can do and what cameras can reasonably be expected to recognize.

Other common, noncontact sensing schemes use ultrasonic or infrared reflections.² These scattering mechanisms depend on surface texture and orientation, thereby affecting the accuracy and reliability of the measurements. Furthermore, environmental background signals, such as sunlight for infrared and mechanical

noise for ultrasound, can become significant sources of interference for most of these systems. Wired interface devices include variants of the Data Glove³ and families of magnetic position trackers.⁴ These have the obvious constraint of encumbering the user with extra hardware, something that can range from inconvenient to impossible for extended or intermittent use.

The ideal interface would be no apparent interface: Users would act as they wish, using either free gestures or manipulating objects that provide tactile cues, and the computer would unobtrusively and continuously know the user's state, at the limits of their physical performance. This desire is not entirely unrealistic. The Theremin,⁵ one of the earliest and most responsive of electronic instruments, hints at one solution. Sensing how the presence of a player's body capacitively loads antennas used to set the pitch and amplitude of an oscillator⁶ makes exquisite musical control possible, approaching the lyrical expression of the human voice. A variation of the Theremin, the Radio Baton,⁷ uses a shaped electrode array and wired batons. Although it adds wires to the interface, it can obtain absolute 3D information.

The Theremin is an old idea, but some aspects of the physics and mathematics behind it remain open questions. A number of charge-transport pathways contribute to the observed signal and have historically been lumped together as "capacitance." These can be separated and measured, posing a new nonlinear inverse problem to go from the measured charges to the conductors perturbing the field. It may surprise you that this inverse problem is nonlinear, since the Laplace equation (like the more general Maxwell equations from which it may be derived) is a linear partial differential equation. This means that for equipotential boundaries at fixed locations, superpositions of solutions are also solutions. However, the electric field strength at the receiver is not a linear function of the location of these equipotential boundaries. The inference problem is to determine the location of the boundaries from knowing the field strength at particular positions. Thus, unlike more familiar linear imaging problems that arise in com-

puted tomography or magnetic resonance imaging, this problem occurs on a nonlinear basis, requiring a much more difficult search. By multiplexing the transmitters and receivers, multiple “projections” can be measured, and we believe this will enable a new form of true 3D imaging. Our first efforts to infer 3D geometrical information from electric field measurements appear elsewhere,⁸ further work on this problem is forthcoming.

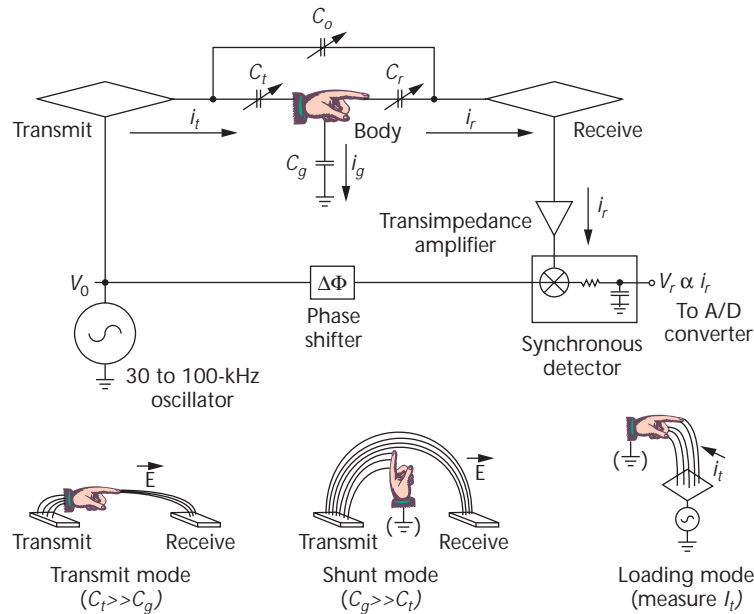
Although extracting 3D images is enticing, one of the most appealing features of this technology is the continuum that extends from imaging with a large array of electrodes to position measurement with small numbers of electrodes, all the way down to a single-electrode “button” that measures distance and contact.

One of the main difficulties with video is that a camera collects too much information. A sensing modality that lets you collect as much or as little information as needed in a particular application is appealing and unusual.

Mechanisms and implementation

Figure 1 depicts the basic implementations of electric field sensing. The top diagram is a model that describes all sensing modes. This simple hand sensor consists of two electrodes: A transmitter driven by a low-frequency, low-voltage radio source (orders of magnitude below Federal Communications Commission and health restrictions) and a receiver that detects the transmitted signal through the capacitive paths given in the figure. To reduce interference from ambient electromagnetic background, the receiver usually has a narrowband response centered at the transmit frequency, generally provided by a synchronous detection scheme.⁶

Before a hand comes between the transmitter and receiver, a signal is received through the intrinsic capacitive coupling (C_0) determined by the electrode’s size and proximity. When the hand enters, the capacitive coupling from the transmitter into the hand (C_t), hand into the receiver (C_r), and body into ground (C_g) alter the amount of signal detected at the receiver. The body is essentially perfectly conductive at these frequencies, especially when compared with the picofarad-level capacitances sketched above. If the body does not lie extremely close to either electrode, C_g dominates, and the body effectively becomes a grounded shield because electric field lines from the transmitter couple into the hand and are directed through the body to the room, away from the receiver. Thus the received signal decreases as the hand approaches. This is called the “shunt mode” (Figure 1, bottom center). A related single-electrode “loading mode” (Figure 1, bottom right) measures the current pulled from the transmitter plate into the body (via C_g) and needs only a single electrode.



1 General implementation of electric field sensing (top) and its three primary application modes (bottom).

This is how the classic Theremin and most other embodiments of “capacitive sensing” work. In “transmit mode” (Figure 1, bottom left) the body lies close to (or makes contact with) the transmitter, hence C_t dominates and the body becomes an extension of the transmit electrode. The received signal now increases as a function of body proximity to the receiver electrode.

All our interface implementations use either transmit or shunt mode, which provide measurements that prove more informative and robust than those given by loading mode. Loading mode measurements can be likened to images formed without a lens, since only one “end” of each field line is constrained by the measurement. Transmit mode works well for tracking the motion of a user in contact with a transmitter. Because the transmit mode’s received signals are a simple function of the distance between the body and receive electrode, we can easily estimate limb position.⁶

Although shunt mode does not require the user to contact a transmitting plate, the nonlinear three-body coupling (transmitter/body/receiver) is more complicated. It requires more mathematics to recover a detailed position,⁸ but we can still easily obtain simple gestural responses.⁹ The shunt mode example in Figure 1 shows a pair of electrodes—one transmitter and one receiver. Generally, by multiplexing transmit and receive functions, you can use an array of N transceiver electrodes to collect $N(N - 1)/2$ independent numbers. For simple proximity and position measurements, or when making constrained decisions about an object’s state,¹⁰ you would choose a small N ; for more complex imaging applications, you would choose a larger N .

The hardware for electric field sensing is simple and inexpensive. We designed several generations of electronics that we term “Fish,” since many species of fish use electric fields to sense objects in murky waters.¹¹ To take a user interface metaphor, mice live on a 2D surface, while fish navigate in three dimensions. Our first device, used in most of the examples in the next section, contained

2 The Flying Fish, which lets a user navigate through a virtual world by moving his hands around a monitor screen. The user sits on a transmitting electrode; the receive electrodes can be seen at the edges of the monitor.



copper tape, copper mesh, and aluminized mylar useful for prototyping electrode geometries. Mylar metalized with indium tin oxide (a transparent conductor) yields transparent electrodes. With such materials, interface design becomes “arts and crafts,” since you can quickly cut the materials to try several geometries.

An electric-field sensor array’s position resolution is a function of the mechanical placement of the electrodes, background pickup environment, and amount of noise in the

one transmitter and four dedicated receivers employing analog filtering and demodulation. The device also had 8-bit digitization and serial communication with a host computer that ran the gesture-tracking algorithms.

Our subsequent circuit, the LazyFish, is a minimal implementation of electric field sensing that can be embedded in handheld devices. The unit has four resonant transmit channels and two receive front-end gain stages. The outputs of the receive front-end channels feed directly into the onboard microcontroller’s analog-to-digital converter inputs. The received signal is then demodulated in software.

The School of Fish is our most general electric field sensing implementation. The School contains a network of intelligent transceiver electrodes, each with a dedicated 8-bit microcontroller and demodulation circuitry. An arbitrarily large array of such electrodes can be assembled by daisy-chaining them onto a common RS-485 serial bus and placing the electrodes as needed. The electrode parameters (sensitivity, filtering, response time, and transmit-receive mode) are all dynamically downloaded from a host computer, hence this system is highly adaptive and configurable. Since the School of Fish electrodes are transceivers, this hardware measures all $N(N - 1)/2$ independent pairwise capacitances. We’re using this platform to develop electrostatic imaging algorithms. By deploying more electrodes, you collect additional geometrical “projections” or increase the size of the working volume in which the body can be tracked.

Any conductor will suffice for electrodes. We found

received signal. The latter likewise depends on the demodulation filter bandwidth (hence response time). The gesture-sensing systems described in the following section have electrodes spaced from 15 to 70 cm apart and can resolve millimeter-level motion on millisecond timescales at up to a meter in range. (By taking particular care in the electrostatic geometry, these techniques are used to attain micron-level positioning across centimeters of range.¹²)

Because nonconductive, nondielectric materials don’t affect electric fields, you can hide electrodes behind furniture or build them into other common insulating objects. Ground planes can also selectively confine and shield the sensitive region—an advantage lacking in magnetic systems, which are not easily shielded.

Applications

Our first position trackers based on electric field sensing technology exploited the transmit mode. One of the original devices was a chair⁶ with a transmit electrode on the seat (providing excellent coupling into the occupant), four receive electrodes to measure hand position mounted at the vertices of a 50 × 70-cm rectangle in front of the chair, and two receive electrodes on the floor to measure foot position. This device has been used for many musical applications, where body motion was mapped in various ways to trigger and shape synthesized and sampled sounds¹³ (see <http://physics.www.media.mit.edu/creative.html>). Shortly thereafter, we adapted this configuration into a computer graphics interface, as shown in

3 The Gesture Wall as a schematic (right) and the actual installation (left) in the Brain Opera at Lincoln Center, New York, July 1996.




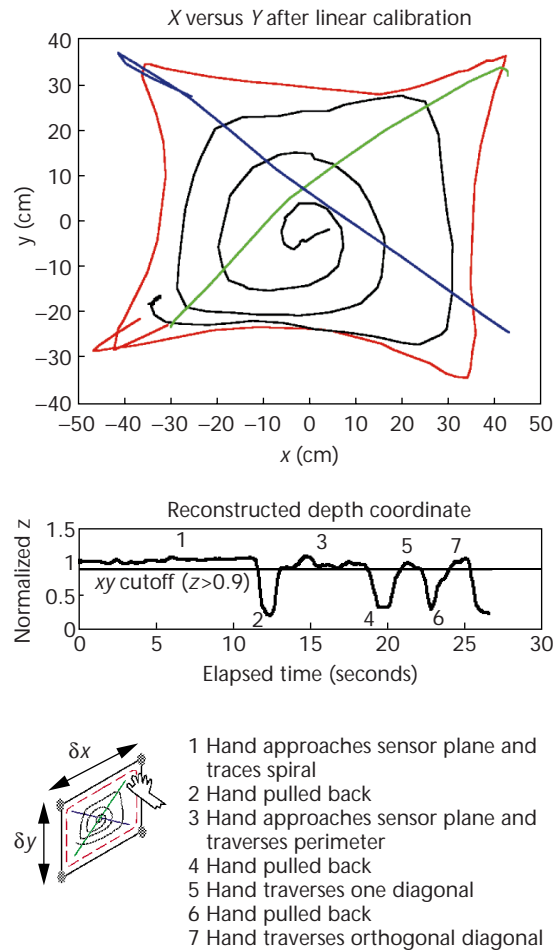
Figure 2. Here, we affixed a set of four curved receive electrodes to the corners of a computer monitor and placed the user on an office chair containing a transmit electrode. The sensors detected the proximity and up-down position of the user's left and right hands. By moving the hands differently in and out, plus collectively up and down (in a pose like that shown in Figure 2), the user can intuitively navigate through a 3D graphic landscape. This same apparatus can easily scale down to finger size for applications in which arm fatigue becomes a problem. More details are given elsewhere.¹⁴

Another set of more recent transmit mode implementations were the Gesture Wall interfaces (Figure 3) created for the Brain Opera,¹⁵ a large, touring interactive multimedia experience for the general public (see also the above URL). Because the Gesture Walls were designed more for a transient audience, they do not use a chair, but rather transmit into the body through a plate below the feet. (The shoe impedances were first calibrated out by touching a reference sensor before using the device.) The gesture sensors consisted of four receive electrodes ("budlike" objects on goosenecks in the photo of Figure 3) placed at the corners of a rear-projection screen. Figure 4 shows the reconstructed (x, y) hand position in the sensor plane (with an outstretched arm) and distance (z) from this plane, as linearly derived from actual Gesture Wall sensor data. In this example, the (x, y) hand position is plotted only when the hand lies close to the sensors' plane, detected from the measured z signal. As the hand moves in and out of the sensor plane, the plotted color changes. These shapes were all made by quickly "drawing in the air" with a hand (the actual hand trajectories taken appear at the bottom of Figure 4) and attest to this interface's utility.

In our actual implementation, the detected body motion controlled a stream of sequenced audio (composed by Tod Machover) and an interactive video clip (designed by videographer Sharon Daniel). The rear-projected video used "watery" imagery (emphasizing soft blue and green colors) as the visual metaphor, so the translucent screens appeared full of water.

The interactive graphics system (described elsewhere¹⁶) ran on a 133-MHz Pentium. We built four different interactive experiences for the Gesture Walls. Our goal was to create a role for the audience that was consistent with the general theme of the visual design, the types of information that the Fish sensor yielded, and the open, tetherless manner of interactivity. The interactive design challenge was to link the motions of a participant's hand to a light and playful interactive environment. Since visual output is such a strong feedback mechanism, participants of interactive visual devices are highly sensitive to cause-and-effect relationships between their actions and the system's reaction. Therefore, we needed to make simple visual cues that linked the physical and virtual worlds together, while keeping the experience nontrivial and as engaging as possible.

Three of the four graphical system designs used a simplified particle system in which the video content—short 10-second looped video sequences—were broken into a collection of 76,800 (320×240) particle agents. These





4 Position of a hand in the (x, y) plane of the gesture wall sensors (top) and corresponding (z) range (below), as determined from Gesture Wall data for a user freely "drawing" in the air. The actual hand trajectories that generated this data are diagrammed and summarized at the bottom.

particle agents were spatially located within the 2D image area and had a set of simple and highly localized behaviors based on Newtonian physics. These agents were part of a physical-based model that included several different "forces," such as momentum, friction, spring forces, and outside (interactive) forces. When initialized, the system placed all the agents so that the original image sequence appeared "normal" and the viewer saw the unperturbed video sequence play out. However, as an interactive participant moved a hand in the sensitive region, the particle agents closest to the hand's (x, y) coordinates were "pushed" along the vector of motion. The dynamics of each particle resulted from four forces as follows:

- momentum (the particle's tendency to travel in the same direction),
- friction (a dampening force),
- a spring force (to bring the particle back to its original spatial coordinate), and
- the user stimulus force (caused by viewers moving their hands).

Out of the simple Newtonian rules comes a swirling flow of pixel elements that float around the screen according to the physical-based model.

5 A Smart Table. Electrodes under the table (left) sense the hand's position (top), allowing free interaction with a simple virtual environment, visible on the monitor (right).



One of the Gesture Wall systems ran an algorithm called “scatter,” where the particles flew apart like a group of gnats that scatter when a person swishes a hand through them, reforming themselves when the outside forces stop. Another system ran a “sweep” algorithm, where the particles had a reduced restoring force, causing them to move in the direction of the participant’s motions and then slow down due to friction. Hence the user swept the image particles off of the screen like dust. The “swarm” system used an attractive mode of interaction, where the agents would be drawn toward the user’s motion, as if they were a swarm of bees.

We named the fourth Gesture Wall “wipe.” Unlike the particle systems described above, it used a spatial-temporal alpha-blend algorithm to allow the participant to gradually clear away one image to reveal another. As the interactive system received the (x, y, z) coordinates of the user’s hand from the Fish sensor, it updated an alpha-blend matrix in proportion to the vigor exerted by the participant, correspondingly revealing the underlying image where the user stimulus occurred. A time-decay system gradually restored the overlying image, to create the impression that the visual surface was a piece of cold glass that fogged up and could be wiped clean, only to have it fog up once again.


We also designed several different graphics interfaces based around shunt-mode sensing. By mounting a transmitter and three receivers underneath a wooden table, we created an active region above the table in which the

naked hand interacted with a virtual space. In this demonstration (see Figure 5), a simple computer graphic image corresponding to the activated region appeared on a monitor. The space contained an iconic representation of a hand and a cubic object. By moving your hand above the table, you can control the hand’s 3D position in the virtual space. When the virtual hand touches the cube, the hand closes, and the cube can be moved around the space. When the cube is set down again, the virtual hand opens, and the cube remains. We’re currently working on algorithms that use additional electrodes to infer the size and orientation of the user’s hand, so the cube can be picked up or released anywhere in the space.

With the sensors in the same under-the-table configuration, we performed simple gesture recognition on the signals to create a “smart breakfast table” that lets users turn the pages of an electronic newspaper. Flipping the hand left turns the page forward, flipping the hand right turns the page back, and moving the hand up or down changes the newspaper’s sections. An electronic newspaper can be considered a virtual environment, but not one that needs to be navigated in an immersive way. An activated breakfast table that lets users flip naturally through a virtual paper is vastly preferable to a cumbersome point-and-click or mouse-driven solution.

A similar gestural interface we developed is called the FishPad—a Fish-based input device for exploring visual space through arm movements. Likewise, the FishPad is a flat, shunt-mode array, made from a central

6 A user interacting with the FishPad using the one-handed configuration to browse through a watch catalog (right) and a two-handed implementation to play a virtual “pin-the-tail on the donkey” (left).



transmit electrode and surrounded by four receive electrodes arranged in a north-south-east-west configuration. These electrodes detect the 3D location of the hand, which zooms and pans through a high-resolution 2D image (Figure 6, left). The device's input space is directly mapped onto a selected region of the image. Unlike a fly-through, the feedback is immediate. Returning the hand to the previous position moves the users back to a known reference point.

You can use FishPads with one or both hands. In the two-handed configuration (see Figure 6, right), the primary hand zooms and pans through a space, while the secondary hand does meta-operations on the image (separate FishPad arrays track each hand). In one example, one hand zooms into a complex discrete automata program (with many cells on the screen dynamically changing), while the other hand examines the contents of any cell currently in the viewing region. In a more entertaining example (Figure 6, lower right), we created a virtual "pin the tail on the donkey" game. The primary hand explores the space looking for tailless donkeys, and the secondary hand pins on tails.

The capabilities of the Fish influenced the FishPad design decisions and pointed to areas of further study. As mentioned earlier, the shunt-mode Fish sensors have a nonlinear response to distance. Due to slight differences in any one person's body capacitance and grounding, they don't respond equally to everyone. The solution was not to recalibrate for each person, as with the Gesture Walls, but to use reasonable defaults and simple transformations on the input to keep the position updates high. Able to operate at 50 Hz or more, the Fish had no problem keeping up with a 20-frame-per-second rate, which was generally fast enough for people to quickly learn the response of the FishPad to their movements.

The two-handed FishPad proved much more challenging than using only one hand, suggesting that the hands would work better together instead of decoupled in separate spaces. Fatigue and loss of accuracy was a problem with prolonged use, which suggested a new FishPad design where the hands could be at rest or supported. We anticipate that a future implementation with the newer Fish hardware and a new physical design will overcome these initial problems. The resulting improvements in bandwidth and accuracy would yield a larger input area that could include both hands in a single sensing region, together with the ability to measure more subtle gestures.

Conclusions

Sensing with electric fields is not new, but we have found that lurking behind what appears to be a trivial exercise in capacitance is a deeper inference problem with significant implications for user interfaces. Given the basic inexpensive instrumentation needed to make these measurements, physical interface design reduces to shaping electrodes with easily available materials. This lets user actions on and around familiar objects be sensed responsively and reliably, activating the space around the objects without introducing any apparent intrusive technology. Users tend to react to such systems by first marveling at the "magical," causal control with-


out any apparent mechanism for the connection. They then quickly forget about the presence of any intervening technology and focus on the application. ■

Acknowledgments

We'd like to acknowledge the many colleagues and collaborators who contributed to the evolution of this work, including Tom Zimmerman, Tod Machover, and John Maeda. This research has been supported by the Things That Think consortium and our other sponsors at the MIT Media Laboratory.

References

1. P. Maes et al., "The Alive System: Full-Body Interaction with Animated Autonomous Agents," *ACM Multimedia Systems*, Vol. 5, No.2, 1997, pp. 105-112.
2. J. Fraden, *AIP Handbook of Modern Sensors*, American Institute of Physics, New York, 1993.
3. T.G. Zimmerman et al., "A Hand Gesture Interface Device," *Proc. CHI+GI 87 Conf. on Human Factors in Computing Systems*, ACM Press, New York, 1987, pp. 189-192.
4. H. Murry and M. Schneider, "Virtual Reality Position Tracking," *Circuit Cellar INK*, No. 60, July 1995, pp. 24-29.
5. L.S. Theremin, "The Design of a Musical Instrument Based on Cathode Relays," reprinted in *Leonardo Music J.*, No. 6, 1996, pp. 49-50.
6. J. Paradiso and N. Gershenfeld, "Musical Applications of Electric Field Sensing," *Computer Music J.*, Vol. 21, No. 2, 1997, pp. 69-89.
7. M.V. Mathews, *Three Dimensional Baton and Gesture Sensor*, US Patent No. 4,980,519, United States Patent Office, Washington, DC, 1990.
8. J.R. Smith, "Field Mice: Extracting Hand Geometry from Electric Field Measurements," *IBM Systems J.*, Vol. 35, No. 3 & 4, 1996, pp. 587-608.
9. T.G. Zimmerman et al., "Applying Electric Field Sensing to Human-Computer Interfaces," *Human Factors in Computing Systems*, ACM Press, New York, 1995, pp. 280-287.
10. T. Costlow, "Safety, Comfort Drive Electronics," *Electronic Engineering Times*, March 31, 1997, pp. 136-138.
11. I. Wickelgren, "The Strange Senses of Other Species," *IEEE Spectrum*, Vol. 33, No. 3, 1996, pp. 32-37.
12. J. Paradiso, "New Technologies for Monitoring the Precision Alignment of Large Detector Systems," *Nuclear Instruments and Methods in Physics Research A*, Vol. 386, 1997, pp. 409-420.
13. T. Machover, *Media Medium* (musical score), Ricordi, Milan and Paris, 1994.
14. D. Allport, E. Rennison, and L. Straussfeld, "Issues of Gestural Navigation in Abstract Information Space," *Proc. CHI 95: Human Factors in Computing Systems*, ACM Press, New York, 1995, pp. 206-207.
15. J. Paradiso, "Electronic Music Interfaces: New Ways to Play," *IEEE Spectrum*, Vol. 34, No. 12, Dec. 1997, pp. 18-30.
16. C. Dodge, *The Abstraction, Transmission, and Reconstruction of Presence: A Proposed Model for Interactive Computer Art*, masters thesis, MIT Media Lab, Cambridge, Mass., Aug. 1997.



Joshua Smith is a PhD candidate and research assistant at the Physics and Media Group of the MIT Media Laboratory. His central research topic is the problem of inferring body geometry from electric field measurements. He has also developed new instrumentation and signal processing techniques for electric field sensing. He has an MS from the Media Laboratory at MIT, an MA in physics from Cambridge University, UK, and a double BA in computer science and philosophy from Williams College, Williamstown, Pennsylvania. He has had internships at the Santa Fe Institute, Los Alamos National Lab, the Yale University Department of Computer Science, and NASA's Goddard Institute for Space Studies.



Tom White is a masters candidate at the MIT Media Laboratory, working in the Aesthetics and Computation group. His work focuses on creating new ways of communicating with computational media, often through the design of new hardware devices. He is currently finishing his masters thesis about a new family of interfaces known as Liquid Haptics. Information about his present research can be found at <http://tom.www.media.mit.edu/>.



Christopher Dodge is an independent interactive media producer and software developer in Cambridge, Massachusetts. He has worked several years in the field of digital signal processing software development. He has a BA from New York University (1991) in an interdisciplinary study of computer science, film/video, and music composition. He earned his MS at the MIT Media Laboratory (1997), focusing on the development of real-time interactive media environments. He was awarded a two-year Artist/Researcher-in-Residency at the Center for Art and Media Technology (ZKM) in Karlsruhe, Germany. As an interactive computer artist, his works have been shown around the world at such festivals as the MultiMediale 4, Ars Electronica 96 and 97, and the Artex 97 Media Art Biennial, in which he won the grand prize.



Joseph Paradiso is a research scientist at the MIT Media Laboratory, where he develops new sensing technologies for responsive environments. He is currently the technology director for the Things That Think program. He received a BS in electrical engineering and physics from Tufts University in 1977 and a PhD in physics at the Massachusetts Institute of Technology in 1981. He has held positions at the Draper Laboratory in Cambridge, Massachusetts and the Swiss Federal Institute of Technology (ETH) in Zurich designing high-energy physics detectors, spacecraft control algorithms, and sensor systems. He has also designed several music synthesizers and electronic musical interfaces. He is a member of the IEEE, American Institute of Aeronautics and Astronautics, American Physical Society, and Sigma Xi.



Neil Gershenfeld leads the Physics and Media Group at the MIT Media Lab, and co-directs the Things That Think industrial research consortium. His laboratory investigates the interface between the content of information and its physical representation. He has a BA in physics with High Honors from Swarthmore College, Pennsylvania and received a PhD from Cornell University, New York studying order-disorder transitions in condensed matter systems. He was a Junior Fellow of the Harvard Society of Fellows, where he organized an international study hosted by the Santa Fe Institute on forecasting the behavior of complex systems.



David Allport is a principal engineer at VSIS, a wholly owned subsidiary of Mitsubishi Electric America in Sunnyvale, California. He received a BA in Sanskrit and Pali from Oxford University, England in 1978 and a PhD in Indian philosophy from Oxford in 1982. After completing his masters thesis in artificial intelligence and knowledge-based systems at Sussex University, he spent several years as a research fellow working on computer natural language understanding. He was a Hewlett-Packard visiting scientist at the MIT Media Lab from 1994 to 1995.

Contact the authors at the Physics and Media Group, MIT Media Lab E15-495A, 20 Ames St., Cambridge, MA 02139, e-mail phmreq@media.mit.edu.