

Clutching at Straws: Using Tangible Interaction to Provide Non-Visual Access to Graphs

David McGookin, Euan Robertson, Stephen Brewster

Department of Computing Science

University of Glasgow

Glasgow G12 8QQ

{mcgookdk, stephen}@dcs.gla.ac.uk, www.dcs.gla.ac.uk/~mcgookdk

ABSTRACT

We present a tangible user interface (TUI) called Tangible Graph Builder, that has been designed to allow visually impaired users to access graph and chart-based data. We describe the current paper-based materials used to allow independent graph construction and browsing, before discussing how researchers have applied virtual haptic and non-speech audio techniques to provide more flexible access. We discuss why, although these technologies overcome many of the problems of non-visual graph access, they also introduce new issues and why the application of TUIs is important. An evaluation of Tangible Graph Builder with 12 participants (8 sight deprived, 4 blind) revealed key design requirements for non-visual TUIs, including phicon design and handling marker detection failure. We finish by presenting future work and improvements to our system.

Author Keywords

Tangible User Interface, Visual Impairment, Haptic Interaction, Graphs

ACM Classification Keywords

H.5.2 Information Interfaces and Presentation: User Interfaces—*Interaction Styles*

General Terms

Design, Experimentation, Human Factors

INTRODUCTION

Understanding spatial visualisations such as mathematical graphs, charts and maps is an important life skill. They are common in newspapers, magazines and within many scientific disciplines. Jones and Careras [10] report that over 2.2 trillion graphs were published in 1996 and this number has been steadily rising since. The inability to access and understand these visualisations can severely limit career choices available to an individual. Such access is particularly problematic for people who are blind or have visual impairments [6]. By their very nature, line graphs, maps, bar and pie charts exploit basic aspects of our visual perceptual and cognitive systems to communicate information quickly and efficiently. Non-visually this means that graphs require effort to

understand, losing many of their advantages. This problem is compounded by a lack of efficient tools to allow construction and browsing by visually impaired users.

The most common way to provide access to a visually impaired user is to print a graph on special heat sensitive paper. When the printed graph is passed through a heat printer, the surface raises up creating a tactile relief that can be explored using the fingers. However, this technique is cumbersome, requiring special formatting of the graph and two passes through a printer. Whilst a visually impaired person can create such graphs by him or herself, the graph cannot be inspected until it has been produced in tactile form, causing any necessary changes to be costly, as the graph cannot be modified after creation. This lack of interactive construction is especially problematic in education settings; students are learning about coordinate systems and graph drawing for the first time and mistakes are common. In such situations a tactile paper grid attached to a corkboard is used. The user inserts map pins into the board, which are connected with rubber bands to create graph features - such as data lines and axes. Whilst this technique is common, inexpensive and flexible, it also brings problems. Firstly, the pins used are sharp to allow them to be pushed into the board and take the tension of the rubber bands. This creates a risk that the pin may be pushed into the finger. As the rubber bands are under tension, it can be difficult to pull them far enough to fit over the pins, with any failure to do so rewarded with the band flying off which the user may, or may not, notice. Additionally, if the graph must be modified, it is necessary to partially deconstruct by removing both bands and pins. The graph, once created, cannot be stored as would be the case for a sighted user drawing on paper, as the materials must be reused [15, 22]. This limits the practicality of the corkboard technique to simple teaching scenarios and makes it unsuitable in business settings such as modifying a graph produced in Microsoft Excel.

CORKBOARD CONSTRUCTION OBSERVATION

To understand further the problems of graph construction using the corkboard technique, we asked a 17 year old blind, male student at the Royal National College (RNC) Hereford UK, to carry out some graph transcription exercises. We asked the student to copy eight line graphs that had been created on raised paper to a grid on a corkboard using the pin and rubber band technique. Each graph had one data series which contained 3-4 control points. We video recorded the interaction and photographed the completed graphs.

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, to publish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2010, April 10-15, 2010, Atlanta, Georgia, USA.

Copyright 2010 ACM 978-1-60558-929-9/10/04....\$10.00.

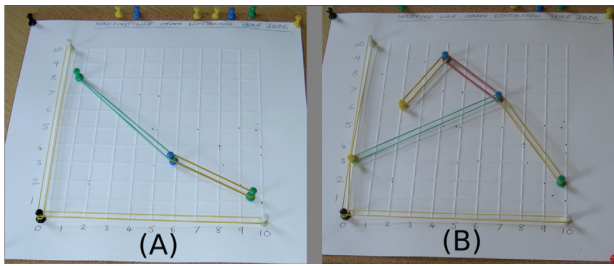


Figure 1. Examples of correctly (A) and incorrectly (B) created line graphs constructed by a visually impaired person using the corkboard construction technique.

When carrying out the tasks the participant explored and interacted with the graph using both hands. To construct, the participant used one hand to mark the origin of the grid and then counted along and up the grid with the other hand to mark a position, before using the origin marking hand to retrieve and insert a pin. The participant also compared by using a hand to mark a point on the corkboard whilst referring to a point on the tactile paper graph. There were several cases where the rubber bands became detached, either when trying to explore the graph and re-orientate, or when the band had not been correctly attached. The participant rarely pushed the pins all the way into the board, making them vulnerable to being pulled out when the rubber bands were applied (as a safety measure the experimenter ensured that all pins were fully inserted into the board). Overall the student correctly completed four out of the eight graphs. The other four graphs had numerous problems including rubber bands which had not been correctly attached and just formed shapes, as well as instances where a rubber band that had become detached had been placed on an incorrect pin, causing the graph to be meaningless. Figure 1 shows both of these situations.

Although we worked with only one participant, the findings graphically illustrate the problems of graph construction that have been identified in group discussions with visually impaired users [15, 22] and how such access is greatly hampered by the relatively primitive technology available.

RELATED WORK

To overcome the limitations of the graph access and creation tools previously discussed, several researchers have investigated the use of virtual haptics and non-speech audio to provide access to graph-based data. The earliest of this work, by Mansur [13], used a pitch based mapping to communicate the y-axis value of a data series, with the x-axis mapped to time. As the sound was played, the user could gain an impression of how the data series changed over time. This work has been extended by Smith and Walker [19], who found that the addition of context through auditory “tick marks”, which serve a similar function to the visual grid lines, improved accuracy in understanding the graph. The work of Mansur, in using a value pitch mapping, lies at the centre of most work in accessing mathematical data through sound and has proven to be robust over several different types of graph. Flowers and Hauer [8] have shown that user understanding of pitch value mappings of box plots and scatter plots is comparable to visual representations of the same data. However,

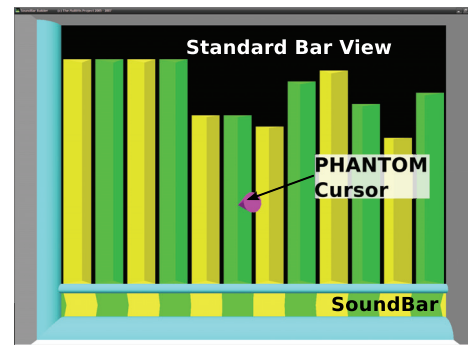


Figure 2. A screenshot of the SoundBar Builder system. Bars are represented by grooves and can be dragged up and down using the PHANTOM Omni.

whilst auditory graphs have proven to be successful and the technology required to generate them is low cost, there are no effective means to manipulate the graphs interactively.

The other major approach to improving graph and chart accessibility is to use haptic interaction. Most attempts have looked at using force-feedback technology. Here the user interacts with a virtual model of the graph via an end-effector, such as a mouse or pen. As the user moves the end-effector around, its position is tracked and compared to the virtual graph model. Resistance via motors is then provided to create the illusion of touching a physical object. Fritz and Barner [9] created an automatic graph construction tool where a user could input a function via the keyboard and the resulting graph was physically “carved” out of a virtual block which could be explored using a SensAble PHANTOM haptic device (www.sensable.com). Yu and Brewster [23] carried out several studies and developed a number of guidelines for designing both line graphs and bar charts for exploration via a PHANTOM. Their work showed that users could effectively answer simple questions with the PHANTOM and that accuracy was significantly higher than raised paper graphs. However, the time taken to complete the tasks was found to be significantly greater. Additionally, users became confused at intersection points in line graphs where two data series crossed and would often unknowingly switch between data series at these points. Yu and Brewster [23] identified no effective way to deal with this problem. More recent work has looked at allowing users to manipulate the graphs. McGookin and Brewster [15] adapted Yu’s and Brewster’s [23] technique to allow users to drag bars in a bar graph up and down and thus allow construction. Bernareggi *et al.* [2] have started to develop a less structured graph construction system, allowing users to place control points and connect them to form graph features via a PHANTOM device.

One limitation of this haptic technology however, is its single point of contact nature. The user interacting with the graph does so only at one position at a time and cannot easily compare different parts of the graph spatially. Rather, he or she must try to remember the relevant information or try to relocate it, losing the current position in the graph as this is done. Lederman and Klatzky [11] have studied how users haptically interact with physical objects. They identified several Exploratory Procedures that were used by par-

ticipants to determine haptic properties such as size, shape, weight, etc. Most of these involved two hands (e.g. contour following) or stimulation of tactile receptors through the skin (e.g. roughness) and are not supported by single point of contact haptic devices. As graph interpretation relies heavily on all relevant information being concurrently held in short term memory (which degrades rapidly over time) [12], this can limit the ability to answer more complex questions, such as those that require comparison between different bars in a bar graph. These limitations mean that any virtual haptic system is severely impoverished in comparison to the paper-based techniques previously discussed.

To ameliorate these issues, several researchers have proposed ways to augment virtual haptic graphs. McGookin and Brewster [16] provided quick overviews of bar graphs by providing a separate auditory view below the x-axis. Their Sound-Bar (see Figure 2) improved accuracy where multiple bars had to be compared, but was only useful for certain types of questions and added complexity to the interaction. Wall and Brewster [21] augmented the haptic graph of Yu and Brewster [23] with “beacons”, small markers that could be spatially set and used to return to previous locations. However, they found that the beacons were often not used and were found to subjectively increase demands on memory, as participants had to remember where a beacon had been placed.

The limited usefulness and relative increase in complexity that overcoming the problems of single point of contact haptic devices requires, has led to the development of systems that incorporate a static tactile guide to aid exploration. Wall and Brewster [22] allowed access to pie charts via a standard graphics tablet. A compact disc (CD) attached to the tablet represented the pie chart. As the user moved the tablet pen around the edge of the CD, the segments of the pie chart were sonified using a pitch value mapping. The tactile chart outline allowed the user to employ his or her other hand and to mark segments for easy return and comparison. In an evaluation with visually impaired users, the tactile element of the discs afforded better orientation within the graph.

Overall, this prior work has shown that the disadvantages of existing tactile graph access techniques, that they are cumbersome and their inability to be stored, can be overcome by non-speech audio and virtual haptic technologies. However, what is also clear, is that many of the advantages of physical tactile graph access, two handed interaction, quick overviews, spatial frame of reference and flexibility to employ fingers for marking, have been lost in trying to overcome the problems. The more recent work of Wall and Brewster [22] has looked at trying to augment tactile diagrams. However, this work still constrains the user to interact with the computer based information via a single point of contact, limiting the available interactions and potentially overloading short term memory. We propose an alternative approach that builds directly from existing tangible technologies whilst incorporating support for computer based data.

TANGIBLE INTERACTION

Although the common systems in use for graph construction and browsing use physical components, the work of tangible

user interfaces (TUIs) as popularised by Ullmer and Ishii [20] has not yet been applied to this area. Whilst TUIs have become more popular in the last decade, there remains significant disagreement as to the definition of a TUI, with some researchers arguing that conventional mice are examples [7]. In this paper we consider the definition of Ullmer and Ishii [20]: “*user interfaces employing physical objects, instruments, surfaces, and spaces as physical interfaces to digital information.*”. Users primarily interact with such systems by manipulating real world objects called physical icons (phicons) which are tracked by computer systems causing physical manipulation to change digital state. We employ this definition as a means of distinguishing between tabletop systems and the virtual haptic and audio solutions previously discussed. Whilst there are many examples of such interfaces, there is little evaluation of their usefulness [14]. There are, however, several reasons why we believe that TUIs are a suitable and promising approach for non-visual graph access. Sharlin *et al.* [18] note that effective tangible user interfaces must incorporate good mappings of spatial information to digital artefacts. One aspect of a successful mapping they argue, is that the cost of changing or manipulating the state of the tangible user interface should be low. That is, it should support trial and error activity, much like the construction and manipulation of graphs previously discussed. Antle, Droumeva and Ha [1] compared children’s performance in completing a jigsaw puzzle when interacting with a tangible system on a digital table and a virtual jigsaw controlled with a mouse. They found that time taken was lower and completion rates were higher with the TUI than with the mouse. On analysis, they determined that this was due to the tangible elements of the task leading to better spatial model construction which is again important in graph reasoning [12].

Whilst TUIs appear to be promising, there is no research that investigates their use for people with visual impairments or the use of tabletop tangible interfaces without vision. We do not know how such interfaces should be designed and there are several key research questions that must be addressed if TUIs are to be used in non-visual scenarios. Can such interfaces be designed to support visually impaired users? How should phicons (physical icons) be designed to allow non-visual use? How should functionality be split between tangible and non-tangible elements of the interface?

TANGIBLE GRAPH BUILDER

To investigate these issues we developed a tabletop TUI system that allows users to browse and construct both line and bar graphs non-visually. Due to the lack of existing non-visual tangible guidelines, we chose as a starting point, to base our system on the corkboard creation technique. This has the advantage of being the most unconstrained technique available, being suitable for the creation of many types of graph. The grid system it employs also forms the basis of fundamental mathematical knowledge of 2D space which is important to understand for mobility in everyday life such as accessing map based information. During development we employed guidelines generated from tactile diagram research or virtual non-visual graph-based systems wherever these seemed appropriate. The following is a discussion of

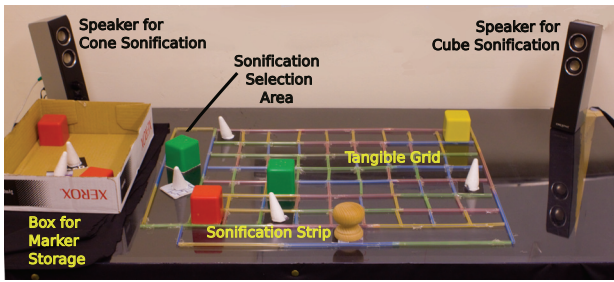


Figure 3. The setup used in Tangible Graph Builder showing the grid, sonification strip and sonification selection area.

the line graph system, we outline the differences between it and the bar graph system later.

A (9 x 7) tangible grid was constructed from drinking straws and attached to a clear perspex topped table. As the grid would not be something a user would want to change during construction or browsing, we followed the guidance of Challis and Edwards [5] in permanently affixing the grid to the table so that it could not be moved. We replaced the pins of the corkboard technique with phicons. Two different shapes (a cube and a cone) were used to represent phicons for two data series (see Figure 5). Each cube was filled with plasticine so that it weighed approximately 110g, in comparison to the polystyrene cones which weighed 10g. These provide distinctly different textures, shapes and weights that should make discrimination between the phicons easier. A 4 x 4 cm cardboard square was attached to the base of each phicon allowing it to be snugly held within the grid.

In the corkboard graph construction technique, connecting the control points (pins) is very much “joining up dots”, but can cause a great deal of problems in ensuring the correct pins are being connected, or when removing the bands to change the position of the pins later. In our system the job of connecting the pins was automatically carried out by the computer. To allow for tracking of the phicons we employed ARToolkit (www.artoolkit.org). ARToolkit tracks fiducial markers (see Figure 5) in 3D space with a camera, and more commonly is used to render 3D graphical objects on top of these markers. Custom software was written in C# to track the markers and determine where in the graph grid they were located. The markers were tracked using a Logitech Webcam Pro 9000. We attached the markers to the cardboard base of the phicons, with the camera placed on the floor under the table. The system maintains a record of where on the table the markers are and what data series they belong to. Therefore a model of the graph can be maintained.

Based on previous research carried out on both haptic graph browsing [16] and non-visual sonification of linear data series, we incorporated a sonification strip into the table (see Figure 4). This strip runs along the base of the x-axis and can be controlled using a special phicon (see Figure 5). As the user drags the phicon along the strip and it passes between the major units of the x-axis, the system calculates the appropriate y-value of the data series at that point and converts it to a Musical Instrument Digital Interface (MIDI) pitch value based on the formula provided by Brown and Brewster [4].

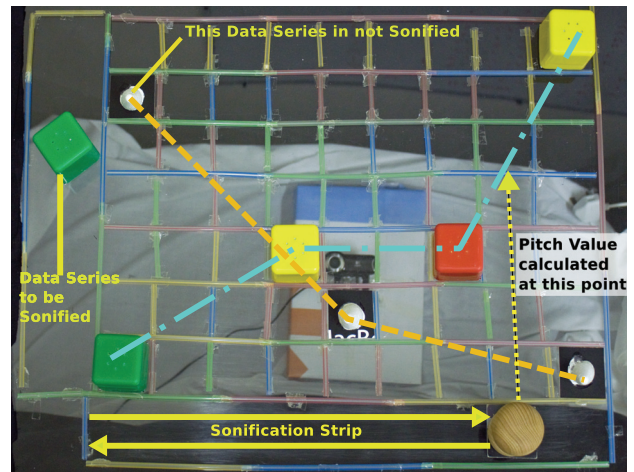


Figure 4. An annotated illustration of an interaction with Tangible Graph Builder.

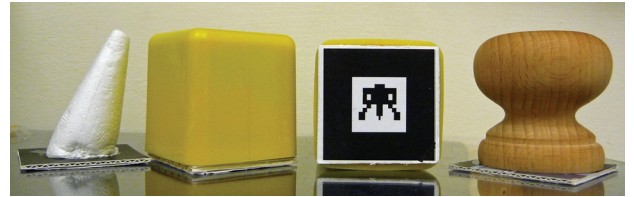


Figure 5. The phicons used in Tangible Graph Builder for the data series, a fiducial marker used to track the phicons and the phicon used for interaction with the sonification strip.

The notes for each data series were played using a piano timbre (General MIDI patch number 000) and each data series was spatially panned to the left or right stereo channels to improve separation (see Figure 3). As noted by Brown and Brewster [4], when identifying crossing points between two data series accuracy is improved if both are concurrently presented. Conversely, when identifying turning points or gradient of a data series, accuracy is improved if only one is presented. To allow for both options, we introduced a further area to the left of the y-axis. Placing a phicon from one or both of the data series in this area caused the corresponding data series to be sonified when the sonification strip phicon was moved. In this way the user can control which data series are played. The sonification strip and the control area were physically demarked on the table in the same way as the tangible grid (see Figure 4).

The bar graph version of the application was similar to the line graph version; Each column of the grid was treated as a separate bar, a phicon in that column at any row position was treated as the top of the bar. We changed the sonification strip slightly to include a distinct tone played with a synthesised drum pad (GM patch number 090) to indicate if the column a user had moved into was set with value 0 (indicated by no phicon in that column). In addition, the bar graph version only supported one data series (each column could have only one phicon at a time), therefore the sonification selection area was not used and the user only had to move the sonification strip phicon to hear the graph. Because of this, both sets of phicons could be used interchangeably.

EVALUATION

We carried out a two part study on Tangible Graph Builder. The main aim was to identify the usefulness of TUIs for visually impaired users by providing answers to the set of research questions previously outlined. The first part involved sighted users who were sight deprived, whilst the second involved a group of blind users. Due to the current problems of graph access technologies, graph knowledge amongst visually impaired users can be variable. As a fairly small group of visually impaired users was available to us, we wanted to be able to compare the results back to a group of users with consistent graph understanding. Hence we compared to a group of sight deprived students with at least a high school level of graph knowledge.

Sight Deprived Study

Eight sighted (5 men, 4 women) participants performed tasks using Tangible Graph Builder. All reported normal hearing and normal or corrected to normal vision. Participants performed four types of task. All of the tasks were typical of those that might be carried out in graph work at school in the U.K by students in the 14-16 age range [3]. The questions were based on those from existing work evaluating virtual haptic and auditory graph systems [4, 15, 16, 21] and allowed flexible use of the sound, referral to the phicons, or a combination of both. The four task types were:

Construct line graphs with two data series: A printed table of control points for two data series was supplied. Participants were asked to reproduce the resulting line graphs using Tangible Graph Builder.

Browse line graphs with two data series: A line graph with two data series was presented with Tangible Graph Builder. Participants had to explore the graph and identify the crossing points between the two data series, as well as the number of times each data series changed direction, either from a positive to negative gradient, or a negative to positive gradient. After these questions had been answered the participants were asked to sketch the graph on a paper grid with the same number of rows and columns as the tangible grid.

Browse bar charts: Participants were given a pre-built bar chart containing 12 bars and had to answer a question about the graph. E.g. “What three bars have the highest values?” or “From bars 1,3 and 5 which has the lowest value?”. The questions and stimuli for this task type were taken from McGookin and Brewster [16].

Construct bar charts: Participants were given a printed table with values for six bars and asked to construct a bar chart. Participants had to scale the y-axis values in order to fit the chart into the tangible grid.

Each participant was given a demonstration of Tangible Graph Builder before commencing the tasks. During this demonstration the participant was allowed to look at the tangible grid and the phicons placed on top of the table. However, during the experimental tasks a black cloth screen was erected between the participant and the table. Participant’s interactions with the table were video taped and their answers

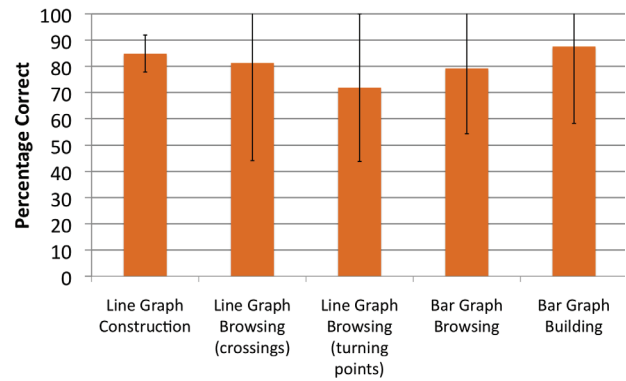


Figure 6. A graph showing the accuracy of the sight deprived users carrying out the tasks. Accuracy is presented as a percentage of total possible score for each category. Shown with standard deviations.

to the questions recorded. After completing all tasks, participants were interviewed about their experiences using Tangible Graph Builder, with emphasis on the research questions previously outlined. These data (video recordings, quantitative results, interview transcripts and experimenter notes) were analysed based on a framework analysis approach [17], using the research questions as initial topics, but allowing other topics to emerge. In the following sections we briefly discuss quantitative performance and the major topics to emerge from the framework analysis.

Results

The graph shown in Figure 6 illustrates the accuracy of users on each of the tasks. Due to participants carrying out a different number of trials for each graph type and task, we express accuracy as a percentage of the total possible score for each task and graph type. The results obtained are comparable to performance by blindfolded sighted users when performing similar tasks in virtual haptic and auditory graph access software [4, 15, 16].

Phicon Design

Of the two phicon types used to represent data points (cubes and cones), seven of the eight participants expressed a preference for the cube phicons. Several reasons for this were stated, but most commonly was the likelihood of knocking over the lighter cone phicons. As one participant stated “If you go on the graph and start feeling around, that’s when the cones can travel”. Another participant felt that the cubes “were better objects to identify with than the other ones (cones). The other ones were like papers and I missed them out. I think the cubes were just bolder. At times I moved my hands over them (cones) and I just moved them. The others were a bit difficult to do that to because they were um, stronger. I think I must have knocked one or two (of the cones) over”. The subjective views of participants were confirmed by the quantitative results and video recordings. Out of the seventy two trials that were performed, there were nine occasions where a cone phicon was dislodged from the grid, compared to one occasion where a cube phicon was dislodged. Additionally, all of these occurred when the user was accessing or constructing a line graph. This may be partly due to the increased use of sound in bar graph questions (see section on exploratory strategies), or that bar gra-

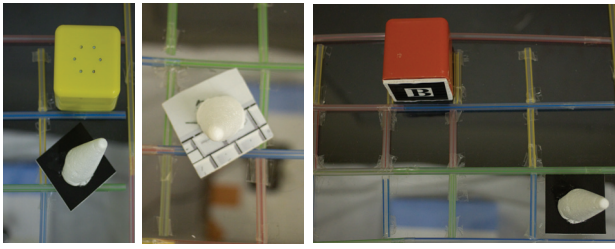


Figure 7. Cone phicons were often dislodged so that they were lost by the camera tracker but not noticed by participants (left and centre). Because of the regular shape, the cube phicons were sometimes placed wrong side down (right).

phs are more predictable in layout (with only one marker in each column). Although participants identified when phicons had been knocked over, it was not immediately recognised if they were only dislodged, even if this was sufficient for the marker to be lost by the tracking system (see Figure 7). Where the phicon had been knocked over, the participant immediately identified this and replaced it in the closest square, even if this was not the grid square where the phicon was originally located. Although the cubes were preferred by almost all of the participants, there were issues. Because the cubes were regular, during the graph construction tasks participants sometimes put the cubes down so that the fiducial marker was not on the table top (see Figure 7). This happened on two occasions but was not picked up by participants, primarily due to the lack of sound usage in the construction tasks (see section on exploratory strategies). However in cases where the user must browse and modify a pre-existing graph, this is likely to be a greater issue.

Exploratory Strategies

All of the participants used a two handed exploration strategy to complete all tasks. The way in which both hands were used varied between the browsing and construction tasks, as did the overall use of the sonification strip, but the strategies broadly followed the techniques observed in our earlier corkboard construction observation.

In construction tasks two main strategies emerged. In the first, which was used by half of the participants, the users would move relatively. The non-dominant hand used the previously positioned phicon as a reference point, whilst the user referred to the printed sheet to identify the next phicon and calculate the relative position of that phicon. He or she would then hold the phicon in the dominant hand and count from the non-dominant hand along the x-axis and then up or down the y-axis to locate the appropriate position. The non-dominant hand then met the dominant hand to mark position and the participant moved onto the next phicon. In the other strategy the participant used two hands to count from the origin along the x-axis, holding the phicon in the dominant hand. The x value was marked with the non-dominant hand and then users counted up with the dominant hand containing the phicon to place it in the correct position in the tangible grid.

In the browsing tasks a number of different strategies emerged. Participants tended to move between them depending on the task. In addition, the sonification strip was more

widely used. There were only 4 out of 24 construction trials where the sonification strip was used, compared to 27 out of 48 browsing trials. In all of the construction trials the participants reported that they used the sonification at the end of the task to confirm that the graph had been correctly created. Arguably the sonification strip does not add useful functionality in purely construction tasks. In the line graph browsing tasks the use of the sonification strip formed part of richer strategies. The sonification strip was more commonly used when accessing line graphs (used in 14 of 16 trials). The ways in which it was used agree with the results of Brown and Brewster [4] and their evaluation of the SoundVis system previously discussed. When trying to find intersection points, participants would sonify both data series together. When trying to identify features of a single data series participants would explore each individually. However, participants often would not be able to make a decision on crossing points or turning points purely from the sound. In such cases they would move along the sonification strip until reaching an area where a crossing might occur and then moving both hands up the graph from that point to explore the phicons. Users would either mark a phicon with the left hand and then find the next phicon of the same data series with the right hand, or if the user was looking for intersection points, mark a phicon with one hand and feel around for phicons of the other data series with the other hand.

Browsing bar graphs, due to their more regular structure, produced more polarised strategies. Participants either used the sonification strip (4 participants), or used a two handed strategy to browse and mark bars (4 participants). There were very few cases where a participant used both techniques on a trial. There was also variation within participants. Notably one participant who did not attempt to use any sound when browsing the more complex line graphs, immediately started using the sonification strip, and only the sonification strip, when asked to browse bar graphs. Participants were unable to fully explain the reasons for this. When using the sonification strip participants would usually use fingers on the non-dominant hand to mark candidate bars (e.g. when trying to find the highest) by touching the base of the grid at the appropriate x position. When using the phicons, participants followed a variant strategy, touching a phicon with the non-dominant hand and spatially comparing it to the next phicon, or using the non-dominant hand to mark candidate bars by touching the phicons.

Division of Functionality

An important aspect of any tangible user interface is how to divide the system functionality between phicons that the user can manipulate, and state information that should be communicated by the system to the user either through visual (usually projection), audio or physical (either by moving the phicons in space or changing their physical aspect) means. In our system we have three distinct categories: the tangible grid (which was fixed and non-manipulable by the user), phicons (which could be manipulated by the user) and the line series (which changed as a consequence of phicon manipulation). We chose this division based on the work of Challis and Edwards [5], who in the design of a non-visual musical score browser recommended that fixed features, which could

not be altered by the user, should be represented in a tangible form, and other features should be represented by sound. We asked all participants about the division of functionality and if they felt this was appropriate. No participants raised issues with the division and no relevant events from the user interaction were identified. We believe therefore that the division of functionality was appropriate.

Tracking Accuracy

Tangible Graph Builder used a fiducial visual camera based tracking system to monitor phicon position. This was primarily motivated to allow for rapid development of the initial system. Modifications during the initial development (wrapping the table in dark material, dimming the room lighting and illuminating the table from below) improved tracking accuracy, but there were still instances during the study where markers would either fail to be detected, or be intermittently detected by the system. Whilst the experimenter stepped in during prolonged periods of intermittent failure and where markers had been totally lost, this was not done initially to observe if users detected the marker loss and how they dealt with it. There were two types of marker loss that were relevant: failure of a marker in the grid and failure of a marker in the sonification area. Where a marker detection failure occurred in the sonification selection area, participants became aware of this quickly due to the lack of any sound output from the system and took remedial action (either twisting or “jiggling” the phicon to try to force detection). The participants acquired these strategies from the experimenter during the initial (sighted) familiarisation phase. However, in cases where the markers in the grid failed to be detected, or were intermittently detected, the system redrew the graph which caused the sonification to change. In such cases the participants attempted to take no corrective action in spite of the sonification strip providing highly inconsistent sounds. In the post study interviews where intermittent detection was most prevalent, participants expressed their lack of confidence in the sonification, their answers and the reliability of the sound. As one participant said of using the sonification strip to browse a line graph: *“but then it went screwy, I don’t think I heard the same thing twice”*. In visual tangible interfaces it is generally straightforward to communicate loss of detection, but non-visually it is much harder. We discuss this further in the future work section.

Another issue surrounding tracking accuracy occurred with marker “jitter”: cases where the detected position of the marker varied between frames of the camera. This usually caused no problems, as the marker would jitter within the tangible grid square it was placed. We deliberately chose each grid square to be 4 x 4cm to ensure that the jitter would not be an issue. However, this still arose with the sonification strip phicon. The sonification strip would play the musical note that represented the current y value of the sonified data series when it moved between x-axis grid squares. Users would often move the sonification phicon slowly to hear and count each note (particularly in the bar graph browsing tasks). If they paused near the transition between grid squares, the marker could jitter to each side, playing a note each time. Participants coped with this by using either a finger on the same hand controlling the sonification phicon to

count the tangible grid, or used their other hand to accomplish the same task. One participant described this issues as it related to browsing a bar graph: *“I knew each sound corresponded to a bar, but sometimes I would listen to two sounds, two pitches that played very fast. So in order to increase accuracy and assign a pitch to a bar, I was counting as I was moving a bar (sonification phicon)”*. Participants also handled this by moving the phicon back to the origin and trying again.

Whilst both of these issues may be reduced or eliminated by more accurate tracking technology, they provide useful insight into the issues of marker detection and the coping strategies employed.

Blind User Study

To confirm the results obtained, we carried out the study again with four blind users (2 men and 2 women) with no residual sight. Two were congenitally blind, whilst the other two were late blind and had been sighted when taught graphs in school. Participants completed the same study as the sight deprived group with the following variations. Participants were initially introduced to the system by free exploration. Participants explored the table and had features explained as and when they came into contact. As the participants were blind rather than visually impaired, the black screen between participant and table was not used. Any printed materials were read out by the experimenter on the request of the participant. As there was no effective way to reproduce the line graphs on paper, participants were only asked for the turning and crossing points in the line graph browsing tasks. Summary accuracy results are shown in Figure 8. The results show overall good performance in all tasks, with the exception of the identification of turning points in the line graph browsing tasks. Reasons for this are discussed in the following sections. In other areas performance exceeded that of the sight deprived group, with no errors at all in the bar graph construction tasks. We noted no significant differences between the congenitally blind and late blind participants.

Usability Problems with Tangible Graph Builder

Strategies employed to explore the graph were similar to the sight deprived group, with participants using two handed exploration strategies with one hand to manipulate the graph and the other to mark context within the graph. In browsing tasks participants again used both hands to explore and mark features in the graph. When carrying out the line graph browsing tasks the distribution of strategies changed. Two of the participants did not use the sonification strip at all, whilst the other two primarily used the sonification strip with a small amount of physical exploration with the phicons. When asked, the participants who had only used the phicons for exploration said that they had been trained on using touch and that the tactile was obvious. However both immediately opted for sound when browsing the bar graphs.

The use of two hands was mentioned as an important aid to completing the tasks, that at times made it unnecessary to try to use the sonification strip: *“In terms of the question when it was a case of how many times the item (data series) cross over each other. I found it easier (using tactile cues only).*

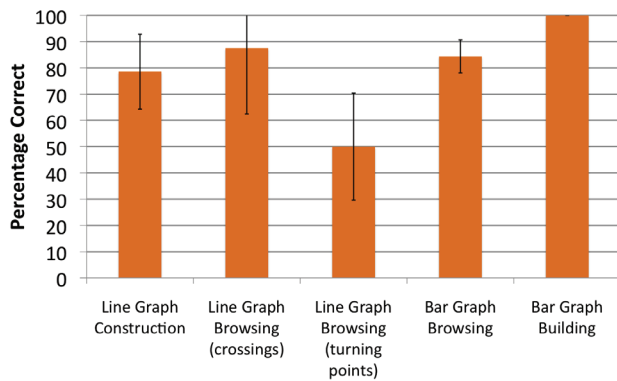


Figure 8. A graph showing the accuracy of the blind users carrying out the tasks. Accuracy is presented as a percentage of total possible score for each category. Shown with standard deviations.

That may have been partly due to not knowing if I got reliable feedback or not. I trusted my mental image with what I felt with my hands". This led to two of the participants to suggest that the audio feedback be simplified in the line graph system and rather than provide a pitch based mapping of the current value of the line series, it should provide only a relative value of the two data series (e.g. a high pitch if one data series was above the other and a low pitch if the reverse). However, two participants, those that had used the sonification strip extensively during the line graph browsing tasks, felt that the sounds were fine if the marker tracking could be made more accurate.

The issues identified with the phicons by the sight deprived group were also present. However, there were less occasions when phicons were knocked or dislodged by participants. On four trials a participant dislodged a phicon from the grid. On all occasions this was a cone phicon which the participant identified but replaced in the wrong grid square. All of the participants expressed that the cube phicons were preferred, as the cones were too light. As one participant said: *"The blocks are nice and solid, the cones being lighter are easier to tip over and knock out of the slot they were in."*. Another participant mentioned that: *"In a thing this size you tend to have a sweep around so you need something fairly solid."*. Overall, one participant commented on the phicons as being: *"nice that they just snug fit. I think you want something that just fits really snugly."*. In addition to phicons being knocked, there was one occasion where the participant put a cube phicon down with the marker facing upwards requiring intervention by the experimenter to correct.

In addition to soliciting comments about issues with Tangible Graph Builder, we explicitly asked participants if there were any features that could be added which would make completing the tasks easier. Primarily comments addressed contextual feedback. Participants mentioned that especially in large grids, the ability to query what the value of a grid square was, or the current position of a phicon, would be important to complete the tasks. One participant mentioned: *"If you had a third type of marker (phicon) that you could use in a similar way, if you wanted the reading at a particular point, you could put the relevant marker on to get*

a reading of the position.". Another participant mentioned that tapping a phicon on the table could be used as a means of triggering spoken feedback on its position.

Comparison with Existing Systems

In addition to confirming the results from the sight deprived users, we wanted to gain more qualitative feedback on how the use of "real" tangible user interfaces compares to virtual haptic feedback (such as via the PHANTOM force feedback device) and a Mansur style soundgraph [13]. To this end, after participants had completed the tasks with Tangible Graph Builder and their views had been elicited, we asked them to complete bar and line graph browsing, as well as bar graph construction tasks using two previously written and evaluated systems (discussed in the following sections). The tasks that we asked participants to complete were of the same complexity as those carried out with Tangible Graph Builder. However, we excluded the line graph construction tasks as the soundgraph software (as previously discussed) only supported graph browsing and not graph manipulation. Data collection was based on interviews conducted after all systems had been used.

SoundVis

SoundVis, as evaluated by Brown and Brewster [4], provides access to soundgraphs via the numeric keypad of a standard computer keyboard. Users use the keypad to move left and right in the graph. SoundVis can sonify up to two data series at a time and users can switch between serial and parallel presentation of two data series, as well as switch between the data series using other keys on the keypad. Each data series is rendered using a general MIDI piano timbre (GM Patch 000) using a pitch value mapping and stereo panned to a different speaker.

SoundBar Builder

SoundBar Builder combines two evaluated techniques [15, 16] to allow construction and overviews of bar graphs. Bars are modelled as recessed grooves that can be explored using the PHANTOM. Below the bars and the x-axis is a SoundBar. This acts in the same way as the sonification strip from Tangible Graph Builder; as the user moves the PHANTOM pen along the strip, the bar immediately above is sonified using the same pitch mapping as the sonification strip. SoundBar Builder makes extensive use of speech, primarily to stop users from getting lost. Touching any feature in the graph and then pressing the button on the PHANTOM pen, yields speech feedback on that feature (the bar index, axis name, etc.). A screenshot of the graph model is shown in Figure 2.

Results and Discussion

The strategies used to explore the graphs in SoundBar Builder and SoundVis were the same as those that have been previously identified in the individual evaluations of these systems [4, 16]. Due to space constraints we concentrate here on the qualitative differences between the three approaches.

All of the participants said that Tangible Graph Builder provided a useful two handed interaction: *"More than one point of contact gives you the ability to see where things are relative to each other which you don't get with the PHANTOM. You don't just touch with one nerve, you touch with*

all nerves, shape, texture etc.”. All participants expressed that when exploring a graph (as was exhibited in Tangible Graph Builder) they would try to obtain an overview of the area first: “The PHANTOM does give you a tangible thing, its probably slightly more work in the mental memory processes to remember how things are relatively. You have to get into the bar and find out if it is higher or lower than the next one. There is quite a methodical process required there. The tactile (Tangible Graph Builder) gives you the ability to fairly quickly move and get an indication of where the lines were.”. When using the PHANTOM version participants could easily become disorientated: “In the tangible one, it is easier for me to locate things in space. In the PHANTOM one I sensed that the degree ... I would be more likely to make a mistake trying to find the points”.

Another important topic to emerge from user discussions was the confidence that participants felt about the system and as such their answers. All of the users felt the PHANTOM to be more reliable than Tangible Graph Builder, primarily down to issues of marker loss. However the combination of different ways of accessing the graph in both SoundBar Builder and Tangible Graph Builder was an important factor in improving confidence: *“The good thing with the tangible one and the PHANTOM one, is that you are integrating senses. That allows you to rely more in your answers.”. Two of the participants felt that SoundVis provided a more accurate representation of the change of shape of a line graph and made identifying turning points easier: “You do get a feel for the shape of the curve (with SoundVis) ... which you don’t get very easily off that (Tangible Graph Builder)”.* However, as SoundVis contained a greater number of divisions on the x-axis (~100) in comparison to the 9 divisions on the Tangible Graph Builder sonification strip (one for each column in the tangible grid), this may be able to be rectified with more units on the sonification strip.

GUIDELINES FOR NON-VISUAL TUIS

Our study has yielded useful information that would be beneficial to future designers of tangible user interfaces for non-visual use. In this section we extract and discuss guidelines that we believe future designers should consider.

Phicons Should be Physically Stable

We originally designed our phicons to be haptically different, so that they could be quickly discriminated using as many Exploratory Procedures as possible [11]. However, the cone phicons were more often knocked over and least preferred by the participants. Phicons should therefore be hard to accidentally move. Varying the weight, as shown by the preference for the cube phicons, is an effective way to do this. The cube phicons at 110g are a good starting point.

Phicons Should have Irregular Forms

Although the cube phicons were preferred, there were occasions, due to their regular shape, where they were placed with the fiducial marker on an incorrect side. Phicons should be irregular so they can only be placed on the table one way. An effective means of doing this, as was suggested by one of the visually impaired users, may be to attach embossed shapes on top of the phicons so “up” can be determined.

Divide Functionality Appropriately

In any non-visual tangible user interface, there are three types of data. Data which are fixed (or very infrequently changed) through the use of the system (e.g. the grid the user constructs the graph on), data which are frequently and directly changed by the user (e.g. the position of data points on the grid) and data which are frequently and indirectly changed by the user (the relationship between consecutive phicons in the grid). In extension of the guideline of Challis and Edwards [5], fixed information should be represented by immovable physical objects, directly manipulated data should be represented by phicons, whilst indirectly changed data should be presented via sound or tactile stimuli. Participants were explicitly invited to comment on this issue during the interview phase and all felt that the distinction employed did not adversely affect their performance, therefore we believe that it is an appropriate way to divide functionality.

Provide Awareness of Phicon Status

In any tangible system phicons can fail to be detected. Even if the system is reliable, the user may put a phicon in the wrong place, or in a way that means it is not detected. In cases where intermittent phicon detection failure caused the sonification to change, users were unsure of their answers and had no way to test the detection status of a phicon. Whilst it is relatively easy to indicate that a phicon is not detected via visual means, it is harder to do so non-visually in a way that does not annoy the user. However, such awareness is important and should be provided.

FUTURE WORK

Tangible Graph Builder has provided much useful guidance in developing non-visual tangible user interfaces, however there are several aspects of its design that still require development. From the results of both evaluations we need to improve the tracking detection of our system as well as provide information about the status of each phicon. We propose to insert a small Arduino (www.arduino.cc) - a programmable microcontroller - into each phicon. Via a bluetooth connection, our system will be able to communicate with the phicon and inform it if it is, or is not, being detected. The microcontroller can then provide feedback to the user. For example, each phicon could play a short musical sound when it was detected and lost by the camera. If either sound was heard when the phicon had not been explicitly moved, the user would know of a problem. To aid finding the problematic phicon a peltier heat pump (an electronic component that can heat up or cool down via an electrical current) could be employed so phicons that were not detected could be cooled. This would allow feedback to be provided in an unobtrusive yet useful way. An additional aspect is to confirm our guidelines in other tangible scenarios. Overviews of maps and geo-data present similar problems for visually impaired users, indeed a grid is the basis of most map systems, and we plan to investigate the role of TUIs there.

CONCLUSIONS

Accessing graphs and charts with a visual impairment not only presents problems due to the translation of a visual representation, but also the impoverished technologies that are available to access and manipulate the graph. Tangible

Graph Builder presents a first step at augmenting those technologies to be more useable, whilst incorporating the benefits of computer-based support. In doing so, we have moved away from current research approaches, such as virtual haptics and sonifications controlled via a keyboard which, whilst overcoming many of the problems of the paper-based and corkboard techniques, restrict the exploration strategies users would naturally employ. Our evaluation with both sight deprived and blind users has shown the benefits this approach can bring. Participants were able to complete graph construction and browsing tasks by employing a range of two-handed exploration strategies. We believe the application of TUI technology to provide non-visual access to graphs is a valuable approach, that will empower users to access information which has up to now been largely inaccessible.

ACKNOWLEDGEMENTS

The authors wish to thank James McGookin for constructing the table and technical assistance. This work is supported by the EU FP7 No.224675 “Haptimap” project.

REFERENCES

1. Antle, A., Droumeva, M., and Ha, D. Hands on what?: comparing children’s mouse-based and tangible-based interaction. In *IDCS’09* (Como, Italy, 2009), vol. 1, ACM Press, pp. 80–88.
2. Bernareggi, C., Comaschi, C., Marcante, A., Mussio, P., Parasiliti Provenza, L., and Vanzi, S. A multimodal interactive system to create and explore graph structures. In *CHI’08* (Florence, Italy, 2008), vol. 2, ACM, pp. 2697–2702.
3. Brown, D., Howat, R. D., Meikle, G., Mullan, E. C. K., Murray, R., and Nisbet, K. *New Mathematics in Action S32*. Nelson Thornes, Cheltenham, 2004.
4. Brown, L., and Brewster, S. A. Drawing by ear: Interpreting sonified line graphs. In *ICAD 2003* (Boston, USA, 2003), vol. 1, ICAD, pp. 152–156.
5. Challis, B. P., and Edwards, A. D. N. Design principles for tactile interaction. *Haptic Human-Computer Interaction LNCS 2058* (2001), 17–24.
6. Clamp, S. Mathematics. In *Visual Impairment: Access to Education for Children and Young People*, H. Mason, S. McCall, C. Arter, M. McLinden, and J. Stone, Eds. David Fulton Publishers, London, 1997, pp. 218–235.
7. Fishkin, K. P. A taxonomy for and analysis of tangible interfaces. *Personal and Ubiquitous Computing 8* (2004), 347–358.
8. Flowers, J. H., and Hauer, T. A. Musical versus visual graphs: Cross-modal equivalence in perception of time series data. *Human Factors 37*, 3 (1995), 553–569.
9. Fritz, J. P., and Barner, K. E. Design of a haptic data visualization system for people with visual impairments. *IEEE Transactions on Rehabilitation Engineering 7*, 3 (1999), 372–384.
10. Jones, R., and Careras, I. The empirical investigation of factors affecting graphical visualization. *Behavior Research Methods, Instruments & Computers 22* (1996), 643–655.
11. Lederman, S. J., and Klatzky, R. J. Hand movements: A window into haptic object recognition. *Cognitive Psychology 19* (1987), 342–368.
12. Lohse, G. L. Models of graphical perception. In *Handbook of Human-Computer Interaction*, M. Helander, T. Landauer, and P. Prabu, Eds., 2nd ed. Elsevier, Amsterdam, 1997, pp. 107–135.
13. Mansur, D. L., Blattner, M., and Joy, K. Sound-graphs: A numerical data analysis method for the blind. *Journal of Medical Systems 9* (1985), 163–174.
14. Marshall, P., Rogers, Y., and Hornecker, E. Are tangible interfaces really any better than other kinds of interfaces. In *Workshop on Tangible User Interfaces in Context and Theory* (San Jose, USA, 2007), vol. 1.
15. McGookin, D. K., and Brewster, S. A. Graphbuilder: Constructing non-visual visualizations. In *BCS-HCI 2006* (London, UK, 2006), vol. 1, Springer, pp. 263–278.
16. McGookin, D. K., and Brewster, S. A. Soundbar: Exploiting multiple views in multimodal graph browsing. In *Nordichi 2006* (Oslo, Norway, 2006), vol. 1, ACM press, pp. 145–154.
17. Ritchie, J., and Spencer, L. Qualitative data analysis for applied policy research. In *Analysing qualitative data*, A. Bryman and R. Burgess, Eds. Routledge, London, 1993, pp. 173–194.
18. Sharlin, E., Watson, B., Kitamura, Y., Kishino, F., and Itoh, Y. On tangible user interfaces, humans and spatiality. *Personal and Ubiquitous Computing 8* (2004), 338–346.
19. Smith, D. R., and Walker, B. N. Tick-marks, axes, and labels: The effects of adding context to auditory graphs. In *ICAD 2002* (Kyoto, Japan, 2002), vol. 1, ICAD, pp. 362–367.
20. Ullmer, B., and Ishii, H. the metadesk: Models and prototypes for tangible user interfaces. In *UIST 97* (Banff, Canada, 1997), vol. 1, ACM Press, pp. 223–232.
21. Wall, S. A., and Brewster, S. Providing external memory aids in haptic visualisations for blind users. *International Journal of Disability and Human Development 4*, 4 (2004), 331–338.
22. Wall, S. A., and Brewster, S. A. Tac-tiles: Multimodal pie charts for visually impaired users. In *Nordichi 2006* (Oslo, Norway, 2006), vol. 1, ACM Press, pp. 9–18.
23. Yu, W., and Brewster, S. A. Multimodal virtual reality versus printed medium in visualization for blind people. In *Assets 2002* (2002), vol. 1, ACM Press, pp. 57–64.