

Sensetable: A Wireless Object Tracking Platform for Tangible User Interfaces

James Patten, Hiroshi Ishii, Jim Hines* and Gian Pangaro

Tangible Media Group
MIT Media Lab

20 Ames Street, Cambridge, MA 02139
{jpatten, ishii, pangaro}@media.mit.edu

*System Dynamics Group
MIT Sloan School of Management
50 Memorial Drive, Cambridge, MA 02139
jhines@mit.edu

ABSTRACT

In this paper we present a system that electromagnetically tracks the positions and orientations of multiple wireless objects on a tabletop display surface. The system offers two types of improvements over existing tracking approaches such as computer vision. First, the system tracks objects quickly and accurately without susceptibility to occlusion or changes in lighting conditions. Second, the tracked objects have state that can be modified by attaching physical dials and modifiers. The system can detect these changes in real-time.

We present several new interaction techniques developed in the context of this system. Finally, we present two applications of the system: chemistry and system dynamics simulation.

Keywords

Tangible user interface, interactive surface, object tracking, two-handed manipulation, system dynamics, augmented reality

INTRODUCTION

A tabletop workspace with mechanisms for display and input is an appealing context for research in Tangible User Interfaces (TUIs) [5] for several reasons. Such a space provides ample room to organize objects spatially, which can be an important part of thinking about the problem solving process [6]. Users can collaborate easily around such a space to solve problems using both hands. Finally, physical objects in this type of environment can be more than just input devices: they can become embodiments of digital information.

The notion of an interactive display surface that is able to sense the positions of objects on top of it has been discussed in the HCI literature for many years [9,12,15]. However, the typical approaches to this object-tracking problem each have some limitations. Computer-vision-based approaches can




Figure 1: A system dynamics application running on top of Sensetable

have problems with robustness due to the need for controlled lighting conditions. [7] Tracking latency can also be an issue when objects are moved around in the sensing space. Magnetic tracker based approaches, such as those made by Polhemus and Ascension require that wires be attached to the objects being tracked.

To support our research in interactive tabletop surfaces, we decided to develop a new platform, called Sensetable, which aimed to improve upon existing methods in two ways. First, we wanted the platform to provide accurate, low-latency wireless tracking of 6-10 objects on a flat surface. Second, we wanted to allow users to modify the tracked objects (using dials or "modifier" tokens), and to map these physical changes to changes in the application running on the platform. All of the technologies we investigated for this platform employed some form of electromagnetic sensing to determine the positions of objects.

After considering several alternatives, we decided to implement our first prototype by extending commercially available sensing tablet technology. Once our first prototype had been completed, we began developing applications and exploring interaction techniques using the system.

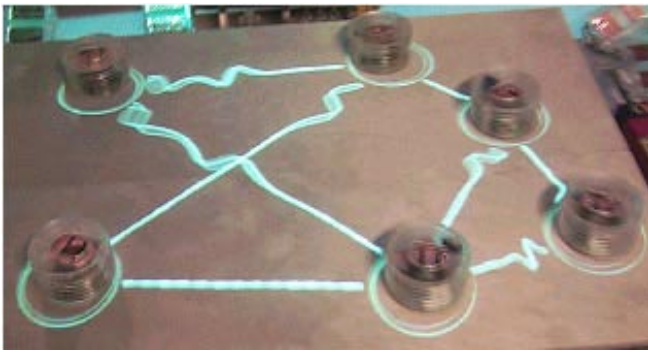


Figure 2: An interactive art piece made with Sensetable

In the next section we describe previous work related to the Sensetable project. In the third section, we describe the implementation of our first Sensetable prototype. We continue by presenting the interaction techniques we have developed using Sensetable. We then present the chemistry and system dynamics applications we have developed on top of Sensetable. Finally, we present our conclusions and plans for the second Sensetable prototype.

RELATED WORK

A series of research has influenced our work and helped us to identify the functional requirements for the Sensetable project. Wellner's Digital Desk [15] system introduced the concept of an interactive tabletop that was both physical and digital. Users interacted with digital content in the system by "touching" projected graphical representations on the desk. The system detected these touches using a camera and microphone. Interactions such as making calculations using a calculator projected on the desk were possible using this system. [16]

The Bricks project [2] pioneered the use of graspable handles for manipulating digital objects directly using two tethered Ascension Flock of Birds™ trackers. This system illustrated some of the powerful things one could do with a platform that tracked objects in real-time, and merged input and output into one physical space. However, this system was limited in that it only provided two physical objects for the user to manipulate, and these objects were connected to the computer with wires.

The metaDESK [12] system built on the ideas presented in the Bricks system by demonstrating the use of "phicons", or physical icons, in the context of an interactive surface. An infrared camera inside of a table tracked these phicons using simple computer vision techniques. Output from the system was projected into the same space using rear video projection.

The I/O bulb [13] system demonstrated the use of an interactive surface for urban planning. This system used an advanced vision technique that involved tracking objects based on unique patterns of colored dots. However, the limitations of computer vision in stability, robustness, and speed were still apparent in this application.

Several commercial platforms can provide robust tracking of physical objects. However, these devices are limited by the number of objects they can track at a time. [7] Usually, a state of the art product such as the Wacom Intuos™ [14] can track at most two input devices.

Zowie Entertainment, now part of the LEGO Group, released a breakthrough toy using multiple-object tracking technology at very low cost. Although their technology allows fast, high resolution tracking, the hardware only provides information about the identity and position of objects in the sensing space. However, we were interested in developing interaction techniques based on allowing the user to physically manipulate the objects using buttons, dials or by attaching modifiers. This led us to develop our own sensing platform.

IMPLEMENTATION

Our current implementation uses a pair of modified commercially available Wacom Intuos™ sensing tablets that are placed next to each other to form a 52cm x 77cm sensing surface. These tablets are an appealing technology to use for the Sensetable project because they can sense the positions of objects with roughly 1000 dpi resolution, and have very low latency compared to computer vision based approaches. As well, the mice used with these tablets each have a 32 bit serial number, which is useful for identifying mice when they move from one sensing surface to another. On the other hand, these tablets can only track two objects at a time. To circumvent this problem, we built the pucks to be tracked by augmenting the mice with a circuit to switch the sensing coils inside of the mouse on and off randomly. The random number generator we use ensures that each puck is turned on about one third of the time.

This duty cycling approach yields a tracking latency of less than a second. To reduce this latency, we added a circuit to sense when the puck is being touched. We built this using a capacitance sensor, which monitors an antenna wire wrapped once around the circumference of the puck. When the puck is touched, the microprocessor inside it detects a capacitance above a certain threshold, and it turns that puck on 100% of the time. In this way, the system can track objects that are being touched at a latency equal to that of an unmodified Wacom™ tablet. Objects that are not being touched are updated with a higher latency.

The pucks have two sockets inside of a crescent shaped recess on their top surfaces, shown in figure 3. These sockets connect to a 16 wire bus which




Figure 3: A Sensetable puck, with a socket for attaching a dial or modifier. A US quarter is shown for scale.



is used to communicate with dials and modifiers which can be placed on top of the pucks. Currently, four of these pins are used to communicate with the dials, four are used to communicate with the modifiers, and eight pins are reserved for later use. The modifiers have a unique digital ID, and bus connectors



on the top and bottom so they can be stacked. Currently the stacking order cannot be detected, but we are adding more intelligence to the modifiers to allow this. Because the dials use the same bus connector as the modifiers, they can be used while attached directly to a puck or while on top of a

Figure 4: The top and bottom of a dial that plugs into a Sensetable puck.

modifier or series of modifiers. A dial is shown in figure 4.

Two dual processor 866MHz Intel® Pentium® III Xeon™ computers are used to drive the system. One receives the data from the sensing surface and displays graphics onto the sensing surface in response. A second computer drives two vertical displays to the rear of the sensing surface, which can




Figure 5: System Architecture of Sensetable for System Dynamics Simulation

provide extra information relevant to the interaction happening on the table. In our system dynamics simulation application, this second machine also performs the actual simulation. In the future we plan to use both machines together to simulate larger system dynamics models in real-time. The system architecture is shown in figure 5.

Limitations

If more than two pucks on top of one of the sensing tablets are touched at the same time, tracking latency increases. In our testing with one and two users, this limitation was not a problem, because users did not typically move more than two objects at a time. However, we have not tested the interface in collaboration scenarios with larger groups of people. Our second generation prototype, which is briefly described in the continuing and future work section, is designed not to have this limitation.

Another limitation is a 3.5 cm gap in the sensing field due to interference between the two boards, where the two sensing elements touch each other. Our second prototype is designed to not have this problem.

INTERACTION TECHNIQUES

Once the underlying Sensetable hardware and software had been constructed, we began experimenting with some interaction techniques for use on top of the platform. The primary mode of interaction with the system is as follows: Graphical representations of digital information are projected onto the tabletop sensing surface. When the user moves a puck close to one of these graphical representations, the puck becomes “bound” to that item, and physical changes to the puck, such as plugging a modifier into the socket on top, cause corresponding changes in the bound information. For example, attaching a modifier to a puck when that puck is bound to a molecule in the chemistry application changes the charge of the molecule. Below we describe in more detail techniques for:

- Binding and unbinding pucks to and from digital information
- Manipulating digital information with pucks
- Visualizing complex information structures
- Sharing information between the tabletop sensing surface and a traditional display screen.

Binding and Unbinding

One of the challenges associated with tangible user interfaces is finding a way to interact with a large amount of information with a finite number of physical objects. One approach we have explored involves mechanisms for easily binding and unbinding physical objects to and from digital information. In applications where there is a low density of digital information that can be bound to pucks, one can attach a physical puck to a digital item just by moving the puck within a certain proximity of the object to be bound. This method is simple and works well, but as the number of digital items to which pucks can be bound increases, it can




Figure 6: Information being projected on top of the Sensetable pucks.

become difficult to select something to be bound without accidentally selecting something else first. We used two measures to address this issue. First, we dynamically adjusted the spacing of digital items displayed near an unbound puck to make it easier for the user to select a particular one. As well, we increased the amount of time necessary for the binding process to occur. If the user moves the puck toward an item on the table, the system displays graphical feedback that indicates the given item will be bound to the puck shortly if the puck is not moved. Before the binding process is complete the user can move the puck to cancel the operation.

To unbind a digital item from a puck, we use a shaking gesture. This approach is appealing because the visual effect seems to suggest that the physical forces being applied to the puck are breaking the bond between it and the digital item. However, when first interacting with the system, many users expected that they could unbind a puck from its associated digital information by picking the puck up off of the sensing surface and placing it down on top of some other digital item on the surface. While this is quite a reasonable expectation, our current hardware prototype has difficulty differentiating the act of lifting a puck off of the sensing surface from a puck switching itself on and off as part of the time-sharing scheme our prototype system uses for communication. Our second generation prototype of the system includes the ability to detect when objects have been lifted off of the sensing surface, so we intend to explore the "paperweight" metaphor offered by this technique in the future.

On one hand, we wanted to make it easy for users to attach and detach the pucks to and from digital items in the system. But in doing so, we did not want to complicate the metaphor that the puck was a physical embodiment of the data itself, and that adjustments to the physical objects would cause the data itself to change. Initially, we projected information about the corresponding digital content in front of the pucks on the table. This led one user to comment that pen or wand shaped objects might make more sense for manipulating the data, because they would not obscure so much of the information in front of them on the table. This comment

suggested to us that the user was not treating the puck as a physical embodiment of the digital data. To address this issue, we experimented with projecting information about the puck onto the puck itself, (as seen in figure 6) rather than in front of the puck. This change cleared up some confusion about what the pucks represented. We are interested in exploring other methods of displaying information about a digital items' state on the puck itself. One such approach involves a fold-down display, which is described in the continuing and future work section.

Dials and Modifiers

Exploring the use of dials and modifiers that could be plugged into tracked objects was one of our primary motivations in developing the Sensetable platform. There has been little exploration of this approach to physically modifying computational parameters. The AlgoBlock [10] system allowed children to adjust simple computer programs by rotating knobs on top of physical bricks. However, each of these dials was permanently attached to its corresponding brick, and could only modify one program parameter. We have explored the use of dials and modifiers on top of the pucks in a more dynamic role. In the chemistry application, modifiers can be placed on top of a puck to change the charge of the atom or molecule to which that puck is bound. In the system dynamics application, users can employ the dials on top of the pucks to adjust parameters in the simulation, as seen in figure 7. Users liked the idea of being able to physically manipulate simulation parameters in this manner. However, when using an early prototype of the dial functionality, users had two criticisms.

First, they wanted the information about the changes caused by manipulating the dials to be displayed on the sensing surface in addition to being displayed on a screen behind the surface. Second, they wanted graphical feedback near the dials themselves to provide a better sense of what the dial setting was at a particular point in time. After we made these changes, one could use the dials by focusing just on the table surface itself, rather than having to divide one's attention between the input on the sensing surface and the output of a rear display screen.




Figure 7: A simple graph of "potential customers" as a function of time. This graph is updated as the "unit sales" dial is adjusted.

Tangible Visualization Techniques

At times, users may wish to interact with more data at one time than can be legibly displayed on the sensing surface. We have explored several techniques to deal with this issue.

First, in the context of the system dynamics application, we developed a layout algorithm that adjusts the prominence with which objects are displayed on the table. Each digital item is assigned an importance according to a "scoring process" based on application specific criteria, and the model is searched for any items that overlap with each other. When a pair of overlapping items is found, the one with less importance is darkened to the point where it is still barely visible, and the graphical information associated with the other item is much easier to read. Being able to see the faint presence of an object provides the user with a cue that more information is available there, so he or she can focus on it using the techniques described below.

Indicating center of attention

While pucks are primarily used to move and manipulate digital items on the table, one can also use them to indicate interest in a particular region of the table. Using the scoring process described above, digital items near a puck are given higher display priority, and thus are made more visible. In a display space crowded with information, this yields a Fisheye [4] like effect where more detail is provided in the areas of user interest. The use of multiple pucks in the interface provides an easy way for the user to simultaneously indicate several areas of interest in the sensing space.

Semantic Zooming

Another technique Sensetable employs to give users intuitive controls over information display on the table is a semantic zooming [8] technique in which the distance between pucks on the table affects the level of detail used to show the information between the two pucks. The metaDESK [12] project demonstrated a technique related to this one for displaying maps. While the metaDESK example involves displaying information with a very literal interpretation of space, we have explored the use of this technique for physical navigation of digital data with no inherent spatial component. One example is the abstract graph structure used to represent simulations in system dynamics. Rather than changing the size of individual items being displayed on the table, we again use the scoring process described above to fade less important items into the background as two pucks come closer together. Nodes are faded into the background when they begin to interfere with the display of a more important node. With this approach, one can show different parts of the model in different levels of detail at the same time on the sensing surface. In contrast, related approaches such as the metaDESK display information at only one level of detail at a time.

Sharing information with an on-screen display

For some tasks, a user might want to share data between the tabletop interaction surface and an on-screen display in order to use tangible and WIMP interaction techniques together. The mediaBlocks system [11] provides a method for moving data between a physical container and an on-screen WIMP interface which involves placing a tagged wooden block in a socket on the side of the screen. More recent augmented surfaces work [9] adds the notion of a spatially continuous connection between the screens of portable computers and

nearby tabletops and wall surfaces. In this work, users can employ their mouse cursor to move objects to and from the physical world. Data can be associated with physical objects, but only with the mouse cursor. Building on the notion of a spatially continuous workspace, we have explored a method for this type of data sharing using Sensetable's physical, tracked objects as the means of transport and control. A flat panel display is aligned with the left side of the rear of the sensing surface, so that the display area of the flat panel begins where the display and sensing surface of the tabletop ends. Digital information that can be moved between the screen and tabletop space is displayed in boxes along the lower edge of the flat panel display, as seen in the top image of figure 8. The top portion of the rear display shows a higher-level view of the information for context. Directly below each of these boxes is a corresponding box projected on the sensing surface itself. When a puck is placed in one of these boxes, the contents of the corresponding on-screen window "slide" down onto the tabletop, highlighting the box with the puck inside it, as seen in the middle image of figure 8. Once the contents of the box have moved into this small portion of the tabletop space, the puck that is now bound to these contents can be used

to move and manipulate them on the table, as seen in the bottom image of figure 8. As the puck is moved, the contents expand to fill a larger part of the tabletop interaction space in a spring-like motion.

APPLICATIONS

We explored the interaction techniques described above in the context of two applications, described below. The




Figure 8: The process of moving information from the screen to the tabletop.

chemistry application is a proof-of-concept application to show some of the types of interactions one might have with Sensetable, while the system dynamics application has been developed in concert with system dynamics researchers at the MIT Sloan School of Management as a means to begin using Sensetable to address a real problem.

Chemistry

Figure 9 shows a tool built on top of Sensetable for teaching students about chemical reactions. The user can map the pucks to atoms or molecules, and then move these around in the workspace. When the atoms and/or molecules which are needed for a particular chemical reaction are brought into close physical proximity, the reaction occurs. The user can then manipulate the reaction products to use them in other reactions. The user can place modifiers on top of the pucks to change the electrical charge of the atom or molecule.

System Dynamics Simulation

Our second application of Sensetable is system dynamics [3] simulation. A picture of this application is shown in figure 1. System dynamics is a method for studying complex feedback systems in fields such as business and the social sciences. It involves the analysis of computer models to conduct "what if" analysis on a system. Using this analysis, one can develop an understanding of how the different parameters in a model affect each other. For example, in a model of the fox and rabbit populations in a forest, the size of each population would have an effect on the size of the other because of the predator/prey relationship between foxes and rabbits. One might hypothesize that an increase in the fox population would lead to a decrease in the rabbit population. One could then adjust the fox population in a simulation of the model to test this hypothesis.

A system dynamics model consists of a series of nodes (such as the rabbit and fox populations above) connected via a series of edges. In the Sensetable system dynamics application, the user can attach pucks to these nodes and use the dials on top of the pucks to adjust the corresponding simulation parameters. He or she can also move the pucks around to reorganize the display of the model. When parameters are changed, the system recomputes the simulation and displays the results on Sensetable itself and on a display to the left rear of the table.

During the design and development process of this application we asked people with varying levels of system dynamics experience to use the system. Their experience ranged from being a professor conducting research in system dynamics to having only a cursory knowledge of the field. Some of these tests were conducted with pairs of users with roughly equivalent experience working together, while others involved a single person using the interface while giving us verbal feedback about it. We conducted ten of these sessions that lasted from 30 to 60 minutes. Eight users participated in these tests, with several trying the interface at two or three stages of the development process.

Initially, users reported having difficulty analyzing models with more than 25 nodes in our system. They commented

that the automatic graph layout algorithms in our system removed some of the information that was encoded in the original layout of the system dynamics model. The person developing a system dynamics model usually carefully designs the layout of the nodes in the graph so that important causal loops in the model can be readily identified and studied. By adjusting the layout of our graphs on the sensing surface, we often removed some of this information. Because our layout algorithms were intended in part to deal with the problem of limited screen real estate, we began to investigate other methods of dealing with more complex graphs.




Figure 9: A chemistry application running on top of Sensetable.

Current on-screen system dynamics simulation packages address the problems stemming from limited screen real-estate by breaking up the model into a larger number of "views," each of which display a certain feature of the model. One can switch between these views using a menu. This approach to interacting with smaller portions of a system dynamics model at a time led us to explore the method of sharing data between the screen and tabletop portions of the interface described in the interaction techniques section. The use of this technique in the system dynamics application is shown in figure 8.

When a user first begins interacting with the system, he or she sees a complete version of the system dynamics model to be analyzed on the vertical display at the left rear of the interface. Directly below this graph is a display of several portions of the model that contribute significantly to the model's overall behavior. These portions have been selected in advance by the author of the model. The user can move one of these subgraphs from the vertical display to the tabletop sensing surface using the data sharing technique described in the "Interaction Techniques" section. As the puck is moved away from the screen, the subgraph expands to fill the TUI space, while one node in the subgraph stays attached to the puck.

Once one is through experimenting with a particular subgraph, he or she can return it to the on-screen space and choose another. When analyzing a system dynamics model with more than 25 nodes, users preferred moving parts of the

model between the GUI and TUI spaces to interacting with the entire model on the Sensetable at one time.

DISCUSSION

One of the things that surprised us while developing the system dynamics application was the different role of the layout of the model in on-screen space and in tabletop space. Traditionally in system dynamics models that are displayed in a WIMP interface, the author uses the spatial organization of the model to communicate information about important structures in the graph. For example, loops in the model sometimes cause patterns of oscillating behavior. One usually arranges the nodes in such loops so that it is very clear that the nodes form a loop. Thus in a WIMP context it can often hinder the process of analyzing the graph to adjust the layout of nodes from their original positions. However, there are also benefits one may achieve from adjusting the layout of the graph. Reorganizing the nodes may make a problem solving process easier by allowing the user to offload computation from his or her mind to the environment. Kirsh discusses organizing objects to help one think in [6]. For example, if one wanted to determine which among a group of nodes had an oscillatory effect on a parameter in the simulation, one might arrange the nodes to be tested in a line, and then adjust the dial on top of each corresponding puck in sequence and see what happened in response. As one tested each node, one might sort the nodes into two groups on the table depending on whether they contributed to the oscillation or not. At the end of this process, the arrangement of the nodes on the table would hold the answer to the original question, without any need on the part of the user to memorize or write anything down during the process.

By providing a static layout of the graph on the left rear display, and a dynamic, manipulable version on the tabletop, we believe Sensetable provides some of the better aspects of both interface styles for the problem domain of system dynamics. The screen provides a frame of reference for the analysis going on, and the tabletop allows the user to look at and manipulate a more manageable portion of the model during the process of analysis. In general, this seamless connection between the screen and tabletop allows one to move pieces of digital content to whichever space is best suited for the task at hand. While the current connection makes little use of the keyboard and mouse, we expect that as the Sensetable applications continue to mature, the keyboard and mouse will be quite useful in the graphical portion of the interface for tasks that are not done well in the tangible part of the interface.

Why Tangible?

We believe that Sensetable provides several benefits over traditional GUI-based techniques for analyzing system dynamics models. First, the ability to manipulate the physical dials and see real-time feedback about the change in simulation results was very exciting to our users. They enjoyed being able to use both hands at the same time to adjust two different parameters simultaneously. One commented that this approach helped him "develop an

intuition more quickly" about what the model would do. This interface often involves one less level of indirection between the human hand and the actual computational change taking place than does a mouse adjusting a slider [1]. We hypothesize that this ease of manipulating parameters may lead to more thorough analysis of models, which may in turn lead to a better understanding of the models' behavior. However, our experience with users thus far is suggestive, but not sufficient to evaluate this claim.

We believe the fact that Sensetable affords collaboration between users is also important. Instead of collaborating verbally while one person adjusts parameters with a keyboard and mouse, Sensetable allows different people to change parameters simultaneously. For example, this feature would be useful if managers of separate manufacturing plants owned by a company wanted to look at how various changes in their respective plants' production would affect the company as a whole. Each could control the parameters associated with his or her factory while observing the aggregate effect on the company.

CONCLUSIONS

We have presented Sensetable, a robust platform for tracking multiple objects wirelessly on a flat surface with high accuracy and low latency. The use of an electromagnetic sensing approach frees one from the problems typically associated with computer vision based approaches to object tracking. These include occlusion, susceptibility to variations in lighting, and higher latency. In addition to overcoming these issues, our sensing approach allows one to give the tracked objects state which can be physically manipulated with controls such as dials and modifiers.

Using this platform, we have explored some new interaction techniques including changing the distance between pucks to control the amount of information displayed between them, using pucks to indicate points of interest for a "fish-eye" like approach to displaying crowded graphs, and using gestures to bind and unbind physical pucks with digital content.

We have implemented an application on top of Sensetable to analyze system dynamics models. Users familiar with system dynamics tested the interface during the development process. For them the most valuable part of the interface was the ability to quickly adjust multiple parameters using the dials and see real-time feedback. While users also valued the ability to move the nodes around using the pucks, they found the association between the pucks and nodes unclear until we began projecting the names of the nodes onto the corresponding pucks themselves.

During the process of developing the system dynamics application, we developed a workspace that included a seamless interface between display screen and tabletop components. The rear display screen preserves the original structure of the system dynamics model and provides a reference frame for the investigations performed using the tangible component of the interface. On the other hand, the tangible component allows the user quickly to investigate the

effect of parameter changes on the model, and to reorganize portions of the model in support of this investigation.

CONTINUING AND FUTURE WORK

We are currently working on developing several aspects of this work more thoroughly. First, we are developing a new sensing board that uses a more scalable tracking technology than the one we currently employ. The new surface is constructed from 25 cm square sensing boards, which can be tiled to form sensing areas of varying size and shape. We anticipate that the primary limit on the number of objects which can be tracked at one time on the new board will be the number of objects which can physically fit on the surface. The tags for the new system are smaller (less than 4 cm on a side) as well. We plan to investigate interaction techniques that become feasible only with this larger number of tags.

In the area of interaction techniques, we plan to continue our investigation of how Sensetable can be combined with other approaches to the user interface, such as WIMP, speech based interfaces, etc. Our hope is that research in this direction will lead to interfaces which can solve problems that cannot be readily solved using just a single one of today's predominant approaches to the human-computer interface.

We are also excited about exploring interaction techniques that relate solely to tangible interfaces. One example is the stacking of modifiers on top of a puck. We anticipate using the stacking of modifiers to allow the user to perform "what if" analysis in a system dynamics simulation. For example, if a certain node represents the population of an animal in a forest, one modifier could mean that natural predators of the animal were removed, another could mean that the population was struck by some sort of disease, and so on. By composing these modifiers on top of the puck representing the animal population, users could experiment with a variety of scenarios within a simulation. We have completed the development of the hardware necessary to support this interaction, and we are currently working on completing the software so that we can begin to experiment with the technique.

Finally, we are interested in exploring various types of controls on the pucks themselves. One example is the use of a fold down display surface attached to the side of a puck. If the puck can sense when the display surface is folded open, the position and orientation of the puck on the sensing surface can be used to project extra information about the puck onto the surface. We plan to use this technique in the system dynamics application to display graphs of various simulation parameters as a function of time. A user will be able to open the display of a puck bound to a node in the simulation to see a plot of that node's behavior over time.

ACKNOWLEDGMENTS

We would like to thank the Intel Corporation and the Things that Think and Digital Life consortia of the MIT Media Lab for their support of this work. We would also like to thank Mary Murphy-Hoye of Intel, Professor Tom Malone,

Professor Rob Jacob, John Quimby, George Herman, Paulo Goncalves, Benjamin Koo, Chalermmon Lertpattarapong, the Tangible Media Group, Joe Panganiban, Gustavo Santos. We also thank the anonymous reviewers for their helpful comments.

REFERENCES

1. Fitzmaurice, G., Graspable User Interfaces. Ph.D. Thesis, University of Toronto, 1996.
2. Fitzmaurice, G. W., Ishii, H. and Buxton, W., "Bricks: Laying the Foundations for Graspable User Interfaces," in *Proceedings of the Conference on Human Factors in Computing Systems CHI'95*, ACM Press, pp. 442-449, 1995.
3. Forrester, Jay Wright, Industrial Dynamics. MIT Press. Cambridge, Mass, 1961.
4. Furnas, G., "Generalized Fisheye Views," in *Proceedings of Conference on Human Factors in Computing Systems CHI '86*, ACM Press, pp. 16-23, 1986.
5. Ishii, H. and Ullmer, B., "Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms," in *Proceedings of Conference on Human Factors in Computing Systems CHI '97*, ACM Press, pp. 234-241, 1997.
6. Kirsh, D., "The intelligent use of space," *Journal of Artificial Intelligence*, 73(1-2), 31-68, 1995.
7. Paradiso, J., et al., "Sensor Systems for Interactive Surfaces," *IBM Systems Journal*, 39(3-4), 892-914, 2000.
8. Perlin, K., and Fox, D., "Pad: An Alternative Approach to the Computer Interface," in *Proceedings of ACM SIGGRAPH '93*, ACM Press, pp. 57-64, 1993.
9. Rekimoto, J., and Masanori, S., "Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments," in *Proceedings of Conference on Human Factors in Computing Systems CHI '99*, ACM Press, pp. 378-385, 1999.
10. Suzuki, H. and Kato, H., "AlgoBlock: A Tangible Programming Language, a Tool for Collaborative Learning," in *Proceedings of the Fourth European Logo Conference*, Athens, Greece, pp. 297-303, 1993.
11. Ullmer, B., et al., "mediaBlocks: Physical Containers, Transports, and Controls for Online Media," in *Proceedings of SIGGRAPH '98*, ACM Press, pp. 379-386, 1998.
12. Ullmer, B. and Ishii, H., "The metaDESK: Models and Prototypes for Tangible User Interfaces," in *Proceedings of Symposium on User Interface Software and Technology UIST '97*, (Banff, Alberta, Canada, October, 1997), ACM Press, pp. 223-232, 1997.
13. Underkoffler, J., and Ishii, H., "Urp: A Luminous-Tangible Workbench for Urban Planning and Design," in *Proceedings of Conference on Human Factors in Computing Systems CHI '99*, ACM Press, pp. 386-393, 1999.
14. Wacom Technology, <http://www.wacom.com>
15. Wellner, P., "Interacting with paper on the Digital Desk," *Communications of the ACM*, 36(7), 86-96, 1993.
16. Wellner, P., "The DigitalDesk calculator: Tangible manipulation on a desk top display," in *Proceedings of ACM Symposium on User Interface Software and Technology, UIST '91*, ACM Press, pp. 27-34, 1991.