

DiamondTouch: A Multi-User Touch Technology

Paul Dietz and Darren Leigh
Mitsubishi Electric Research Laboratories
201 Broadway
Cambridge, MA 02139 USA
+1-617-621-7500
{dietz,leigh}@merl.com

ABSTRACT

A technique for creating a touch-sensitive input device is proposed which allows multiple, simultaneous users to interact in an intuitive fashion. Touch location information is determined independently for each user, allowing each touch on a common surface to be associated with a particular user. The surface generates location dependent, modulated electric fields which are capacitively coupled through the users to receivers installed in the work environment. We describe the design of these systems and their applications. Finally, we present results we have obtained with a small prototype device.

KEYWORDS: DiamondTouch, multi-user, touch, collaborative input, single display groupware

INTRODUCTION

DiamondTouch is a multi-user touch technology for tabletop front-projected displays. It enables several different people to use the same touch-surface simultaneously without interfering with each other, or being affected by foreign objects. It also allows the computer to identify which person is touching where.

During the course of research on Human-Guided Simple Search [1] some of our colleagues have constructed a collaborative workspace in which multiple users work on the same data set. The environment consists of a ceiling-mounted video projector displaying onto a white table around which the users sit. A single wireless mouse is passed around as different users take the initiative. Our colleagues proposed that the collaboration would be improved if the users could independently and simultaneously interact with the table, and considered using multiple mice.

The use of multiple mice in a collaborative environment is particularly problematic. It can be challenging for users to keep track of one pointer on a large surface with lots of activ-



Figure 1: The collaborative work environment for Human-Guided Simple Search.

ity. Keeping track of many mice is nearly impossible. This leaves users physically pointing at their virtual pointers to tell other users where they are. Also, relying on a separate physical device keeps us from utilizing the natural human tendencies of reaching, touching and grasping.¹

Using a large touch-screen as the table surface would seem to be an answer, but existing touch technologies were inadequate. Most allow only a single touch and do not identify users. While schemes have been developed where users take turns [3], we wanted the interaction to be simultaneous and spontaneous.

Unlike electronic whiteboards or other vertical touch systems, the tabletop nature of our display creates a problem: people tend to put things on tables. With a pressure-sensitive surface, foreign objects create spurious touch-points causing single touch systems to malfunction.

Optimally, we would like a multi-user touch surface to have the following characteristics:

¹Plus see the discussion in [2] for advantages of touch tablets over mice.




Figure 2: DiamondTouch works by transmitting signals through antennas in the table. These signals are capacitively coupled through the users and chairs to receivers, which identify the parts of the table each user is touching. This information can then be used by a computer in the same way as mouse or tablet data.

1. Multipoint: Detects multiple, simultaneous touches.
2. Identifying: Detects which user is touching each point.
3. Debris Tolerant: Objects left on the surface do not interfere with normal operation.
4. Durable: Able to withstand normal use without frequent repair or re-calibration.
5. Unencumbering: No additional devices should be required for use – e.g. no special stylus, body transmitters, etc.
6. Inexpensive to manufacture.

The DiamondTouch technology meets all of these requirements. In the following sections, we describe its operating principles, the sorts of interactions that are possible, and the results of our experience with a small prototype device. We also present some ideas for future work and applications.

DiamondTouch

DiamondTouch works by transmitting a different electrical signal to each part of the table surface that we wish to uniquely identify. When a user touches the table, signals are capacitively coupled from directly beneath the touch point, through the user, and into a receiver unit associated with that user. The receiver can then determine which parts of the table surface the user is touching.




Figure 3: A set of antennas is embedded in the table-top. The antennas are insulated from each other and from the users.

The table surface is constructed with a set of embedded antennas which can be of arbitrary shape and size. The antennas are thin pieces of an electrically conductive material which are insulated from each other. Since the coupling of signals to the users is done capacitively, the antennas are also insulated from the users, and the entire table surface can be covered by a layer of insulating, protective material as shown in Figure 3. Each antenna extends over a single area of the table to be unambiguously identified: the system cannot tell where on the antenna a user touches, just that the user touches that antenna. A transmitter unit drives each antenna with its own signal that can be distinguished from the signals of the other antennas. Users are capacitively coupled to their receivers through their chairs, and the receivers are connected back to the transmitter through a shared electrical ground reference.

When a user touches the table, a capacitively coupled circuit is completed. The circuit runs from the transmitter, through the touch point on the table surface, through the user to the user's receiver and back to the transmitter.

With proper design, capacitive coupling [4] through the human body [5] can be quite reliable. The first consideration is that we wish to operate via “near field” (i.e. capacitive) coupling. By limiting the transmitting frequencies so that the dimensions of the entire system are very short compared to a wavelength, very little energy is radiated. This reduces problems with radio frequency interference and with unwanted coupling between the antennas. For reasonably sized tables, frequencies should be in the sub-MHz range.²

The system can be understood with the aid of a simplified equivalent circuit as shown in Figure 4. C_{table} represents the capacitance between the user's finger and a transmitting antenna in the table. C_{chair} represents the capacitance between the user and a conducting chair. The coupling capacitance is

²At a frequency of 1 MHz, the wavelength is about 300 meters.




Figure 4: The equivalent circuit for the DiamondTouch system.

the series combination of these two capacitances:

$$C_{coupling} = \frac{C_{table} \times C_{chair}}{C_{table} + C_{chair}}$$

Since the coupling area of a finger is very small compared to the entire body in a conducting chair, C_{table} tends to be very small compared to C_{chair} . Thus, the equation reduces to $C_{coupling} \approx C_{table}$. This means that the precise capacitive coupling via the chair is inconsequential, as long as it is large enough. If it were desirable to have the users stand on conductive floor plates instead of sitting in conductive chairs, the coupling area would be substantially smaller, but still very large compared to a finger. Experience has shown that even thick-soled shoes do not present a problem in this scenario.

When a user’s finger is far from the table, C_{table} is very small and little or no signal is coupled from the transmitter to the receiver. As the user’s finger approaches the table, C_{table} increases, the coupling to the receiver increases and so the received signal strength increases. The signal strength is also proportional to the area of the touch: using a thumb or the heel of a hand will produce a higher received signal strength than using a little finger.

DiamondTouch requires reasonable electrical isolation between the users. This constraint is violated if two or more users (or their chairs) are touching, or are in very close physical proximity. In this regard, social norms of “personal space” have been sufficient to keep the inter-user coupling acceptably small. However, this behavior can be explicitly exploited. By touching another user (or their chair), the touches of either user are interpreted as touches for both users. Typically, the coupling “through” a second user is somewhat weaker, and so it is usually possible to determine the “primary” user versus “shared” users. This provides a simple and intuitive mechanism for users to jointly indicate a selection.

UNIQUE SIGNALS

Because a user may touch several antennas at once, it is important that the receiver be able to distinguish between and identify any mix of incoming signals. We can do this if the signals are “separable”, or in signal processing terms “mutually orthogonal”.

There are many ways of generating such signals.³ For example, each antenna could be driven with a sinusoid of a different frequency. A receiver that is coupled to several antennas could determine which ones they are by examining the frequency spectrum of the received signal. Unfortunately, generating the numerous frequencies required for a large array is complicated and relatively expensive, so we rejected simple “frequency-division multiplexing” in our prototype.

Time-division multiplexing is another option. In this case, each antenna is separately driven in turn by a given signal while the other antennas are not. The timing of the received signals is used to determine which antennas are coupled to the receiver. While this system is very simple to implement, it may not be appropriate for larger arrays. The problem is caused by the interplay of various constraints. To provide good response time, the entire array must be scanned 10 to 100 times per second. However, as noted previously, practical modulating frequencies are limited to the sub-MHz range. This leaves very few modulation cycles per antenna, making receiver design difficult, especially in the face of other interfering noise sources. There are some clever ways of reducing the scan time [6] that help to extend the practicality of time-division multiplexing schemes, but these are beyond the scope of this paper.

Another way to construct a set of orthogonal signals is by code-division multiplexing, which is a spread spectrum technique. In fact, this turns out to be a particularly elegant approach for large arrays because very simple hardware operations (shifts, XORs, etc.) can be used to generate the large number of spreading codes. The simple hardware can even be cascaded, so that smaller touch devices can be tiled to make much larger ones. The spread spectrum approach will actually provide a significant gain in signal-to-noise ratio for large arrays.

ANTENNA PATTERNS

As we stated before, the antennas embedded in the tabletop can be of arbitrary shape and size. A designer may choose to implement just a few large “buttons” or a much more complicated array. Of course, a general, configurable solution is more desirable than a particular one that is designed into the hardware.

The most general solution is a “full matrix” pattern, in which a very large number of antennas are arranged in a rectan-

³A concise explanation of these various types of multiplexing can be found in [7]. For more general information on orthogonal signals, and spread spectrum information, see [8].




Figure 5: The row-column antenna pattern that our prototype uses. Each row or column is composed of diamond shapes connected in one direction and isolated in the other. This allows the maximum surface area for both layers without the upper one shielding too much of the lower one.

gular grid. Such a matrix of individually driven antenna “pixels” allows an unambiguous determination of multiple touch locations, even for a single user. Unfortunately, this is also the most difficult pattern to manufacture due to the very large number of connections required and the correspondingly large amount of supporting circuitry. In reality, the full matrix pattern may be unnecessary for many applications. Although the simultaneous, multi-user feature is essential, it is usually sufficient for each user to indicate at most a single touch point or bounding box. This functionality can be obtained with a simple row/column pattern that drastically reduces the number of antennas.

The rows and columns will usually be on two different layers. Due to shielding effects, there is some subtlety to creating a good row/column antenna pattern. A simple rectangular pattern of columns on the upper layer will overlap and cover too much of the equivalent set of rows on the lower layer. This will decrease the amount of area through which the rows can capacitively couple signals, weakening their sensitivity. A good antenna pattern will minimize the area in which the rows and columns overlap, while maximizing their total areas. We have found the connected diamond pattern shown in Figure 5 to be a good choice. This pattern has the interesting property that the row conductors are identical to the column conductors, rotated ninety degrees. In our prototypes, this allowed us to create a single conductor pattern and use it for both rows and columns, saving manufacturing costs.

In use, a touch will most likely span multiple rows and columns with different degrees of coupling. The received signal strengths can be used to estimate a centroid for the touch, obtaining positioning finer than the row and column spacing. However, an alternative way of using this information is to present a bounding box for the touch event, defined by the outermost rows and columns that have significant cou-

pling. This leads to an interesting use of the device – a single user might touch two points to define a bounding box. This is a very natural way of selecting a rectangular area. In practice, we have found two modes of operation to be useful: when the coupled area is small, assume that the user is indicating a point. When it spans a larger area, assume that the user is trying to specify a bounding box. The end result is that even this simplified row/column design allows some multi-touch capability for single users.⁴

PROTOTYPE

In order to test these concepts, we have created a small DiamondTouch prototype, part of which is shown in Figure 6. The prototype has an active area of approximately 20 cen-




Figure 6: Part of the prototype's antenna array. Compare with Figure 5.

timeters by 20 centimeters containing 80 antennas arranged as 40 rows and 40 columns. The half-centimeter pitch was chosen so that a typical finger touch would span at least two rows and two columns. A 0.5 millimeter thick double-sided printed circuit board was designed to be either the row or the column array, depending upon the rotation. Since we would like the coupling to either rows or columns to be about the same, the boards were arranged with the antenna arrays sandwiched in the middle of the stacked row and column boards with a very thin insulator in-between. Thus the gap to the top surface was very similar, varying only by the thickness of the insulator.

The antenna arrays are driven by a transmitter board that appears in Figure 7. For the moment, we have implemented time-division multiplexing where each antenna, in turn, is driven with 10 cycles of a 100 kHz square wave. While this board is capable of driving the antennas with a 60 volt swing,

⁴Of course, it would be better if a row/column pattern could distinguish multiple touches from a single user. The problem is that, given two X and two Y coordinates, the system cannot tell if the intended touches are (X_1, Y_1) and (X_2, Y_2) or (X_1, Y_2) and (X_2, Y_1) . In most cases, timing information might be used to disambiguate the two cases. If you had (X_1, Y_1) and then (X_2, Y_2) appeared later, you could safely guess the pairings. A case where this method fails is if the two touch points come together and then separate.




Figure 7: The prototype's transmitter board based around a PIC microcontroller. It is small and uncomplicated.

we have found 5 volts to be quite sufficient. Using a higher voltage produces a better signal-to-noise ratio which can be useful in electrically noisy environments.

The receivers are attached (via shielded cables) to padded, folding metal chairs that serve as the user coupling devices. Just about any conductive chair can be used for this application as long as there is sufficient capacitive coupling between the occupant and the receiver cable. Non-conductive chairs will work if a conductive “cushion” (a layer of metal foil, perhaps padded for comfort) is used to couple the user to the receiver.

Figure 8 shows one of the prototype receivers. For maxi-




Figure 8: One of the prototype's receiver boards, based around a PIC microcontroller. One is needed for each user.

imum noise immunity, the receivers use synchronous demodulation, and thus require appropriate synchronization signals from the transmitter board. The receivers digitize the results

and send them in raw form to a PC via fast RS-232 serial connections. There is a separate receiver board for each user. The entire table is scanned 75 times per second and the PC receives a coupling value for each user for each row and each column. The 75 Hz update rate and negligible latency to the computer allow the prototype to be very responsive.

The table is considered to be “touched” when the received signal at an antenna is high enough. In theory, we could use a simple threshold to determine this. However, given component drift, user variations, and varying noise levels, we have found it more practical to adapt a threshold based on current estimates of minimum coupling and noise levels. This works satisfactorily, but more sophisticated methods may yield better results. A problem case arises when the rubber-footed chairs are dragged across the carpet. “Static electricity” causes large noise spikes that require better filtering.

The transmitter and receiver boards are based on PIC microcontrollers and other inexpensive, off-the-shelf electronic components. The most expensive parts we used were the printed circuit boards for the table itself, and these would be much cheaper in a mass produced product.

We have written test software that generates a bar graph display of the coupling level, for each row and column and each user, along the appropriate axes. Different colors are used for each user. The calculated touch points are graphically displayed: a cross-hair cursor is shown for small touch areas, and a bounding box is shown for larger ones.

RESULTS

The prototype DiamondTouch system works quite well. Figures 9 and 10 show the results for two people touching the




Figure 9: Two users are interacting with the table independently.

table at once. The functionality of each user is quite independent.

We have stated that DiamondTouch operation is largely unaffected by objects carelessly left on the surface. Figure 11




Figure 10: Here two users are creating bounding boxes. Note that the operations are independent and they do not interfere with each other.

shows that a conducting object left on the surface does not cause a problem. While normal objects do not affect the table, it is possible to design special ones that do. This could be very useful in applications that use tangible and graspable objects as part of their user interface.




Figure 11: An aluminum can has been placed on the table, but it does not affect operation, despite being electrically conductive.

Because the insulating layer between the antenna array and users does not require any special properties, it can be manufactured from a variety of materials to make the table robust under different environmental conditions. For example, glass or plastic could be used to make the table resistant to liquid and chemical spills. Our prototype was made from a fiberglass laminate called GML1000 [9], whose thermal properties allowed us to operate the table temporarily (and without damage) while it was covered with burning alcohol.

A Game

We have implemented a simple game to demonstrate some of the capabilities of DiamondTouch. Multi-player “Pop-a-Bubble” pits up to four players against each other in a real-time game of reflexes. Colored “bubbles” appear and disap-




Figure 12: Two users are playing the Pop-A-Bubble game. It is possible to implement this game only because the table can identify who is touching where.

pear on the table. The four players are each assigned a color (red, green, blue or yellow) and they get points for “popping” bubbles of their color. They lose points for trying to pop other players’ bubbles. Cyan bubbles are “wild” and any player will receive points for popping one. Magenta bubbles are “poison” and players who pop them lose points.

The game shows off DiamondTouch’s main features. The interaction is both spontaneous and simultaneous: the players reach out and act naturally, without having to worry about turn-taking or dealing with extra “gadgets”. Scoring is easily handled because the game can identify which player touched the table at what location and at what time. And it’s fun to play!

RELATED WORK

Many different technologies have been developed for sensing the position of objects in two (or three) dimensions. Here we discuss some of the technologies that we investigated in our work on DiamondTouch, as well as some that others suggested that we compare to it.

Resistive and capacitive touch screens have been sold for decades, but are confused by multiple touches. Those that are pressure sensitive cannot tolerate any debris objects left on them.

Ultrasonic systems [10] [11] have recently become popular for creating electronic whiteboards, but they require active pen holders, and do not generally allow multiple touches. Larger debris objects may cause “shadowing” which will degrade performance.

One system that can support more than one touch while identifying the tool used is the Wacom Intuos graphics tablet [12]. This has a feature called “Dual Track” that allows two tools (styluses or mice) to be used simultaneously. Unfortunately, the Intuos is smaller and more expensive than we wanted and is limited to two touches only.

Other multi-touch systems that cannot identify users include the FingerWorks FingerBoard [13] and Tactex [14] smart fabric technology. Although the FingerBoard is not shipping as of this date, it appears to use a two-dimensional array of capacitance sensors to obtain a 2-D “image” of the object placed on it. The Tactex technology senses pressure by changes in the optical properties of the material.

Some optical-based input systems have been designed which track hands or other objects around a 2-D area. HoloWall [15] uses a camera and infrared illumination to find objects near a glass wall. Strickon and Paradiso [16] have done something similar in free space using a scanning laser rangefinder. Both system can sense multiple “touches” but cannot easily distinguish between different users.

Near-field electric field (capacitive) sensing has been used for decades in applications as simple as touch switches. More elaborate forms of capacitive sensing were introduced to the user interface community in recent years. Zimmerman, et al described this technology in depth in [17] and introduced the Fish, a device used to measure the position of a hand in three space using electric fields. Related work can be found in [18] and [19]. These systems attempt to detect a hand or other object that is several centimeters from one of the electrodes, and use field strength to determine the position. DiamondTouch differs by requiring that the sensed object be very close (millimeters or less) to an electrode, but uses a large array of these to sense the position.

CONCLUSIONS AND FUTURE WORK

DiamondTouch multi-user touch technology achieves all of our stated goals. It detects simultaneous, multiple touches, identifying which user is touching each point. It is largely unaffected by objects left on the surface, and is extremely durable. There is no stylus to lose, and the entire system can be manufactured inexpensively.

Larger and Different Systems

We are interested in building units much larger than our prototype and see no barriers to doing so. Scaling the electronics should not present a problem. The prototype was small because it was made from printed circuit boards, and these are expensive to make in larger sizes and small quantities. Large antenna arrays could be manufactured very cheaply by etching sheets of metalized plastic. We believe that these could be so inexpensive that we can envision a day when most white-boards sold will include a DiamondTouch antenna array under the writing surface, ready to plug into a separately-sold electronics package if the owner wishes to have touch-input capability.

We have designed and are having manufactured a small run of larger prototype DiamondTouch units. These will be made by silk-screening conductive ink onto flexible plastic, and will measure 80 cm by 48 cm with the same 0.5 cm row and column pitch as the original prototype. They will connect

to the host computer via a USB interface instead of several serial ports.

While we have described DiamondTouch’s use in a front-projected format, the technology is certainly not limited to this. Because the signals are capacitively coupled, very little electric current flows through the antennas so these can be made of a relatively high-resistivity material. This means that transparent conductors such as indium tin oxide can be used, and that the technology will be useful for rear-projection applications as well. Our experiments with such materials are just beginning but show promise.

New Applications

The ability for simultaneous, identifying interaction opens some interesting possibilities. One of the more intriguing ideas is the ability to create virtual personal work areas. We originally envisioned DiamondTouch as a method to allow group collaboration on a common surface, but in practice, individuals will sometimes want to “break away” to briefly address some subset of the problem, and then wish to integrate their result into the whole. When these situations arise, DiamondTouch can create a virtual personal work area in front of the appropriate user that only responds to that user. The user can be manipulating objects in this space, without impacting the larger work effort of other users but for the loss of some table space. Since these virtual personal work areas are software defined, they can be created and destroyed on the fly, in any shape as desired.

The concept of virtual personal work areas can be extended to special “privileged objects”. A privileged object is an icon that allows only certain classes of users to perform certain operations with that object. For example, a plumber and an electrician may be viewing the same house plan, but only the plumber can modify the pipes and only the electrician can modify the wiring.

DiamondTouch’s capability of providing public and private spaces is the input dual of “Single Display Privacyware” [20], which does the same thing with displayed output. Meshing these two technologies could provide some interesting user interface abilities. We are doing research into public/private display systems here at MERL and plan to experiment with a combination of these and DiamondTouch.

Undoubtedly, more new and interesting applications will arise as we gain experience with more and larger DiamondTouch devices.

ACKNOWLEDGMENTS

The authors would like to thank Joe Marks and Kathy Ryall of MERL for both motivating this work, and encouraging us to publish in this forum.

REFERENCES

1. Anderson, Anderson, Lesh, Marks, Mirtich, Ratajczak and Ryall, *Human-Guided Simple Search* in Proceed-

- ings of AAAI 2000, July 30–August 3, 2000. Austin, TX, USA.
2. Buxton, W., Hill, R. and Rowley, P., *Issues and Techniques in Touch-Sensitive Tablet Input* in Proceedings of Siggraph '85 (July 22-26, San Francisco), ACM/Siggraph, 1985, pp. 215–224.
 3. Inkpen, K., McGrenere, J., Booth, K. S., and Klawe, M., *Turn-Taking Protocols for Mouse-Driven Collaborative Environments* in Proceedings of Graphics Interface '97, Kelowna, BC, May 1997, pp. 138–145.
 4. Baxter, L. *Capacitive Sensors Design and Applications*, IEEE Press, Piscataway, NJ, USA, 1997.
 5. Zimmerman, T. G., *Personal Area Networks: Near-field intrabody communication* in IBM Systems Journal, vol. 35, nos. 3&4, 1996, pp. 609–617.
 6. Lee, S. K., Buxton, W. and Smith, K. C., *A Multi-Touch Three Dimensional Touch-Sensitive Tablet* in Proceedings of CHI '85 (April 1985), ACM/SIGCHI, NY, 1985, pp. 21–25.
 7. Glisic, S. and Vucetic, B. *Spread Spectrum CDMA Systems for Wireless Communications*, Artech House, Norwood, MA, USA, 1997, pp. 340–341.
 8. Dixon, R. *Spread Spectrum Systems*, 2nd edition, Wiley-Interscience, NY, USA, 1984.
 9. GML1000 is a copper clad laminate designed for use in antennas and high frequency circuits requiring low loss. It is available from GIL Technologies of Collierville, TN USA. <http://www.gilam.com>
 10. The Mimio electronic white board. <http://www.mimio.com>
 11. The eBeam Presenter. <http://www.e-beam.com>
 12. The Wacom Intuos graphics tablet. <http://www.wacom.com>
 13. The FingerWorks FingerBoard. <http://www.fingerworks.com>
 14. Tactex Smart Fabric Technology. <http://www.tactex.com>
 15. Matsushita, N. and Rekimoto, J., *HoloWall: Designing a Finger, Hand, Body and Object Sensitive Wall* in Proceeding of the ACM UIST '97 Symposium on User Interface Software and Technology, 209-210.
 16. Strickon, J. and Paradiso, J., *Tracking Hands Above Large Interactive Surfaces with a Low-cost Scanning Laser Rangefinder* in Proceedings of CHI '98 Extended Abstracts, ACM Press, NY, April 1998, pp. 231–232.
 17. Zimmerman, T., Smith, J., Paradiso, J., Allport, D. and Gershenfeld, N. *Applying Electric Field Sensing to Human-Computer Interfaces* in Proceedings of CHI '95; ACM Conference on Human Factors in Computing Systems, pp. 280–287.
 18. Smith, J. R. *Field Mice: Extracting hand Geometry from electric field measurements* in IBM Systems Journal, vol. 35, nos. 3&4, 1996, pp. 587–608.
 19. Smith, J., White, T., Dodge, C., Paradiso, J., and Gershenfeld, N. *Electric Field Sensing for Graphical Interfaces* in IEEE Computer Graphics and Applications, vol. 18, no. 3, May-June 1998, pp. 54–60.
 20. Shoemaker, G. and Inkpen, K., *Single Display Privacyware* in Proceedings of CHI 2001, pages 522-529, ACM Press, 2001.