

A Tangible Interface for Organizing Information Using a Grid

Robert J.K. Jacob*, Hiroshi Ishii, Gian Pangaro, and James Patten

Tangible Media Group, MIT Media Laboratory, Cambridge, Mass. USA
{rjacob, ishii, pangaro, jpatten}@media.mit.edu

ABSTRACT

The task of organizing information is typically performed either by physically manipulating note cards or sticky notes or by arranging icons on a computer with a graphical user interface. We present a new tangible interface platform for manipulating discrete pieces of abstract information, which attempts to combine the benefits of each of these two alternatives into a single system. We developed interaction techniques and an example application for organizing conference papers. We assessed the effectiveness of our system by experimentally comparing it to both graphical and paper interfaces. The results suggest that our tangible interface can provide a more effective means of organizing, grouping, and manipulating data than either physical operations or graphical computer interaction alone.

Keywords

Tangible user interfaces, physical interaction, computer-supported cooperative work, sensing board, conference planning, scheduling, direct manipulation, experiment.

INTRODUCTION

Tangible user interfaces (TUI) have been most compelling in spatial or geometric application domains such as urban planning[15], where the physical arrangement of the objects to be manipulated has an obvious, inherent meaning in the application. We want to explore the use of tangible user interfaces in a wider range of more abstract information tasks, where they have been less fully developed. We have developed a new platform and tangible user interface for manipulating, organizing, and grouping pieces of information, which we believe to be especially suited to tasks involving discrete data, abstract or non-geometric data, and collaborative group work. We present our complete system, Senseboard, and an application built with it (see Figure 1). We then use it in a more abstract, simplified form in an experiment to evaluate the performance of our tangible interface.

The basic task our system addresses is common to many application areas: manipulating a set of information items or nodes. People often perform this task by arranging paper

notecards on a desk or, by collaboratively arranging sticky notes on a board. Such arrangement often begins in a free-form way, by accreting small groups of related items, and, later, develops into a larger structure or framework. Tasks like this have thus far been surprisingly resistant to computer support, perhaps because notecards or sticky notes allow manipulation that is more natural and fluid and particularly, free form without a predefined framework. Even when the information to be organized already exists in electronic form and the final output must be produced in digital form, many people find it advantageous to copy the information onto pieces of paper, manipulate them manually, and then re-enter the resulting organization into the computer.

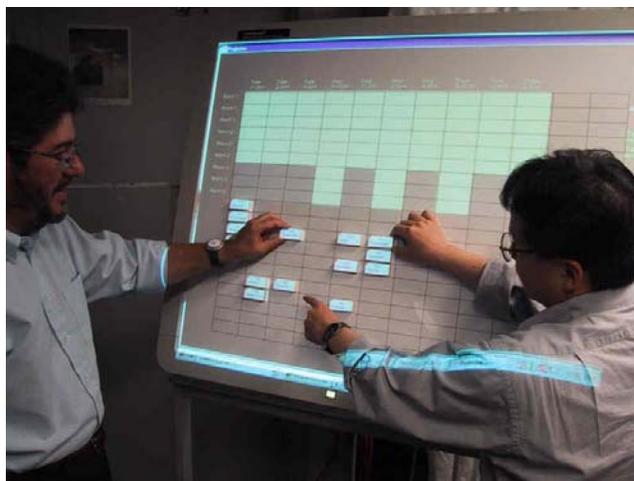


Figure 1. The Senseboard in use.

By providing a tangible user interface for this task, we seek to blend some of the benefits of manual interaction with those of computer augmentation. We expect the basic manual physical approach to give us: (1) a natural, free-form way to perform organizing and grouping; (2) rapid, fluid, two-handed manipulation, including the ability to grab and move a handful of items at once (in contrast to mouse interaction); and (3) a platform that easily extends to collaboration (unlike the conventional mouse and keyboard interface). In addition to these features, our system adds functionality that only the computer can provide, such as: (1) selectively revealing details about each item; (2) displaying alternate views of the data; (3) sorting and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2002, April 20-25, 2002, Minneapolis, Minnesota, USA.
Copyright 2001 ACM 1-58113-453-3/02/0004...\$5.00.

*and Department of Electrical Engineering and Computer Science, Tufts University, Medford, Mass. USA

searching; (4) saving and restoring alternative arrangements of the data; (5) exporting the arrangement to some final output medium (e.g., schedule, shelf list, web pages); and (6) allowing easy backup of an arrangement, which is not possible with a set of physical sticky notes on a board in a meeting room. Current approaches give one or the other of these two sets of advantages. We want to provide a point in the design space between physical and digital, which blends some of the benefits of each in a single platform.

We also present an experimental study that compares our tangible interface to both a purely physical and a purely graphical counterpart and suggests better performance with the tangible interface than either of the two alternatives. We then investigate this issue further, since we expect that a tangible interface will contain selective benefits of physical and digital interfaces rather than all the benefits of both. We believe our interface preserves some of the good qualities of physical paper, but not all, because the TUI is inevitably an imperfect simulation of the real world. We attempt to quantify this tradeoff by comparing to a tangible interface with all benefits of computer enhancement removed. This allows us to measure the performance penalty paid for the imperfect simulation inherent in the TUI separately from the benefit of the computer enhancements; our goal is to design a system in which the latter outweighs the former.

TANGIBLE USER INTERFACES

Tangible user interfaces are a growing area of user interface research that use physical forms to represent data. They were so named by Ishii and Ullmer[3], with roots in seminal work of Wellner[9], and Fitzmaurice[1], and continue a growing range of more recent work[2, 7, 8, 10, 11, 14, 15]. Instead of a generic screen, mouse, and keyboard capable of representing all types of data, TUI uses specific physical forms to represent and manipulate the pieces of data in the system. TUI often uses simple, transparent mechanical structures, so the user can use his or her existing knowledge of the physical world to figure out how to operate them and what they mean. TUI also often involves the augmentation of existing physical objects by adding digital meaning to the objects and their manipulation[11, 16]. Some information and state is thus represented directly by the physical objects, while additional information is provided digitally, typically by video projection onto the physical objects (e.g., a digital shadow of a physical building model). We attempt to combine the physical and digital representations to exploit the advantages of each.

Tangible user interfaces can also be viewed as the further evolution of GUI or direct manipulation style interaction. The essence of a GUI is that the user seems to operate directly *on* the objects in the computer rather than carrying on a dialogue about them[12]. However, while the objects are depicted realistically on the GUI display, the input and output devices themselves are generic and malleable, typically a mouse and a raster display. The price paid is

that their physical form gives no useful information to the user about the objects, the task, the possible commands, or the state of the system. TUI increases the realism of the objects by allowing the user to interact even more directly with them.

A concept underlying these interaction styles is to build on the “natural” equipment and skills humans have acquired through evolution and experience and exploit these for communicating with the computer. For example, we posit that grabbing and moving real objects directly is a more natural skill than pointing or referring to them indirectly. Direct manipulation moved user interfaces in this direction; TUI pushes further toward such “natural” interaction. However, an interface typically will rarely be *entirely* “natural” or physical (or else it would not require a computer at all); instead the design goal is to try to do as much as possible with natural interactions, before adding “artificial” ones[13].

APPLICATION DOMAIN

We want to suggest a plausible future tangible user interface for performing a realistic knowledge worker or office task. The basic application we consider is the generic task of arranging, organizing, or grouping data objects. These data items might be messages, files, bookmarks, citations, papers for a literature search, slides for a presentation, scenes for a movie, newspaper stories, pages for a web site, employees to be reorganized, MP3 files to be played by a disc jockey, ideas from a brainstorming session, or papers to be presented at a conference.

For specificity, we focus here on a task familiar to some readers: the job of organizing conference papers into sessions and scheduling them. It shares most of the key properties of other more general information manipulation tasks, and we will use it as a stand-in for this general class of tasks, without loss of generality. The conference paper task we address begins after the papers have been accepted or rejected, and involves two phases. First, it requires the accepted papers to be grouped into sessions of three papers each, related to each other by topic. Since the papers are accepted or rejected independently, without regard to whether they will fit into groups of three, the resulting papers do not fall neatly into such groups, and this task often requires some false starts and tradeoffs to fit them into groups. The second phase of the task requires that the sessions be assigned to time slots. In our example, there are two or three sessions at the same time, so sessions occurring in parallel should not overlap each other in topic areas nor in the speakers who must be present for the sessions.

The first part of the task thus provides an example of an unstructured organizing or grouping task (building up small groups by accretion), while the second one is a more structured task (assigning items to predefined slots). Together, they cover the important features of the wider range of example applications mentioned above. Though we draw on one of the authors' (Jacob) experience as Papers Co-chair of the ACM CHI 2001 conference, note

that this system was not actually used to produce the CHI 2001 papers program. In fact, the task was performed for CHI 2001, as in other years and in many other conferences, using manual, “tangible” interaction. For the first phase, we printed out each of the 69 accepted papers (which had originally been submitted electronically) and spread them out on the floor. Committee members worked together to move the papers around on the floor, accumulating small groupings of related papers, exchanging papers to improve groupings, and gradually converging on 23 groups of three each. For the second phase, we made a single index card for each group of three papers (rather awkwardly, by printing the list of accepted papers, cutting the printout into little strips, and pasting the appropriate strips onto cards). We then stuck the cards onto a large schedule grid drawn onto a whiteboard and examined and rearranged them several times. When the arrangement was finalized, we manually transcribed the data from the schedule board onto a computer. Both phases were performed by groups of people (a subset of the Associate Papers Chairs) working collaboratively.

Despite having excellent computer support at this meeting (held in a computer-enhanced meeting room at Microsoft Research), having all the relevant data already in electronic form, and needing to produce our final results in electronic form, the committee found it more effective to use manual/tangible interaction to perform this task this year, as in previous years. All of the input data were “born digital.” The CHI papers were submitted electronically as PDF files, along with a web-based form for entering title, author, and other information; and the final destination was a digital database from which the CHI hardcopy advance program and web site were generated. Our system shows an example of how the benefits of physical interaction can be extended to a computer system.

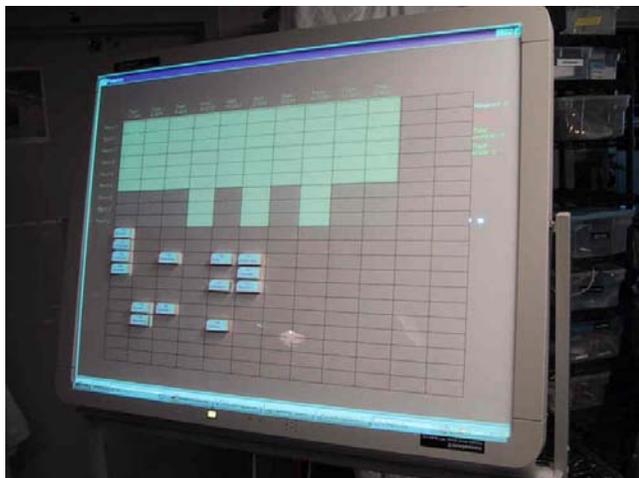


Figure 2. The Senseboard device.

DESIGN

Senseboard consists of a vertical panel 1.1 m. wide x 0.8 m. high, mounted rather like a portable whiteboard, as seen in Figure 2. It is marked with a rectangular grid of 12 columns x 20 rows, covering most of its surface, giving it the appearance of a spreadsheet layout. Small rectangular

plastic tags or “pucks,” resembling white refrigerator magnets, can be placed into these cells and stick there magnetically. Each puck is 5.8 cm. wide x 2.5 cm. high and is magnetized so that it can only be placed squarely into a grid cell. It contains a small RFID (radio frequency identification) tag, which is a passive device polled by a radio frequency signal transmitted by the board beneath it. Each time the user moves a puck, the board sends the identity and the grid location of each of the pucks in the grid to a computer on a serial port. The board is a component of a Bannou Pro Intelligent White Board™ (Uchida Yoko Ltd., Japan). The pucks are our own design, based on those of the Bannou Pro device, but modified so that the user can send additional commands to the system by pressing the surface of the puck or by briefly placing one of several special command pucks over the puck. We coupled this system to a computer (Dell Pentium II PC) driving a video projector, which projects general information onto the board and specific information onto each puck. Our software that generates the projected data is written as a Java application running on Windows 98, receiving input from the board over a serial port and sending its output to the video projector.

Our system combines projection with physical manipulation in ways similar to the early work of Wellner[9] and Fitzmaurice[1] and more recent work such as DataTiles[11], metaDESK[14], and Sensetable[10]. Unlike several previous devices that used computer vision, the Senseboard uses RFID sensing, which provides greater reliability and speed. Other related projects include LiveBoard[6], LegoWall (described in [2]), and Outpost[5]. Our device is also unusual in that it supports discrete, semi-structured interaction, since the pucks can only be placed directly in grid cells. It is best used for tasks where completely free-form input would not be appropriate—where the physical constraint of placing pucks directly into grid cells mirrors the constraints of the task itself and thus provides a useful affordance for the task. Scheduling with respect to a grid of discrete time slots is one such task. In our application, the available time slots and the number of papers that can be assigned to each are represented naturally and implicitly by the arrangement of the grid cells and the fact that only one puck can be physically inserted into each grid cell. Finally, because the pucks are held to the board by the grid of magnets, the board can be angled vertically, rather than horizontally, allowing easier use by a larger group of people working together. In sum, the use of unique RFID tags, multiple pucks operating on the same platform, a vertical work surface, and a constrained grid layout together form a new tangible interface platform, which we use to explore novel interaction techniques.

Rationale

Our overall goal is to make the data items we present tangible and graspable; the user should feel like he or she is touching the real data, as if they were “tangible bits.” Our interface design decisions are guided by an attempt to make each puck seem like a real computer data object, rather

than like a handle that is controlling remotely-located displayed data. Our platform accommodates plenty of pucks, and we normally provide a separate puck for each data item, permanently attached to it, although exceptions are permitted for some operations. Interacting using a mouse could be viewed as the opposite extreme, in that there is only one physical “puck,” and it can be attached to only one data object at a time, usually for very brief periods. Bricks[1] provides an intermediate point. We also exploit physical representation, with specially shaped physical objects to depict commands, and use purely digital representation where appropriate, such as continuous display of the implications of the user's physical actions.

Data Objects

In the basic interface, each data object (a conference paper) is projected onto a cell with a puck. We begin with a design intentionally devoid of specific interaction techniques: the user simply interacts “naturally” with the objects, by seeing them, grabbing them, and moving them. To organize the papers, the user need only move them around on the grid to form any desired arrangement or grouping. In our sample task each object is a full paper accepted for the CHI 2001 conference. A magnetic puck represents each item with some information about the paper (number, title, author) projected onto its smooth white surface. If the user moves a puck to a new cell, its data moves with it. Figure 1 shows the board in use, with each paper represented by a puck on the lower region of the board. The users have moved the papers around the board, and are beginning to arrange them by topic similarity. Here, the board grid is divided into two regions, one for free-form topic grouping at the bottom and one representing the time schedule of the conference at the top.

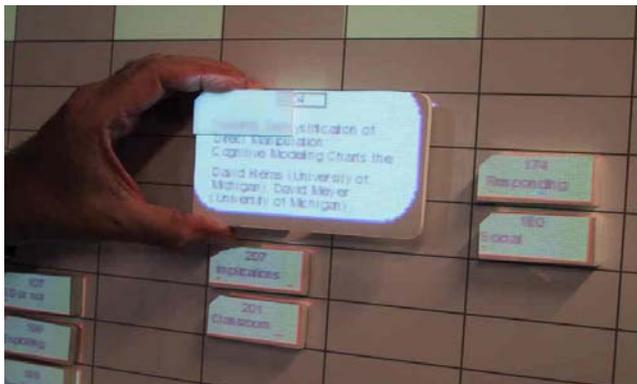


Figure 3. Using the “View Details” command puck to see more information about one of the data items by placing it over that item. Note how the details are projected onto the larger projectable surface of the command puck, temporarily obscuring several adjacent pucks.

Commands

Pucks represent operands or nouns in a very straightforward way. We add operators or verbs by embodying each command in a special puck, shaped differently from the others, to be thought of as a tool or operator or stamper, which can be momentarily placed over a regular data puck to operate on it like a rubber stamp.

These pucks are noticeably thicker than the regular pucks, with distinctive shapes, suggesting the specific function of each tool. Thus flat pucks represent data, and tall, specially shaped pucks represent tools; the tools are placed over the data pucks.

The “View Details” command puck has a larger projectable surface than the data pucks. It allows the user to view additional information about a given item, temporarily overcoming the limitation of the small size of the data pucks. We display information that would not normally fit onto a puck by temporarily obscuring the adjacent pucks. The “View Details” command puck is placed over the puck to be viewed, and the information is displayed until the command puck is removed. The detail information is projected onto the larger command puck area, and the fact that it physically obscures the adjacent cells below it suggests to the user, in a direct physical way, that this temporary information is being placed over those cells, and that the cells below it are still present, but temporarily obscured (see Figure 3).

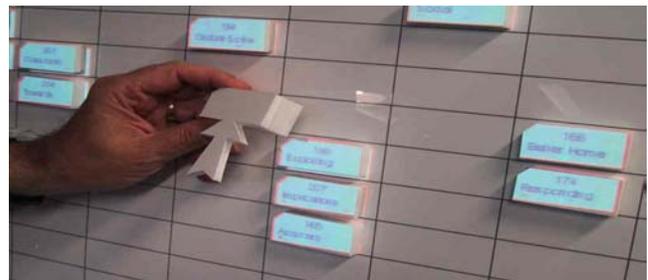


Figure 4. “Group” command: The user is about to execute the command by placing its command puck (with the arrow) over the topmost data item of the column to be grouped. The items will then be grouped into a single item, at the location of the command puck.

The “Group” command allows the user to take several items and combine them into a new, single item, representing the group. For the CHI papers task, this is like grouping (usually three) papers into a single conference session. You first arrange the items on the board to form small groups of interest, placing the items you wish to group near each other on the board, in a single contiguous column. You then apply the Group command by placing the command puck over the first of the items to be grouped, rather like first arranging papers together and then stapling them. The arrow shape of this command puck suggests its function of swallowing the pucks below into a new group (see Figure 4). An “Ungroup” command operates analogously and explodes a displayed group item back into its component items. Its command puck is a down arrow, suggesting it will explode the current item into cells below.

To show how a user might enter new data from a keyboard, we provide a “Type-in” command. The typed text appears at the bottom right of the board, outside of the cell grid. You then place the corresponding command puck on a blank cell and a new item appears there containing the newly-entered data.

The “Copy” or “Link” command displays the original data item and the copy, with a line drawn between them; each can be manipulated with its own puck. This demonstrates how a user might create a line or graph edge rather a new item. The command is intended to allow a user to explore alternative organizations of the data, where one item might need to appear in two places at once. You place the command puck on the original item and then on a blank puck where you want the copy.

In addition to the information projected onto individual pucks, we display general information on the board. This is where computer augmentation provides a clear advantage over sticky notes. When the user places a puck representing a paper into a time slot on the schedule grid, any resulting conflicts are shown graphically. For example, when papers on similar topics are presented in parallel, color-coded lines are drawn between conflicting papers. This information is updated continuously, so the user can see the implications of choices as he or she makes them (Figure 5). Perhaps the most obvious advantage of computer augmentation is the “Export” command, which allows users to convert the final arrangement of the items into a digital form that can be exported to another program, avoiding the need to transcribe the results after arranging sticky notes on a board.

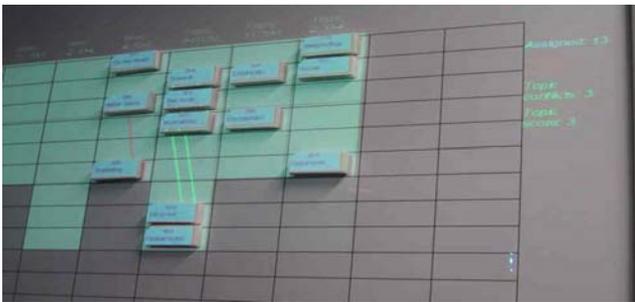


Figure 5. Displaying statistics and conflict information: As the user places papers into the slots of the time schedule in the upper area of the board, the system displays conflicts between papers scheduled for the same time and involving the same author or similar topics. Lines represent conflicts and overall statistics are shown at the right edge of the board.

Alternative Designs

We also implemented a second approach for representing commands. Instead of a separate puck for each command, we display the commands in a reserved area of the board grid, along the bottom. To apply a command to a puck, you bring the puck to the cell that represents that command. For example, the View details command is a single cell at the bottom left of the board. When you place a puck there, its information is shown, just below it, outside of the grid. The Group command area has cells for several member items plus one for a newly created group item. You place the pucks for the objects to be grouped into the member item area, and then place a new puck in the group item area, and a new group item appears there, attached to the new puck. This approach seems conceptually simpler, and somewhat more visible to the user, but we generally prefer the command puck alternative, because it allows the user to

operate on data items without moving them from their current locations. It also exploits the physicality of the command puck to communicate its function.

We also provide a third approach, in which the user simply presses on the surface of a puck to execute a command on it. This is convenient, but more limited, because it can handle only one command, much like the GUI double-click. It can coexist with the others, just as GUIs use double-click for one command, combined with other means for executing the remaining commands. We use this as an alternative View details command.

The overall design begins to define an interaction language for using pucks on a grid. Its syntactic elements include thin pucks for nouns and thick ones for verbs, stamping one puck over another, contiguous groups of pucks, pressable pucks, and representing commands by special puck shapes or reserved spatial locations. Our system shows these for the conference papers task as a concrete example of a broad class of information organizing tasks. The data items could equally well be newspaper stories, movie scenes, music clips, paper citations, or web pages.

EXPERIMENT

We conducted an experiment to attempt to quantify the costs and benefits possible from tangible interaction, compared to graphical and to purely physical interaction. Because the TUI provides natural interaction with real physical objects with their physical affordances and constraints, we expect TUI to provide some benefit compared to GUI. Moving the objects around with your hands can help you think about a problem, because the objects and their configuration embody and retain information about the task state and constraints directly[4]. However, we expect that our TUI cannot match all the benefits of either physical objects or GUI perfectly, but it can provide an otherwise unobtainable blend of benefits of both and, possibly, a performance improvement over each. We also expect TUI to suffer a penalty, compared to pure physical interaction arising from the imperfections in the way that the tangible interface simulates the physical world. However, the TUI can provide additional benefits not possible with plain sticky notes, such as displaying computed information as the user interacts with the objects, so we then expect to gain a performance improvement from such enhancements. The design goal is that, overall, these additional benefits outweigh the penalties paid in going from paper to TUI; our experiment attempts to measure these two components separately.

Task

For this experiment, we simplified the previous application considerably. Since subjects would not be familiar with the intricacies of CHI paper scheduling, we needed a simple, self-contained task that could be learned and performed in a single experimental session. Ours required subjects to plan a working schedule for a group of employees within a set of constraints. We sought to emulate, from the subject's point of view, a more difficult, open-ended, real-world task (such as the CHI papers application), where the rules might

not all be given explicitly, and the task would require more judgment. Our task was designed to exercise similar skills on the part of the subject, but in a more artificial and simplified way for a closed-end experiment. In pilot testing, we determined that the task was sufficiently easy and self-contained and that most subjects could learn it and complete it correctly fairly quickly. Accuracy was high enough across all conditions that speed, rather than quality, could provide a unified measure of performance.

The subject's task was to create a work schedule. Each of 5 days required 3 workers from a pool of 6 and had to meet a set of constraints, e.g., workers had different skills and each day required a mix of skills, some workers could not work on consecutive days. The schedule was presented as a grid on the board (5 days x 3 workers per day); workers were represented by pucks. The subject could assign a worker to a slot in the schedule by placing a puck into a schedule slot, and rearrange them as needed to meet the constraints. The scheduling constraints were listed on a sheet of paper which was read to the subject and then posted on the board for reference. In some of the experimental conditions, the computer continuously checked the subject's schedule and displayed text and graphical messages if it violated any of the constraints.

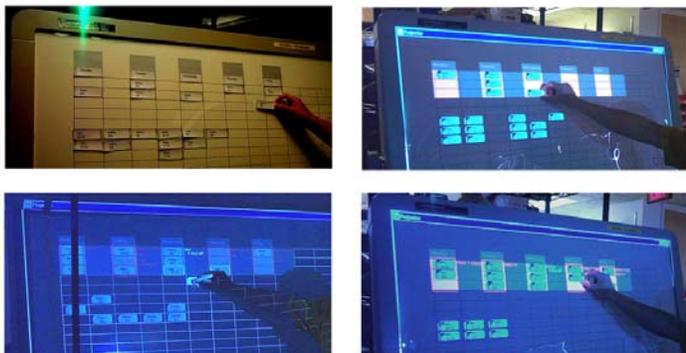


Figure 6. The four experimental conditions: Paper and Reduced-Senseboard (top row, left to right), and Pen-GUI and Senseboard (bottom).

Experimental Design

Subjects performed our task under four different conditions, which we designed to be as similar as possible in overall size, orientation, and general physical setup. Figure 6 shows the four conditions. We used a within-subjects design, with each subject performing the tasks under all four conditions and analyzed the results with a two-factor (condition x subject) ANOVA. We systematically varied the orders in which subjects performed the four conditions as well as the assignment of the four schedule variations to the four conditions. We created four schedule variations of our task, with different names and randomly-generated differences in their skill sets, to reduce the learning effect for each subject, and assigned the variations differently to the four different conditions for each subject.

Paper Condition: One current practice for scheduling tasks is to use pieces of paper or sticky notes, as was done at

CHI. To compare tangible interaction with real physical interaction, we created a version of our experimental task that used conventional paper sticky notes placed on the same vertical board. We had designed the experimental task to require only moving pucks on the board, without pressing or stamping, so that we could have a paper version of the interface for comparison. We pre-printed sticky notes with the same information that was projected onto the pucks (name and skills), in the same size as the pucks, and placed them on the Senseboard itself, which is inactive in this condition, but provides a vertical surface, physical arrangement, and grid identical to the Senseboard condition.

Reduced-Senseboard Condition: An issue with most tangible user interfaces is that their simulation of the real world is imperfect, and we expect this to have some performance cost. In our case, compared to paper sticky notes, there is latency between the user's action and the visible result, possible misregistration of the data projected onto the physical object, lower resolution of the projected display compared to real paper, and loss of the displayed data when the pucks are off the board. Each of these can be addressed with technological improvements (including ways for the computer to write more permanently on the pucks), but, for now, we see our TUI as having many—but not all—of the advantages of paper. We expect some performance decrement due to this imperfect simulation. To measure it, we introduce a condition in which the Senseboard merely emulates the paper sticky notes with the computational enhancements removed. That is, the user moves the same pucks, with their projected labels, but the computer does not display messages if the schedule violates the constraints. We expect this condition to perform worse than paper and worse than regular Senseboard, but to provide a way to tease apart these components of performance.

Pen-GUI Condition: We also compared our system to a more conventional GUI-like computer interface. We wanted a condition closely matched to the physical arrangement of the Senseboard, so we used a digital whiteboard, rather than a regular mouse and keyboard. We used a Microfield Graphics Inc. Softboard 201™, a whiteboard that can track the location of a (non-marking) pen across its surface. We placed the projector at the same height as the Senseboard projector, and tilted the board slightly upward just like the Senseboard. We developed a GUI, drag-and-drop style interface for our task, using the same Java program, modified to allow interaction by dragging the pen on the board instead of the Senseboard pucks, but with the same projected data display, and same size and physical arrangement of the board.

We tried to give this system nearly all the properties of the Senseboard except for tangible pucks. The result was an interface that was not quite a conventional mouse-and-keyboard GUI, but was closer to the Senseboard in ways that probably improve its performance on this task over a conventional GUI. Like a touchscreen, it provides more

direct interaction than a mouse, and it had somewhat better latency and registration than the Senseboard. For these reasons, we expected it to be an attractive choice. What it lacks is the physical manipulation of the pucks or sticky notes, representing data items that can be held in the user's hand, which can help him or her think about the problem.

Senseboard Condition: Finally, we used our Senseboard as described in the previous section. The projector was placed 2 m. above the floor, and the board tilted slightly upward. We used the original Bannou Pro pucks, since this task did not use the pressing and stamping functionality of our new pucks.

Subjects

We used 13 subjects, 6 male and 7 female, recruited with posters on bulletin boards around the MIT campus. Sessions lasted 30-45 minutes; we paid each subject \$10.

Procedure

The task was presented to subjects as a set of sticky notes, pucks, or projected GUI data, arranged in the same initial configuration. For each condition, we explained the task to the subject, demonstrated the operation of the equipment, and then asked the subject to perform the task. Our system measured elapsed time to perform the task and recorded the final schedule they created (and also recorded every move of a puck or data item, for further investigation). For the paper condition, the computer simply recorded elapsed time for uniformity, and we transcribed the arrangement of the sticky notes manually. We then asked the subject to fill out a questionnaire about the task and the four conditions and interviewed them to elicit any further comments or thoughts.

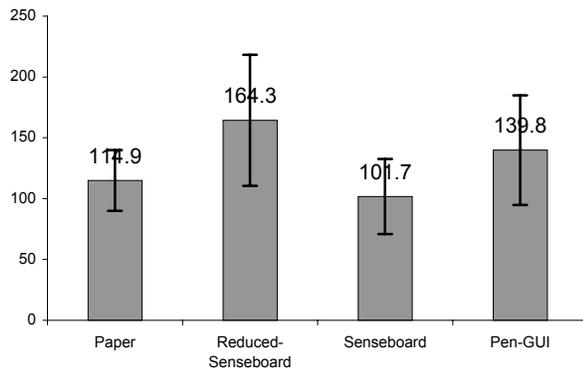


Figure 7. Experimental results: time (in seconds) to complete the task for the four experimental conditions, mean and 95% confidence interval. Note that the first three conditions use tangible interaction to varying degrees; the last two use computer augmentation; only the Senseboard condition uses both.

Results

We collected subjects' performance on each of the four tasks, questionnaire answers, and interview comments. We found that subjects completed nearly all tasks correctly (8 errors out of 780 assignments) so we used the time to complete the task as an overall performance measure. Figure 7 shows the mean time for each condition. The

trends we expected are suggested by these data, though the results are only weakly statistically significant (ANOVA for condition effect gives $F(3,36) = 2.147$, $p = 0.11$). We also analyzed the data for the effect of order of presentation or learning, and found a weak effect ($p = 0.08$), but as noted we had counterbalanced the orders across subjects and conditions for this. Finally we checked to make sure there was no effect of the four schedule variations and found none ($p = 0.32$).

Our questionnaire asked subjects to rate how well they liked each of the four conditions on a 7-point Likert scale. We found a significant overall effect of condition ($F(3,36) = 3.368$, $p = 0.03$), a weak preference for Senseboard over each of the other three conditions, and a substantial dislike of the Paper condition. The questionnaire also asked subjects about the strong and weak points of each of the conditions. Many commented on the value of manipulating the physical tags to aid in thinking; a typical response was: "I like the idea of manipulating something, makes it easier to tell who you're scheduling where." Some subjects were also attracted to the Pen-GUI because they had never seen such a device. Finally, some subjects complained that their head and hands sometimes blocked the projected data; this applied uniformly to all our conditions except Paper. Our plans to mount the projector on the ceiling, aimed downward at 30-45 degrees, and tilt the board upward to match will greatly reduce this problem.

Discussion

The two current user interfaces most likely to be used today for a task like ours are paper notes or computer GUI. Each has strengths and weaknesses, but it is difficult to blend their strengths—to combine fluid, physical manipulation or "tangible thinking" with the features of computer automation to yield better performance than either alternative alone. Our results suggest, though at weak statistical significance levels, that the new tangible interface (Senseboard condition) may indeed yield better performance than either pure physical (Paper) or GUI (Pen).

In comparing TUI to GUI, we opted for a GUI condition that closely matched the physical setup of the Senseboard, to eliminate extraneous factors from the comparison. It was thus not a pure GUI, but a hybrid that included some properties of the TUI. It allowed more direct interaction (touching the screen) than a mouse (at a cost of likely arm fatigue for a longer task than ours). Our particular implementation also had somewhat less latency than the Senseboard and less critical registration requirements for the projected image. We believe that these factors enhanced its performance relative to the TUI and that a conventional GUI with a mouse and desktop display would have given slightly worse performance, but it might be a valuable condition for future study because it is so widely used. It would also be interesting to compare a "Reduced-GUI" condition (graphical manipulation but no computer augmentation), analogous to our Reduced-Senseboard

condition, though we would expect still worse performance.

In comparing TUI to paper, we believe the tangible pucks preserve some of the fluidity and “tangible thinking” qualities of paper, but do so incompletely. We see a small improvement for TUI over paper and then use the Reduced-Senseboard condition to attempt to decompose that improvement into two larger component parts. We decompose this difference into (1) the price paid for the imperfect way TUI simulates paper (difference between Paper and Reduced-Senseboard) and (2) the benefit gained from automation (difference between Reduced-Senseboard and complete Senseboard). We also view this decomposition as measuring the value of “natural” interaction (Paper), minus the cost of simulating it (Reduced-Senseboard), plus the benefit of “artificial” additions. As the price paid for the imperfect simulation is reduced by future technological developments, such as lower latency, better display technology, and new bistable display materials that would let the pucks retain their displayed information, the advantage for TUI can become stronger, while the benefit of the automation will be retained.

CONCLUSIONS

Senseboard shows how features of physical manipulation (collaboration and rapid, fluid, two-handed free-form manipulation) and digital interaction (display statistics and implications, check constraints, view details, group and ungroup, import, export, save, restore, sort) can be blended into a single platform that suggests better performance than either alternative alone. We presented our system design, interaction techniques, design rationale, example application, and an experiment to study it. Our conference paper task is representative of a larger class of generic arranging, organizing, or grouping tasks for which this type of interface should be useful, such as papers for a literature search, pages for a web site design, files on a computer disk, slides for a presentation, employees to be reorganized, or ideas in a brainstorming session. We demonstrate how a tangible user interface can be provided for this kind of abstract or non-geometric problem domain to yield better performance than existing alternatives.

ACKNOWLEDGMENTS

We thank Kenji Alt, Katherine Butler, Angela Chang, Flora Chiu, Zahra Kanji, and Gustavo Santos for their work on hardware and software for the Senseboard; and our colleagues in the Tangible Media Group at MIT, including Brygg Ullmer, Mike Ananny, and Dan Maynes-Aminzade; and in the Supply Chain Visualization project at Intel and at MIT Sloan School for discussions, suggestions, and help. This work was supported by the Things That Think and Digital Life Consortia of the MIT Media Lab and the Intel Corporation.

REFERENCES

1. G.W. Fitzmaurice, H. Ishii, and W. Buxton, “Bricks: Laying the Foundations for Graspable User Interfaces,” *Proc. ACM CHI'95 Human Factors in Computing Systems Conference*, pp. 442-449, Addison-Wesley/ACM Press, 1995.
2. G.W. Fitzmaurice and W. Buxton, “An Empirical Evaluation of Graspable User Interfaces: Towards Specialized, Space-multiplexed Input,” *Proc. ACM CHI'97 Human Factors in Computing Systems Conference*, pp. 43-50, Addison-Wesley/ACM Press, 1997.
3. H. Ishii and B. Ullmer, “Tangible Bits: Towards Seamless Interfaces between People, Bits, and Atoms,” *Proc. ACM CHI'97 Human Factors in Computing Systems Conference*, pp. 234-241, Addison-Wesley/ACM Press, 1997.
4. D. Kirsh, “Complementary Strategies: Why we use our hands when we think,” *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, Lawrence Erlbaum., Hillsdale, N.J., 1995.
5. S.R. Klemmer, M.W. Newman, R. Farrell, M. Bilezikjian, and J.A. Landay, “The Designers' Outpost: A Tangible Interface for Collaborative Web Site Design,” *Proc. ACM UIST'01 Symposium on User Interface Software and Technology*, pp. 1-10, ACM Press, 2001.
6. T.P. Moran, P. Chiu, W. van Melle, and G. Kurtenbach, “Implicit Structure for Pen-based Systems Within a Freeform Interaction Paradigm,” *Proc. ACM CHI'95 Human Factors in Computing Systems Conference*, pp. 487-494, Addison-Wesley/ACM Press, 1995.
7. T.P. Moran, E. Saund, W. Van Melle, A.U. Gujar, K.P. Fishkin, and B.L. Harrison, “Design and Technology for Collaborative Collages of Information on Physical Walls,” *Proc. ACM UIST'99 Symposium on User Interface Software and Technology*, pp. 197-206, Addison-Wesley/ACM Press, 1999.
8. L. Nelson, S. Ichimura, E.R. Pedersen, and L. Adams, “Palette: A Paper Interface for Giving Presentations,” *Proc. ACM CHI'99 Human Factors in Computing Systems Conference*, pp. 354-361, Addison-Wesley/ACM Press, 1999.
9. W. Newman and P. Wellner, “A Desk Supporting Computer-based Interaction with Paper Documents,” *Proc. ACM CHI'92 Human Factors in Computing Systems Conference*, pp. 587-592, Addison-Wesley/ACM Press, 1992.
10. J. Patten, H. Ishii, J. Hines, and G. Pangaro, “Senseable: A Wireless Object Tracking Platform for Tangible User Interfaces,” *Proc. ACM CHI 2001 Human Factors in Computing Systems Conference*, pp. 253-260, ACM Press, 2001.
11. J. Rekimoto, B. Ullmer, and H. Oba, “DataFiles: A Modular Platform for Mixed Physical and Graphical Interactions,” *Proc. ACM CHI 2001 Human Factors in Computing Systems Conference*, pp. 269-276, Addison-Wesley/ACM Press, 2001.
12. B. Shneiderman, “Direct Manipulation: A Step Beyond Programming Languages,” *IEEE Computer*, vol. 16, no. 8, pp. 57-69, 1983.
13. R.B. Smith, “Experiences with the Alternate Reality Kit: An Example of the Tension Between Literalism and Magic,” *Proc. ACM CHI+GI'87 Human Factors in Computing Systems Conference*, pp. 61-67, 1987.
14. B. Ullmer and H. Ishii, “The metaDESK: Models and Prototypes for Tangible User Interfaces,” *Proc. ACM UIST'97 Symposium on User Interface Software and Technology*, pp. 223-232, Addison-Wesley/ACM Press, Banff, Canada, 1997.
15. J. Underkoffler and H. Ishii, “Urp: A Luminous-Tangible Workbench for Urban Planning and Design,” *Proc. ACM CHI'99 Human Factors in Computing Systems Conference*, pp. 386-393, Addison-Wesley/ACM Press, 1999.
16. R. Want, K.P. Fishkin, A. Gujar, and B.L. Harrison, “Bridging Physical and Virtual Worlds with Electronic Tags,” *Proc. ACM CHI'99 Human Factors in Computing Systems Conference*, pp. 370-377, Addison-Wesley/ACM Press, 1999.