

An installation of interactive furniture



by O. Omojola B. Schoner
E. R. Post P. R. Russo
M. D. Hancher R. Fletcher
Y. Maguire N. Gershenfeld
R. Pappu

We report on a project that explored emerging technologies for intuitive and unobtrusive information interfaces in a compelling setting. An installation at the Museum of Modern Art, New York, was part of a public exhibit and included an interactive table that presented information associated with the exhibit to the gallery visitors without visible conventional computing elements. The enabling devices included noncontact sensing of low-cost tags in physical icons, electrostatic detection of hand location in three dimensions, and sensor fusion through lightweight Internet Protocol access.

Advances in computing technology and the advent of the Internet have brought about dramatic increases in the amounts of information available to people. However, the interfaces used to handle this information have not changed much with time: the mouse, the keyboard, and the computer monitor, connected to a box containing the processor and its peripherals. The mouse and the keyboard are nonintuitive interfaces. Touch-typing, as an example, is a skill that takes significant experience or training to acquire, because the layout of the standard computer keyboard groups letters in a nonintuitive fashion. Also, the boxes for the processor and the monitors have, over time, gotten smaller, quieter, and cheaper, but are essentially the same, typically unsightly, and occupy visible space. More importantly perhaps, they cut users off from the rest of the world: sitting in front of a computer often means ignoring the surrounding environment.

Much technology has been developed in disparate pieces to enable the vision of embedding computing

power into everyday objects such as clothing and furniture. There have been some projects such as the Brain Opera^{1,2} that cross the boundaries of individual, small-scale demonstrations and have explored the issues of interconnecting these disparate pieces, but these have been few and far between.

Based on common interests in the potential application of these technological goals, the Museum of Modern Art (MoMA), New York, and the MIT Media Laboratory agreed to a collaboration driven by the desire to increasingly use technology in their exhibits without making aesthetic concessions. Their main motivation was to use *smart spaces*³ in museum exhibits, without the obvious elements of the associated computing. They offered a very useful error metric that is lacking in a laboratory: a very high aesthetic sensibility. In many Media Lab demonstrations, the visual appearance is a distant second to the functionality of the technology. The museum was unwilling to compromise the appearance of the exhibit spaces as a result of any such limitations.

The collaboration centered on the Un-Private House, an exhibition of architectural projects that offered interesting blends of public and private spaces. A significant amount of information about the projects was in electronic form, and, traditionally, this information is presented through computer

©Copyright 2000 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

kiosks placed away from the exhibit. The curator, Terence Riley, and the exhibition's production manager, Andrew Davies, were eager to have this information presented to the exhibit's visitors within the rich context of the exhibit itself. They wanted to spark more social interaction and discussion, and wanted the space to react to the presence of the visitors.

These ideas had elements in common with much of the work done within the Media Lab on smart spaces³ and provided a very interesting environment to apply much of our technology. From the beginning there were a number of significant design objectives. There obviously could not be any exposed technology, since the clash with the appearance and layout of the rest of the exhibit would detract from the museum experience. The museum wanted the technology to be used in ways that enhanced conventional functions within the home (as an example, the interactive welcome mat, described later, is not so much a new idea, but an enhancement of the traditional welcome mat). The installation needed to be reliable enough to handle the crowds of museum visitors without permanent supervision. Finally, the information to be displayed was in a range of formats that included text, images, and video, and the chosen interface had to provide display and browse access to all of these in a straightforward fashion.

During one of the visits Terence Riley paid to the Media Lab, he was exposed to the work of Hiroshi Ishii on tangible interfaces⁴ and was led to an interest in similarly tactile interfaces. Given the desire for a simple, intuitive interface, the first proposals planned to use a tagged object as a physical icon to start the display of information.

A dining table offered the most potential as the physical surface on which the information would be displayed. It also had the added benefit of both symbolizing, and providing a good framework for, the kind of social interaction that was desired.

Initially, a vision-based system similar to that designed by Flavia Sparacino and Kent Larson for the Unbuilt Ruins exhibit at the MIT Compton Gallery⁵ would track the position of various objects on the table. The user could place these at certain points on the table to interact with the information. This idea was abandoned because of concerns about a cluttered interface that would be confusing or require extensive supervision to prevent people from walking away with the objects.

Another of the initial ideas had the welcome mat track the motion of people across it and respond to this in some way. This was first tried using a mat with electric-field-sensing electrodes and a projected image that followed the path of the person walking across it. Issues of robustness for this tracking method caused this idea to be dropped as well.

Of some interest is the fact that in the final design the tracking methods used were the same, but used for different applications: the interactive table interface used electric-field sensing, whereas the welcome mat installation used a vision-based system.

There was a significant amount of information to be displayed, and there were both human resource limitations and a lack of interest in content development, editing, and formatting for the project. MoMA and NearLife, Inc., a Boston-based interactive entertainment company, handled those functions as well as the purchase, setup, and maintenance of the computers that drove the various parts of the installation.

Overview

The gallery arrangement for the Un-Private House was modeled after a house, with sections of the gallery and the associated furniture imitating the various rooms and functions present in a house. The exhibit had 26 architectural projects on display, and there were models and other display pieces for each house distributed throughout the gallery on the tables and other pieces of furniture. Figure 1 shows the layout of the gallery, with the entrance to the gallery in the upper right corner and the dining table as a round gray and white circle in the lower left.

The interactive installation consisted of two parts. The first part was the interactive welcome mat stationed at the gallery entrance, as shown in Figure 2. An image of a traditional welcome mat appeared on the floor outside the entrance, and the image moved in response to the motion of the visitors across it.

The main part of the installation was the interactive table, built around the gallery's dining table. The table was an 8-foot-diameter, five-legged table with a solid-white, 3/4-inch-thick Corian** surface that seated eight people. Each of the eight place settings featured a display projected from above that was 18 inches wide by 14 inches high on the tabletop. In the center of the table was a 5-foot-diameter lazy Susan, also made of solid white Corian, with recessed

Figure 1 The Un-Private House gallery floor plan



holes on its outer edge. Placed in these indentations were a set of 26 coasters (one in each indentation) with an image of a different architectural project on the face of each coaster. The outer edge of the lazy Susan was illuminated with a number of fiber optic light sources from the drop ceiling above (which also served as the projector housing), and there was a top-projected display above it.

The interface presented by an inactive place setting is shown in Figure 3, with the blue circle in the upper right corner displaying the words “Place coaster here to begin.” A pictorial animation of picking a coaster and placing it on the active spot was periodically displayed in the center of each place setting.

After the user chooses a specific coaster from the lazy Susan and places it over the blue circle, the display changes to show a floor plan of the project in traditional blue print. Highlighted on the display are a number of “hot spots,” places within the specific house with information that may be of interest to the user. When a hot spot is touched the associated information is displayed in the lower left corner, as shown in Figure 4. This information included text, images, and animations that showed interesting parts of the house, comments from the museum curator, the house’s architect, and the client, as well as information about the technology used to build the table itself.

Depending on the content, an extra hot spot would appear that allowed a user who wished to share a certain image with the other users at the table to send

Figure 2 Welcome mat



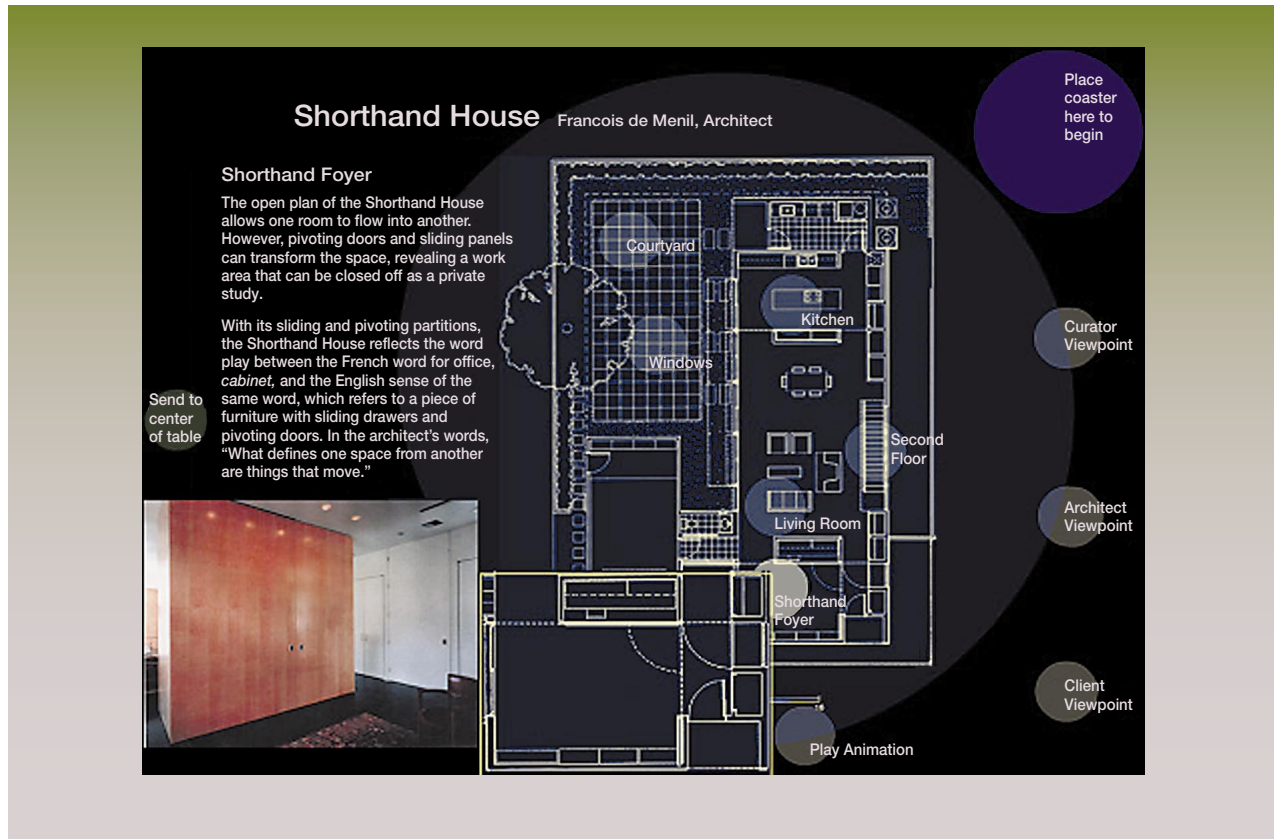
Figure 3 Photograph of default place setting



it to the center. When the user touched the “send-to-center” hot spot, the associated image was projected onto the center of the lazy Susan, facing that user. Another user at the table interested in looking at the image could rotate the lazy Susan, which would cause the image to also rotate. Comments about the exhibit entered on the museum’s Web site by the public were displayed in the center of the table when not in use by any of the place settings. The lazy Susan with a projected image is shown in Figure 5.

The remainder of this paper is organized as follows. The next section describes the welcome mat installation and its use of computer vision technology. Then, the physical icons used in the dining table exhibit are described and the use of noncontact sensing technology. The following section focuses on the

Figure 4 Photograph of activated place setting



gestural interface and the use of electric-field sensing for enabling the use of the human hand as a pointing device. Next, the lazy Susan installation is described including the mechanism, based on quadrature encoding, for sensing its speed and direction of rotation. A section follows describing the local area network and its role in the communication infrastructure for the installation. The following section on system software describes the browser-based client software for displaying the information on the various exhibits. A brief description of the display hardware setup is given in the next section. The final section summarizes the lessons learned from this project.

Welcome mat

The interactive welcome mat was included in the installation to give visitors a hint of what they were about to encounter inside the gallery. It was installed

at the entrance to the gallery and consisted of an image projected from the ceiling that responded to the motion of the visitors as they walked across it. A similar installation is reported in Reference 6.

As noted earlier, the original design called for electric-field sensing as the core technology supporting the welcome mat application, whereas computer vision was used as the gesture-sensing technology for the dining table. However, further reflection led us to implement the gesture detection using electric-field sensing, and to use the computer vision for the welcome mat.

The schematic drawing of the installation is shown in Figure 6. A video camera and a projector were arranged in an approximately coaxial configuration directly above the projected area. The camera detects changes in the visual field, changes caused by

Figure 5 Photograph of lazy Susan with projected image



visitors moving across the carpet, and presents these changes to a simple physical model.


The physical model behind the interactive image is based on an $M*N$ lattice of unit masses, all located above a reference plane. Each mass is connected to the reference plane by a spring with a spring constant K and a damping constant D . The input image from the projector is averaged down to a resolution of $M*N$.

As a visitor moves across the area of the mat, a mass proportional to the difference between the images of two adjacent frames is created in a second plane located above the lattice. This additional mass gravitationally attracts the unit masses and causes them to move toward regions of activity in the lattice. Figure 7 shows a representative image of the lattice of masses. The image activity causing the distortion is in the upper right corner, and the red dots are the individual masses.

The equation representing the dynamics of the physical model is:

$$ma = F_g - F_s - F_d$$

Figure 6 Schematic of the welcome mat installation




Here the mass m is unity, a is the acceleration experienced by the mass, F_g is the force due to gravity, F_s is the restoring force, and F_d is the force due to damping. Furthermore,

$$F_g = -\frac{Gm}{R^\rho}$$

where G is a gravitational constant and R is the distance of the mass from its reference location. The exponent ρ controls the interaction distance and thus the range that the masses move. Short-range inter-

Figure 7 Image of the welcome mat showing image activity and the unit mass lattice



action requires a higher value of ρ than a longer range one.

$$F_s = -Kx$$

where K is the spring constant, x is the distance the spring is extended, and

$$F_d = -Dv$$

where D is the damping constant and v is the velocity of the mass.

Thus

$$ma = -\frac{Gm}{R^\rho} + K(x - x_r) + Dv$$

where x_r is the reference point for the system (chosen to be 0 for the ensuing derivation). This equation may be solved numerically using a method of differences to yield the new position (x_{t+1}) of each mass in terms of its current position (x_t) and its position 1 time period earlier (x_{t-1}) and this is given by

$$x_{t+1} = \left(\frac{D}{m\Delta t} + \frac{K}{m(\Delta t)^2} + 2 \right) x_t + \left(\frac{D}{m\Delta t} - 1 \right) x_{t-1} - \frac{G}{R^\rho} \left(\frac{1}{\Delta t} \right)^2$$

where Δt is the time difference between each position (which we arbitrarily chose to be unity). An identical derivation can be performed for the change in position in the y direction.

The final image of a welcome mat is texture mapped onto the distorted lattice of unit masses and rendered using OpenGL**. Figure 8 depicts the data pipeline of the welcome mat system.

Physical icons

The coasters serving as physical icons are depicted in Figure 9. The person holding the coaster, and only that person, has access to the information on the associated exhibit. Each coaster is approximately 3.5 inches in diameter, 0.75 inches thick, and is made of a hollowed-out Corian base and a clear acrylic cap. An image of the associated exhibit is attached to the bottom of the cap, as shown in Figure 9.

One component of the tagging system used to identify the coasters and retrieve the associated data is a radio frequency identification device (RFID) tag embedded in the coaster. The tag, shown in the foreground of Figure 9, is a simple LC (inductor-capacitor circuit) resonator built from a 3-inch diameter printed copper coil and a ceramic chip capacitor soldered onto the coil. Using the same coil size with a different capacitance for each coaster results in resonant frequencies ranging between 5.4 MHz and 13.2 MHz, in increments of 300 kHz. When compared with conventional silicon RFID tags, our tags have lower cost and are easier to replace. The 300-kHz value chosen for the resonant frequency separation takes into account the manufacturing tolerances of the capacitors and the safety margins desired.

Each place setting at the table is equipped with a tag reader as shown in Figure 10. The tag reader, whose design is based on a design developed at the Media Lab,⁷ measures 6.5 inches wide by 5 inches long, and is located underneath the upper right corner of the display area at each place setting. The reader is mounted above the tauFish electrode array in order to minimize signal losses through the copper electrodes. (The tauFish array is described in the next section.)

The tag reader operates in a frequency range between 5 MHz and 40 MHz, because tags operating in this range are relatively cheap, and the transmit antenna can be implemented as a single-turn coil etched onto the printed circuit board. The tag reader

Figure 8 Functional block diagram of the welcome mat system




Figure 9 Photograph of the coasters with an exposed resonator in the foreground




is powered by a 9-volt power supply. Access to the measurement data is through an RS-232 serial port.

The functional block diagram of the tag reader is included in Figure 11. The design is based on a PIC16C76 microcontroller. The frequency is generated by a single direct digital synthesis (DDS) device, able to generate any frequency from DC up to 60 MHz using a lookup table and a digital-to-analog converter. The output of the DDS is passed through a four-pole low-pass filter to remove the higher-order alias frequency harmonics, and the filtered signal is then amplified by a fast operational amplifier.

The presence of a resonator is detected by a 50 Ω directional coupler designed to measure antenna loading. In this configuration, the untuned single-turn antenna presents an unmatched load and the reflected power is channeled through the directional coupler output. When a resonant structure is placed within range of the antenna, however, the impedance of the resonator dominates the response. Moreover, when the frequency signal equals the resonant

Figure 10 Photograph of the tag reader



frequency a dip in the reflected power results. This method of detection eliminates the need for a broadband tuned antenna, and because the antenna is untuned, the tag reader is immune to spurious responses caused by stray reactance in the environment.

The directional coupler's output is rectified by a diode detector and the signal is sampled by the analog-to-digital converter built in the microcontroller. The microcontroller scans the output range of the DDS device, from 5.4 MHz to 13.3 MHz in increments of 100 kHz, and records an 8-bit sample at each frequency. The sampled data are transmitted through the serial port as 80-byte data packets. A sampling cycle of 105 milliseconds was found adequate.

The serial port output is connected directly to a Filament card (described later in the subsection on net-

Figure 11 Functional block diagram of the dining table installation

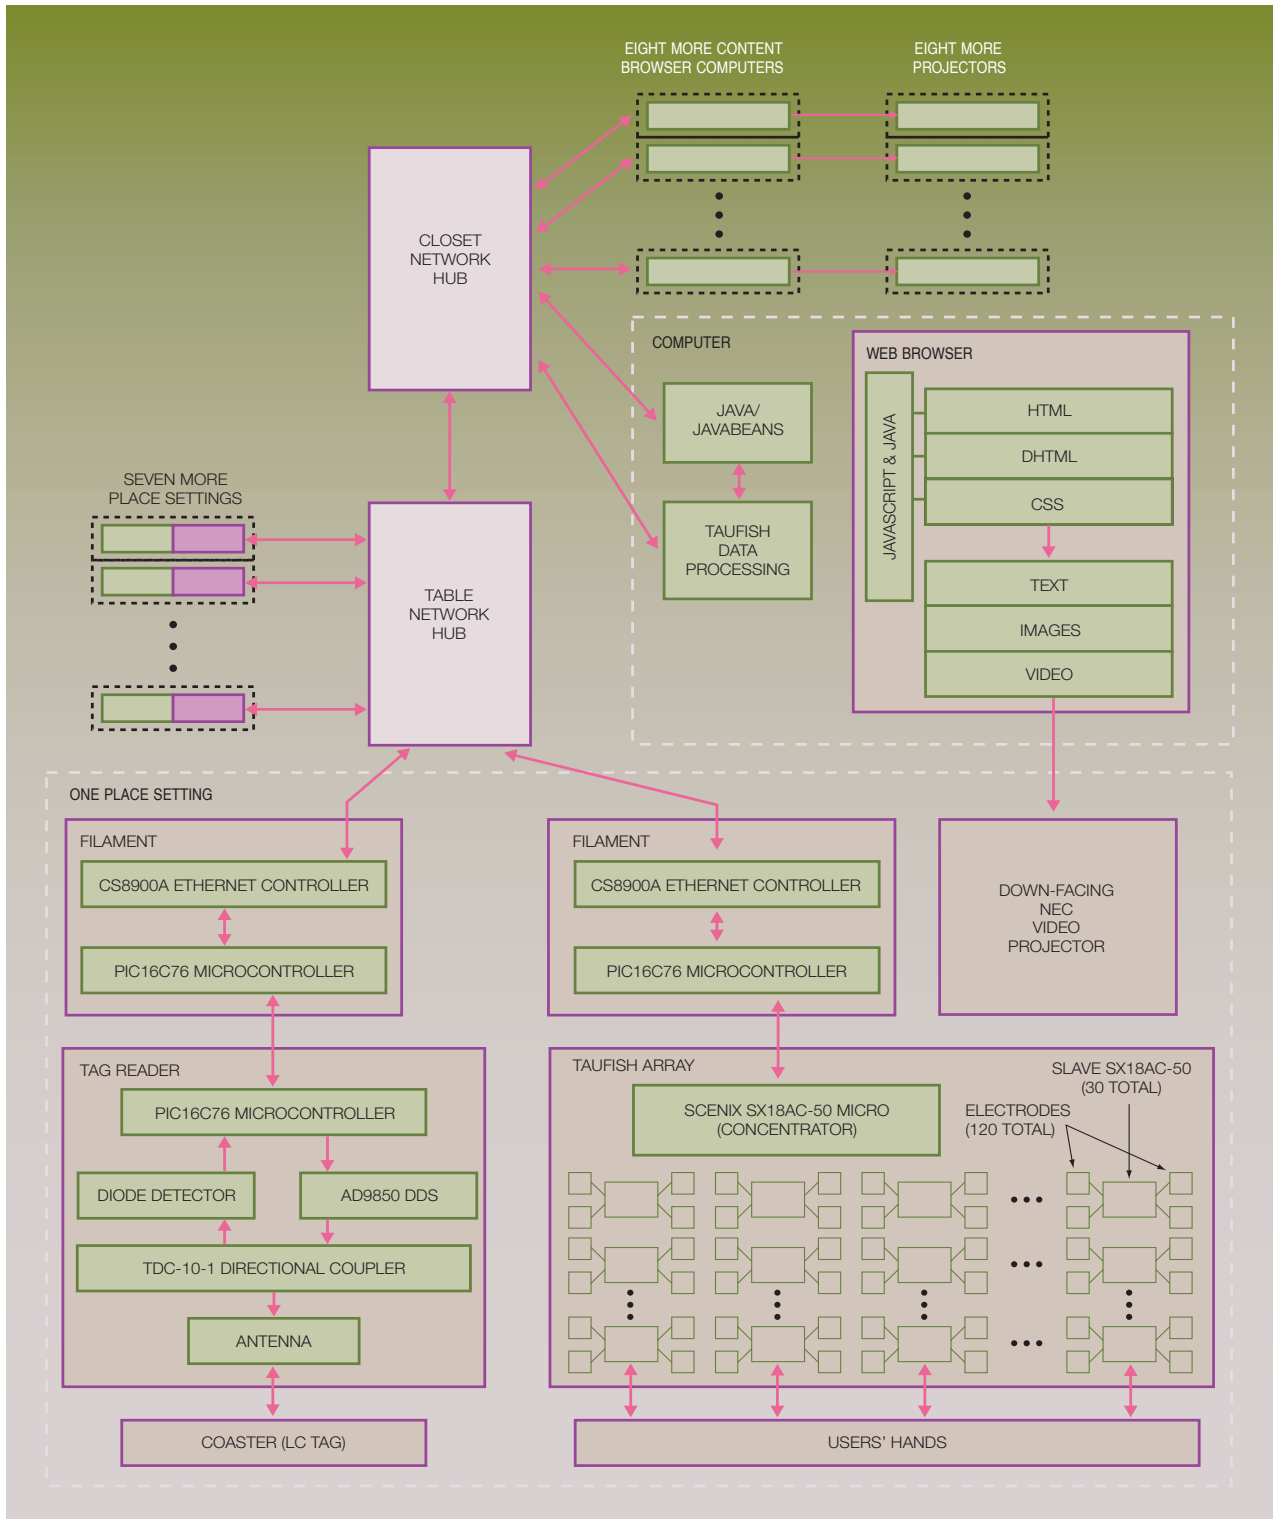
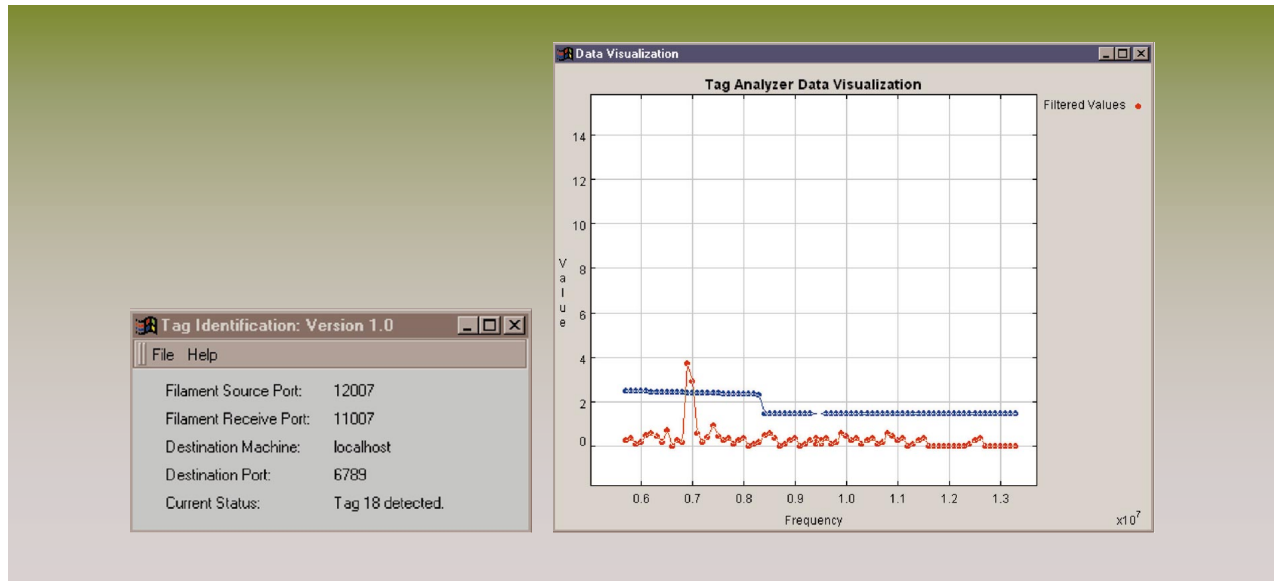


Figure 12 Tag reader interpreter and data visualization



work infrastructure), which broadcasts the data on the local Ethernet network. Remote computers connected to the network take the data off the network and run the data through a tag reader interpreter program, running as a background process.

The interpreter, whose graphical user interface is shown in Figure 12, is written in Java** and takes the following input parameters: the TCP/IP (Transmission Control Protocol/Internet Protocol) port number the Filament transmits to the port number on which the Filament receives, the network address of the machine running the specific place setting's display system software, and the port number on which the system software listens for updates. The mapping between tags and exhibits is specified in a file accessible to the interpreter.


The interpreter receives the broadcast data originating at the tag reader and filters the data using a four-tap wavelet filter. The filter is needed to compensate for certain hardware limitations in the tag reader, remove environmental noise, and eliminate the DC bias added by the analog-to-digital converter. The output of the filter is a 76-point sequence scaled between 0 and 15. The centroid of this sequence is calculated and the resultant number is checked

against the contents of the database file. The database file contains pairs of numbers: the first value in each pair is a tag number; the second value is the associated centroid value. The closest match to the calculated centroid is found, and this identifies the information content to be called up (or switched off). A simple state machine allows the interpreter to track only updates (i.e., changes in the tag reader's data that indicate either introduction or removal of a tag from the area of the antenna). The interpreter, upon detecting an event, sends a data packet to the appropriate place setting's display system software, which takes the appropriate action (either loading or unloading a particular project's information).

Gestural interface

To maintain the illusion of direct interaction with a projected image, we chose a form of electric-field sensing (EFS)⁸ to determine the position of the user's hand over the image. The musical instrument known as the Theremin⁹ is a classic example of an EFS device. Invented in 1917 by Lev. S. Theremin, this eponymous instrument senses small variations in the capacitance of two antennae induced by the user's body and converts these measurements to an audible form.

Figure 13 Four generations of electric-field-sensing instrumentation



In this section, we describe the instrumentation and algorithm used to detect the user's interaction with an image projected onto a surface. The instrumentation is based on the tauFish capacitive sensor (described below). Capacitance is measured at points uniformly covering the image region. The measured capacitance vector is then compared to vectors calculated using a forward electrostatic model, which we describe. The calculated vectors are the charge or capacitance distributions expected when the user's finger is placed at various points on the image surface. These finger locations are referred to as "hot spots." If sufficient total capacitance is observed over the entire image, then the calculated vector most closely matching the observed vector indicates the user's interaction with a particular hot spot.

Our early work with EFS^{10,11} employed synchronous detection to measure the perturbation of weak, quasi-static (50–500 kHz) electric fields. The evolution of the EFS hardware is shown in Figure 13. The first prototype transmitted on two electrodes, received on one electrode and comprised several boards' worth of hand-wired analog circuitry tucked into a 19-inch rack-mount cabinet. The first production-level system was dubbed the *Fish* (named after the aquatic animals that use electric fields to sense their envi-

ronments, including the gymnotiforms and mormyri-formes). It consisted of two transmitters, four programmable-gain receivers, and a microcontroller equipped with MIDI (Musical Instrument Digital Interface) and RS-232 serial interfaces used for data acquisition. The second system (*SmartFish*) incorporated two transmitters, eight programmable-gain receivers, a fast 8-channel analog/digital (A/D) converter, and a 32-bit, 20-MHz DSP used to perform synchronous detection in software. This complex design was ultimately abandoned because of its sensitivity to weak fringe fields and, more importantly, to its own digital noise.


The most recent design, named *LazyFish* and shown as the smallest board in Figure 13, represents a shift in our design of synchronous EFS instrumentation. Instead of using a powerful DSP and several configurable analog input stages, the *LazyFish* incorporates a simple microcontroller, four transmitters, two fixed-gain receivers, and a synchronous undersampling algorithm implemented in firmware, a design which outperforms the earliest Fish boards in speed, sensitivity, and simplicity.

In the quest for an even simpler EFS sensor we eschewed synchronous detection, choosing instead to implement time-domain measurement of capacitive loading. The resulting module, known as the *tauFish* and shown in Figure 14, is even smaller than its predecessors, and controls four channels in a loading-mode measurement.

The 5×6 array of tauFish modules positioned underneath each place setting, as shown in Figure 15, measures 18 inches by 16 inches. Each unit senses the capacitance of four 1.5-inch-by-1.5-inch electrodes. The units are connected by a multidrop serial bus attached to a microcontroller located on the array board. The microcontroller sequences the sensing and readout operations of the tauFish units and produces a sensor data stream. The entire array board outputs data at 115 Kb/s, and full frames of data are generated at a rate of about 10 Hz (where a full frame of data is a 24-bit value for the capacitance measurement for each one of 120 electrodes).

The physical model underlying the EFS application is a forward model¹² of induced charge on the electrodes. Consider the case of a point charge above an infinitely conducting plane. If a charge q is located at (x, y, z) , then the method of images can be used to find the electric field at the surface of the conducting plane. Since the plane is considered a per-

Figure 14 Top and bottom views of four-channel tauFish



fect conductor, there will be no transverse field, and the perpendicular electric field will be determined by the presence of the charge above the plane and its inverse reflection below the plane. The electric field at a position $(\mu_x, \mu_y, 0)$ will be

$$E_z = \frac{qz}{2\pi\epsilon r^3}$$

$$= \frac{qz}{2\pi\epsilon[(x - \mu_x)^2 + (y - \mu_y)^2 + z^2]^{3/2}}$$

where ϵ is the electric permittivity.

Integrating this field over an infinitesimal Gaussian “pillbox” at the surface gives the surface charge density σ :

$$\int \rho dV = \int \nabla \cdot \vec{D} dV$$


$$Q = \int \epsilon E \cdot d\vec{A}$$

$$= \epsilon E_z A$$

$$\frac{Q}{A} = \epsilon E_z$$

$$\sigma = \frac{qz}{2\pi[(x - \mu_x)^2 + (y - \mu_y)^2 + z^2]^{3/2}}$$

Figure 15 tauFish array with 30 tauFish (120 electrodes)

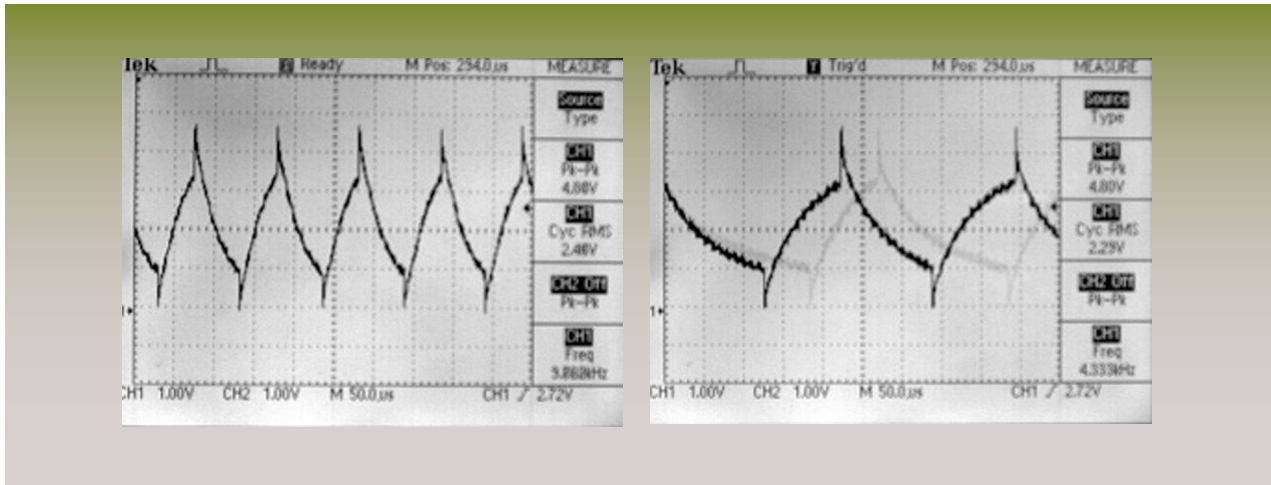


Now assume that (infinitesimal) square electrodes tile the plane. The induced surface charge for each electrode will be

$$Q = \int_{x_0}^{x_0+\delta x} \int_{y_0}^{y_0+\delta y} \sigma(x, y) dx dy$$

$$= \frac{q}{2\pi} \arctan \left[\frac{(x_0 - \mu_x)(y_0 - \mu_y)}{z \sqrt{(x_0 - \mu_x)^2 + (y_0 - \mu_y)^2 + z^2}} \right]$$

Figure 16 Electrode voltage vs time during tauFish capacitance measurement cycle. The first plot is measuring the parasitics and the electrode self-capacitances. The second plot is measuring the parasitics and the electrode-hand capacitances.



Note that the charge q is only a leading term, while the surface charge distribution is shaped strongly by z ; when z is small, the distribution is narrow, and as z increases the charge distribution broadens. The total charge varies linearly as the capacitance ($Q = CV$) and is proportional to the capacitive load due to charge induction.

The tauFish measurement strategy is deceptively simple. One pin of a microcontroller is alternately charged and discharged by using a second pin to pulse a voltage source through a known resistance R . The first pin is also attached to an electrode that serves as the capacitance C under test. The CMOS (complementary metal-oxide semiconductor) input at the pin is configured to have hysteretic thresholds at $0.3V_{DD}$ and $0.7V_{DD}$ (where V_{DD} is the positive supply voltage of the microprocessor). The electrode charging cycle covers the range $[0 \dots 0.7V_{DD}]$, while the discharging cycle covers the range $[V_{DD} \dots 0.3V_{DD}]$. The capacitor voltage evolves in time as

$$V = Ae^{-t/\tau}$$

where τ is the charging time constant $\tau = RC$. It is the measurement of this time constant that gives the tauFish its name. The charging and discharging thresholds are as described above, so the measurement time is found from

$$1 - e^{-t/\tau} = 0.7$$

$$\text{to be } t \approx 1.21\tau$$


In this implementation, charge is transferred during 40 nanosecond intervals (two instructions at an execution rate of 50 MHz) through a $2M\Omega$ resistance. Since the charging time is linear in C , we find that one current pulse corresponds to charging a capacitance of

$$C = \frac{\delta t}{R} = \frac{(40 \times 10^{-9}s)}{(2 \times 10^6\Omega)} = 2 \times 10^{-15}F$$

or 20 femto-Farads. Because this measurement is done in the time domain with minimal analog circuitry (i.e., one additional resistor), its precision depends on the stability of the clock as well as the microprocessor's timing jitter. The measurement can be made arbitrarily precise by averaging several measurement cycles. Figure 16 shows the voltage on an electrode as it is charged and discharged during a measurement cycle. Finally, to minimize the first-order effects of 60-Hz noise that couples into the measurement, the values of the charging and discharging cycles are summed.

While it is simple to evaluate the forward electrostatic model going from a known set of charges to


Figure 17 Display of tauFish array data. The first image is the array's response to the index finger touching the tabletop. The second image is the array's response to a palm touching the tabletop.



the induced surface charge on the plane, the general case of the inverse problem is not so easily solved. Instead, we use *a priori* knowledge of the target activation regions or hot spots that represent hyperlinks in the presented content, and use the forward model from above to predict an expected induced charge distribution for each hot spot. Each page of content specifies a list of hot spots by their positions on the image plane. When a page is presented, the content browser transfers that page's hot spot list to the *hot spot interpreter*, a separate program that receives tauFish array data and reports activation events back to the browser. The hot spot interpreter can also be used to display the sensor data in real time, as shown in Figure 17. The intensity of the large red squares in these images corresponds to variations in capacitance, while the small purple or green dots indicate hot spots (where green denotes a triggered hot spot).

To determine hot spot activation, the observed sensor data are normalized and linearly projected onto the charge distributions calculated for each hot spot by the forward model. This may be thought of as a distance metric that ignores the total charge. If the best fit exceeds a certain threshold (or rather, the angle between the observed and predicted charge vectors is less than a certain value) then the hot spot is activated and an activation event is sent to the browser.


Figure 18 Lazy Susan aperture ring and sensor board



Lazy Susan

The sensing system in the lazy Susan (shown in schematic form in Figure 19) consists of three parts: a sensor board, a reflective aperture ring, and an interface board. The sensor board, shown in Figure 18, includes two Omron phototransistor/LED pairs and a pair of AD790 comparators. The aperture ring, also shown in Figure 18, is a perforated acrylic ring with a reflective Mylar** coating, mounted on the

Figure 19 Schematic of lazy Susan installation



underside of the lazy Susan. The sensor board is stationary and its spacing below the aperture ring can be adjusted with pieces of card stock.

A schematic diagram of the sensing system operation is shown in Figure 19. The sensing system operates on the principle of quadrature encoding, using the pattern of alternating reflective strips and nonreflective apertures caused by the rotation of the aperture ring. The two phototransistors on the sensor board are offset in space by half the distance between a reflective strip and a nonreflective aperture.

The interface board consists of a PIC16F84 microprocessor, a MAX202 serial line driver, and a voltage regulator for the power supplied to the sensor board. The high/low outputs of the detectors are received as digital inputs by the interface board, and a simple state-machine-driven program in the microcontroller determines the rate and the direction of rotation of the lazy Susan.

The output of the microcontroller is a string of “R” and “L” ASCII characters. An “R” is generated when the lazy Susan rotates through an angular distance of one aperture unit to the right, and an “L” when the rotation is to the left. These data are transmitted to a Filament card and onto the network. A computer in the back room receives the data packets and rotates the coordinates of the center image accordingly.

In the course of testing the setup, it was discovered that the redraw time for the center image was slow enough for it to visibly lag behind the rotation of the

lazy Susan, catching up only after the lazy Susan had slowed significantly. Rather than trying to redisplay the image for every aperture unit of motion, the location of the lazy Susan at the end of a redraw was identified and used as target for the next redraw. As a result, when the lazy Susan spun faster, the rotation of the image was closer to the actual motion of the lazy Susan and more pleasing to the eye.

Network infrastructure

The seventeen sensors used in the dining table installation (eight tauFish arrays, eight tag readers, and the lazy Susan sensor board) were connected to each other and the controlling computers via an Ethernet data network. A single Ethernet cable brought network service to the table, and each sensor was connected to the network by means of a custom-made embedded network adapter, which we call the Filament. The Filament represents the minimum level of resources required to connect a device to the network.

Despite the increasing ubiquity of data networks, most devices in the computing industry are still designed for use with a single computer. Unfortunately, this model meshes poorly with the mounting demand for delocalized information access and distributed processing. The Filament was designed to address those needs. It consists primarily of an Ethernet controller (the CrystalLAN CS8900A) and a microcontroller (the PIC16F876) that contains a minimal implementation of UDP/IP, the User Datagram Protocol over the Internet Protocol. When connected to an arbitrary device (via a standard serial port) the Filament allows that device to communicate with any other device or computer on the network.


This combination (Ethernet and IP) was chosen in large part because it has become the *de facto* standard for building-level networking. Computers can communicate with Filament-enabled devices as easily as they can communicate with other computers on the network. No unusual protocols or expensive and cumbersome translation or gateway systems need to be used, as is the case with existing commercial offerings such as those of emWare. Furthermore, because this system communicates directly using IP, it is capable of easily interoperating with a wide range of networks and can benefit from the over 30 years of experience the networking industry has had in using that protocol.

For all those reasons, IP has recently gained acceptance as the appropriate protocol for embedded systems, and a number of engineers and companies have introduced IP-compatible embedded products. These existing devices fall broadly into two classes. Representatives of the first class simply use IP as a packet-data protocol on top of a traditional serial link, using SLIP (Serial Line Interface Protocol) or PPP (Point-to-Point Protocol). However, a serial multiplexing system is still needed to connect a large number of devices. The second class of systems does use Ethernet or another network technology as the physical transport layer. Unfortunately, the existing devices tend to be bulky or overcomplicated, rendering them inappropriate for many simple embedded applications. For this installation, what was needed was an extremely simple Ethernet/IP implementation that could be used to connect the various devices.

The particular device we used, the Filament F3r3, was designed to perform this function in an inexpensive, highly scalable manner. The device itself has a cost of between \$10 and \$20 per unit in large quantities, but the F3r3 is really an eight-wire device (power, ground, two serial pins, and four Ethernet pins) and so the limit of simplification for the F3r3 is a single eight-pin chip capable of connecting any serial device directly to the network. The device is shown in Figure 20.

Using an Ethernet/IP network to connect the various devices offered a wealth of advantages over more traditional serial multiplexing systems. Cabling was trivial, and in particular there were none of the distance constraints often associated with serial multiplexers. Also, by plugging a laptop computer into the Ethernet hub in the table, network debugging tools could monitor the flow of information in real time and in great detail, without interfering with the normal operation of any of the devices or the overall system. The sensors, data processing programs, and display programs were also engineered to communicate using UDP as a form of network-compatible interprocess communication. The end result was perhaps the most important advantage of using an Ethernet-based architecture: everything was completely delocalized. Individual software subsystems could be moved from one computer to another to take advantage of different debugging environments or to rule out computer-specific troubles and other software interactions.

Figure 20 Photograph of the Filament F3r3



System software

In designing the interactive dining table we wanted to allow visitors seamless access to the multimedia content within the exhibit environment. We wanted to create a space where people without computer skills could interact with the museum exhibit as fluidly and easily as they would with a piece of furniture or a book.

In implementing the system software for the dining table we had three major constraints. First, time to completion for this part of the project was a short couple of months. Second, since our group knew how to build hardware and software, but knew little about authoring content, we needed to create a software environment that enabled the design company Near-life to focus on content creation. Third, our system software needed to interact with our sensing devices over Ethernet in real time, by human perception standards.

We chose to use an Internet browser as our software environment because of the built-in support for a variety of interactive content and the ability to rapidly prototype using the HyperText Markup Language (HTML). There were two problems with using HTML alone. First, straight HTML does not give designers enough control over layout and content. Second, and a much larger problem, Web browsers are based on a client/server model. In this model a browser makes an HTTP request to a server and in a time usually long enough to be perceived by the user, the server responds. This information exchange had too much overhead for our system to work in real time. In addition, there was no built-in means for browsers to respond to User Datagram Protocol (UDP) packets. Fortunately, modern browsers interpret a number of other languages that enrich the experience of our modern Internet. In the end we

were able to make images, videos, and text dynamically change in response to physical icons and touch.

Figure 11 shows the overall architecture. The first task was to enable the browser to receive content from our local input devices on the Ethernet. To achieve this, we wrote a small Java applet that communicates with our devices using UDP. Modern browsers make it trivial to embed Java applets in a Web page, but the events within an applet and the events within the Web page itself are usually kept separate. JavaScript** is a scripting language built into modern browsers that allowed us to make this connection. To do this, the Java applet was encapsulated in a JavaBeans** component. This encapsulation creates a standard interface (a bean) that allows two software objects to communicate without the need to know the internals of the other object. This solved our real-time user input problem. This could be an interesting pathway for future projects to receive real-time sensor input to a Web page without needing remote servers. This could also be a means for people with browsers to send information to Web servers via nonstandard sensor or input devices.

With user input worked out, we then turned to the content design aspect. Dynamic HTML (DHTML) and Cascading Style Sheets (CSS) are the industry's response to creating dynamic content with full control of the layout. One can specify the location and visual attributes of objects precisely on the screen. The other important piece is that these attributes can change—by using JavaScript, one can quickly script a complex series of events given a user's input. For example, when a user touches a hot spot, visual feedback is given by brightening the color of hot spots through a transparency effect of one image over another. We were able to display text, images, and videos dynamically to delve into more information about a building. We chose Microsoft Internet Explorer** 5.0 running on Microsoft Windows 98** because its vast, albeit nonstandard, set of multimedia effects allowed us to create a rich environment that could not have been done with any other browser.

JavaScript, as a scripting language, suffers in performance relative to compiled code. It was paramount to the overall control of the software project, but there were certain pieces such as animations, which needed to be precompiled to maximize performance. We turned to Macromedia Flash** to solve this problem. Macromedia has made a set of intuitive tools

for creating quick, beautiful animations that are Web-ready.

The software environment we created spanned many levels: there is microcode on RISC (reduced instruction-set computer) microcontrollers that communicate over Ethernet to a Java applet encapsulated as a JavaBeans component that communicates to JavaScript, which controls the HTML and DHTML of a Web page that then renders text, images, and video. This seems complex, but the beauty of having so many levels of description was that the end product was exactly as intended: seamless.

Projection display hardware

The ten projectors used in the welcome mat and the dining table installations were NEC MultiSync** MT1035 projectors, with light output of 1300 ANSI (American National Standards Institute) lumens. For the dining table, the projectors were mounted in a circular drop ceiling directly above the place settings, with an additional projector over the lazy Susan. They were all side-mounted with lenses facing downwards. Special lenses were used to focus the display area to the size of each place setting. The apertures in the ceiling for the projectors were designed to create the impression of spotlights from the ceiling. Cabling (the VGA, or video graphics adapter, cables for the display and the remote control cables for the projectors) was run through the roof of the gallery to the back room used to store the computers running the exhibit. Figure 21 shows a section of the table and projector housing with the display paths for the projectors in red.

As mentioned in the introduction to this paper, the forward projection display was chosen over flat-panel touch screens or rear projection screens. In the kiosk environment, where the panel is vertical and users are viewing it head-on, traditional LCD technology makes a great deal of sense. But to enable the user to drop things, cut things, and generally use the table as a table, the forward projection method is far superior to any conventional touch-screen technology. Indeed, an embedded flat panel touch screen would require a thick protective layer of glass, which would make viewing the screen at an angle difficult.

As for the rear projection method, the projection hardware would occupy space underneath the table. This might impair the look or the functionality of the table. The use of forward projection did not in any way change the feel of the table or the ability

to sit comfortably at it. There is a definite psychological difference between an LCD screen that is indestructible because it has a thick slab of glass in front of it, and a table that is indestructible because it is a simple, solid piece of furniture. There was some concern that display occlusion (which would occur when a user places a hand into the path of the projected image onto a place setting) would have a negative impact, but the area occluded by the hand was not visible to the user anyway. The projectors themselves were well concealed, and the overall effect was closer to a magical glowing panel than a traditional projected image.


Results and conclusions

The technology development cycle for this project was significantly different from any other internal project of similar scale. One of the major observations from the time spent setting up the installation in the gallery was the difficulty of doing significant hardware debugging on site (see Figure 22): the level of developmental support that can be obtained in art museums is much less than in a conventional laboratory, and it took much longer than initially expected to complete the installation. Integrating the array boards and tag readers with the low-level networking hardware and software proceeded without a hitch. However, the development and integration of the software on the back-room computers was extremely difficult, despite the fact that the interpreters being used were performing relatively simple-to-describe computations. The lion's share of the development time was spent debugging the development environments rather than the actual algorithms. Coming out of these experiences was a stronger realization of the need for a system that would allow limited arbitrary computation to be performed without the overhead of a full operating system, and further development of this concept is under way.

Once the development was complete, it was most gratifying to observe visitors use the various exhibits. The reactions to the welcome mat ranged from the mildly amused to the completely visceral, and even with the fairly simple mapping described in the section on the welcome mat, the behavior was still novel enough to engage some of the younger users for some time. Most users seemed to grasp the concept that the space was, in some way, "aware" of their presence.

The concepts of attaching information to physical objects, navigating through this information with

Figure 21 Section of the table showing projector mountings and display paths



hand gestures, and sharing information using the lazy Susan were met with widespread understanding and acceptance by all the visitors to the gallery. The speed with which visitors acquired an understanding of how to use and navigate the interface told us we succeeded in making the interface simple and intuitive. There seemed to be many people who had vague ideas that there had to be a better way to get at information than a monitor and mouse, but no concept of what that could be; seeing and using the dining table really brought those ideas into focus for them. Perhaps the best reaction was that of an elderly museum benefactor who pulled Neil Gershenfeld aside and commented that she loved the installation and the fact that there were no computers anywhere near the table, because she hated computers but she loved being able to look at all the information. With all the fast microprocessors in the dining table installation she was probably closer to more computing power than she had ever been in her life, but the interface was so simple that it did not matter, and that more than anything else was the desired goal.

Acknowledgments

Flavia Sparacino built the initial prototype for the interactive table. The projectors used in the project were supplied by NEC Technologies. Furniture Co., New York, supplied the table. At MoMA, Andrew Davies, project manager and interactive environment

Figure 22 Photograph of gallery during installation



producer, and Paul Niebuhr, the director of information systems, gave invaluable assistance during the installation. Greg Van Alstyne, the art director, directed the team from NearLife, Inc., that designed the final interface and handled the content editing and formatting for the project. The images of the gallery floor plan, the projector schematic, and the place setting and lazy Susan photographs are used courtesy of Andrew Davies, Greg Van Alstyne, Paul Niebuhr, and MoMA. The installation would not have been possible without the work of Kelly Dobson and John Paul Strachan. The technology development was conducted in the Physics and Media Group and supported by the Things That Think consortium at the MIT Media Laboratory.

**Trademark or registered trademark of E. I. du Pont de Nemours and Company, Silicon Graphics, Inc., Sun Microsystems, Inc., Microsoft Corporation, Macromedia Inc., or NEC Home Electronics (U.S.A.) Inc.

Cited references and note

1. S. Wilkinson, "Phantom of the Brain Opera," *Electronic Musician* **13**, No. 1 (1997).
2. J. Paradiso, "The Brain Opera Technology: New Instruments and Gestural Sensors for Musical Interaction and Performance," *Journal of New Music Research* **28**, No. 2, 130–149 (1999).
3. C. R. Wren, F. Sparacino, A. J. Azarbayejani, T. J. Darrell, T. E. Starner, A. Kotani, C. M. Chao, M. Hlavac, K. B. Russell, and A. P. Pentland, "Perceptive Spaces for Performance and Entertainment: Untethered Interaction Using Computer Vision and Audition," *Applied Artificial Intelligence* **11**, No. 4, 267–284 (1997).
4. H. Ishii and B. Ullmer, "Tangible Bits: Towards Seamless Interfaces Between People, Bits and Atoms," *Proceedings of*

Conference on Human Factors in Computing Systems (CHI '97), Atlanta, March 1997, ACM (1997), pp. 234–241.

5. F. Sparacino, K. Larson, R. MacNeil, G. Davenport, and A. Pentland, "Technologies and Methods for Interactive Exhibit Design: From Wireless Object and Body Tracking to Wearable Computers," *International Cultural Heritage Informatics Meeting (ICHIM 99)*, Washington, DC, September 1999, Archive and Museum Informatics, Pittsburgh, PA (1999).
6. See the Boundary Functions installation by Scott S. Snibbe, <http://www.snibbe.com/scott/bf/index.htm>.
7. R. R. Fletcher, O. Omojola, E. Boyden, and N. Gershenfeld, "Reconfigurable Agile Tag Reader Technologies with Combined EAS and RFID Capability," *2nd Workshop on Automatic Identification Advanced Technologies (AutoID '99)*, Summit, NJ, October 1999, IEEE, New York (1999), pp. 65–69.
8. J. R. Smith, "Field Mice: Extracting Hand Geometry from Electric Field Measurements," *IBM Systems Journal* **35**, Nos. 3&4, 587–608 (1996).
9. L. S. Theremin, "Method of an Apparatus for the Generation of Sounds," U.S. Patent No. 1,661,058 (February 28, 1928).
10. T. G. Zimmerman, "Applying Electric Field Sensing to Human-Computer Interfaces," *Proceedings of Conference on Human Factors in Computing Systems (CHI '95)*, Denver, 1995, ACM (1995), pp. 280–287.
11. N. Gershenfeld, "Method and Apparatus for Electromagnetic Non-Contact Position Measurement with Respect to One or More Axes," U.S. Patent No. 5,247,261 (September 21, 1993).
12. J. D. Jackson, *Classical Electrodynamics*, John Wiley & Sons, New York (1999).

Accepted for publication May 15, 2000.

Olufemi Omojola MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: kingpin@media.mit.edu). Mr. Omojola is a research assistant in the Physics and Media Group at the MIT Media Laboratory. He received the S.B.

degree in electrical engineering and computer science from MIT in 2000 and is currently working toward an M.S. degree. His interests are in flexible reconfigurable hardware architectures for signal processing, and applications for human-computer interfaces.

E. Rehmi Post *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: rehmi@media.mit.edu).* Mr. Post is presently a doctoral candidate in the Physics and Media Group of the MIT Media Laboratory. He received an M.Sc. in media arts and sciences from MIT in 1998, and a B.Sc. in physics from the University of Massachusetts, Amherst.

Matthew D. Hancher *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: mdh@mit.edu).* Mr. Hancher is a research assistant in the Physics and Media Group of the MIT Media Laboratory and a candidate for an M.S. degree. Mr. Hancher's research focuses on embedded networking and autonomous robotics.

Yael Maguire *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: yael@media.mit.edu).* Mr. Maguire is a Ph.D. candidate at the MIT Media Laboratory. He is an IBM Student Fellow holding a master's degree in media arts and sciences from MIT. Mr. Maguire received a degree in engineering physics with Highest Honors from Queen's University in Canada. He is currently working on building a tabletop device for quantum computing, which holds the promise of computing faster than traditional, semiconductor-based computers. He currently holds an NSERC fellowship from Canada and is generally interested in the fundamental ties between information processing and physics.

Ravikanth Pappu *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: pappu@media.mit.edu).* Mr. Pappu is an IBM Fellow in the Physics and Media Group at the MIT Media Laboratory. His doctoral work focuses on designing and implementing inexpensive systems for cryptographic authentication schemes. Prior to this, he worked on algorithms for computing holograms for the MIT second-generation holographic video system and co-created the first dynamic holographic video system with haptic interaction. Mr. Pappu holds a bachelor's degree in electronics and communication engineering from Osmania University, India. He received a master's degree in electrical engineering from Villanova University in 1993, and a master of science in media arts and sciences from MIT in 1995.

Bernd Schoner *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: schoner@media.mit.edu).* Mr. Schoner holds engineering degrees from the University of Technology in Aachen, Germany, and the École Centrale de Paris, France. He is currently a Ph.D. candidate at the MIT Media Laboratory, working on the sensing, modeling, and synthesis of acoustical instruments.

Peter R. Russo *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: prusso@media.mit.edu).* Mr. Russo is currently a sophomore at MIT studying physics and electrical engineering. He is an undergraduate researcher at the MIT Media Laboratory, where he works on conductive fabric touch sensors, and near-field electric and magnetic field sensors.

Richard Fletcher *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307 (electronic mail: fletcher@media.mit.edu).* Mr. Fletcher is a Ph.D. candidate in the Physics and Media Group at the MIT Media Laboratory, working in the field of electromagnetic tagging. His current thesis research is devoted to developing new ways of encoding information in smart material structures. Prior to coming to the Media Lab, Mr. Fletcher worked for five years at the U.S. Air Force Materials Laboratory designing passive microwave structures. He holds undergraduate degrees from MIT in physics and electrical engineering.

Neil Gershenfeld *MIT Media Laboratory, 20 Ames Street, Cambridge, Massachusetts 02139-4307.* Dr. Gershenfeld leads the Physics and Media Group at the MIT Media Laboratory and directs the Things That Think research consortium. His laboratory investigates the relationship between the content of information and its physical representation, from developing molecular quantum computers, to smart furniture, to virtuosic musical instruments. Author of the books *When Things Start to Think*, *The Nature of Mathematical Modeling*, and *The Physics of Information Technology*, Dr. Gershenfeld has a B.A. degree in physics with High Honors from Swarthmore College, a Ph.D. from Cornell University, was a Junior Fellow of the Harvard University Society of Fellows, and a member of the research staff at AT&T Bell Laboratories.