New Views on Haptic Graph Visualisation

David K. McGookin, Johan Kildal, Stephen A. Brewster
The MultiVis Project
Glasgow Interactive Systems Group
Department of Computing Science
University of Glasgow
G12 8QQ
{mcgookdk, johank, stephen}@dcs.gla.ac.uk
http://www.multivis.org

ABSTRACT

This paper proposes that current approaches to haptic graph visualisation (by simply presenting a simple haptic equivalent of a visual graph) are inadequate. We present a short background to graph theory perception and identify the problems of current graph haptic visualisation. We then propose ideas influenced by both information visualisation research and Lederman and Klatzky's EPs in order to overcome the problems of graph perception described. We then identify the research questions that must be answered about our technique before concluding.

Keywords

Haptilisation, Haptics, Force Feedback, Graph Perception, Information Visualisation

INTRODUCTION

The MultiVis project is investigating how to make simple mathematical visualisations (such as graphs and tables) accessible to those who are blind or visually impaired. This work investigates how sound and haptics can be used to communicate information of both detailed graph based information and general trend information. In this paper we will argue that current work in the use of haptics (including our own) to provide effective tools to communicate information is insufficient to communicate graphs to visually impaired people. The reasons for this failure will be discussed and new theories drawn from information visualisation literature will be outlined.

WHY IS GRAPH "HAPTILISATION" DIFFICULT?

Current technology to provide graphic information to visually impaired users is mainly via the use of raised paper ("swell paper") diagrams. These provide raised lines which users can run their fingers over and by using their haptic system to build up a "picture" of the graph. The diagrams however are inflexible and cannot be dynamically created or modified. The MultiVis project has therefore sought to investigate the use of the PHANToM haptic device to allow for users to be able to browse dynamically creatable haptic graphs, which is a step towards allowing users to construct their own graphs.

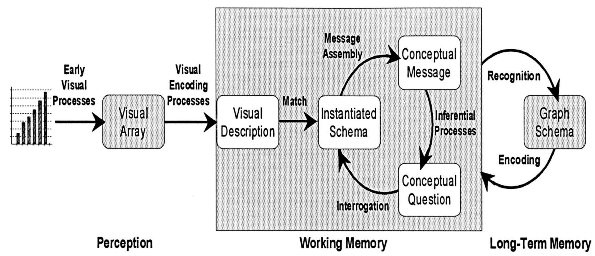

Figure 1: A diagram illustrating the major processes involved in Pinker's theory of graph perception [8]. Reproduced from Lohse [5].

Yu and Brewster [10, 11] have compared the use of PHANToM and raised paper bar graphs. They found that whilst the accuracy of participants' responses to questions such as "which bar is highest?", and "what is the trend of the graph?", were high, the subjective mental demands on users were also heavy. The time taken to answer the questions was also high, with over two minutes per graph. This time is likely to increase as the complexity of the graph is increased, and as such errors would be likely to increase as would the resources required by users. As was pointed out to the authors in a recent discussion with a mathematics teacher of visually impaired students, if students are fully blind they are likely not to be able to perform simple high school level mathematics such as graph transposition along an axis which significantly reduces their abilities to solve more complex unstructured problems. If haptic graphs are to be useable by visually impaired people to the same degree as visual graphs are to the sighted then ways must be found to effectively communicate salient graph information to users via alternative modalities.

In order to properly explain the problems with current haptic graph approaches we must consider how graphs are interpreted. Figure 1 shows a diagrammatic overview of a theory of visual graph perception by Pinker [8]. Pinker's theory considers that the image on the page is perceived and manipulated by the visual system into what he calls a visual description, which describes the primitives of the graph such as rectangle, line etc., the properties of these primitives (e.g. colour, texture etc.) and how they are related to each other (e.g. above, below etc.) This visual description is produced in a few milliseconds which is important given that the description is held in working (short term) memory which has both a limited capacity and information that degrades over time [7]. The visual description is then used to fill in a blank "graph schema" which has been selected from long term memory as the best "fit" for the visual description. The graph schema can, in theory, be infinitely complex and provides links between different pieces of information. For example, the schema may hold the height of each bar in a bar chart and provide easy mental links between the bars to allow a user to quickly identify trend information in the graph. The user interrogates the instantiated graph schema and seeks to answer a question which is likely to rely on only part of the schema. For example, if the user is trying to identify the trend of a graph the user will only need to consult relevant parts, such as the relative heights of the bars from left to right. Conversely, if some data are not available in the graph schema the user will need to trace back from the visual description in order to retrieve the information, with higher cost searches the further back the user must go, ultimately leading to an active visual search of the graph in order to locate the information required.

From Pinker's theory we can consider that haptic graph visualisation has two main problems that must be addressed. Firstly the speed at which information is loaded into a "visual description" needs to be increased. This is a problem given the relatively lower bandwidth of the human haptic system compared to the visual system. These problems are further complicated in "virtual" haptics due to the one point of contact available from devices such as the PHANToM, which reduces the bandwidth available. Since the visual description (and the instantiated schema) are held in working memory the information already processed by the user will degrade over time, meaning that less can be held at any one time and more reference back to the original source will be required in order for users to answer questions. In a comparison of a number of studies which compared different visual representations of graphs, Lohse [5] identified that different graphs more effectively support different types of question. Those representations that minimised the demand on short term memory, by making the information required to answer the question explicit, tended to outperform other representations that did not make such information as obviously available. In order to make graphs fully accessible to visually impaired users it is required to both maximise the use of the bandwidth available and support users' working memory.

MULTIPLE VIEWS OF HAPTIC GRAPHS

Although visual graphs are simple to understand visualisations, from the discussion of the previous section we can see that presenting such graphs in a non-visual way presents several problems, mostly in placing unreasonable demands on working memory. Whilst such problems have not been considered in haptic visualisation, much classical visualisation work has considered that in some cases it may be better to provide multiple views onto the same data with each view presenting a visualisation designed to provide some specific information not easily available in another view [3]. We propose that the same approach could be applied to haptic graph visualisation, with specific optimised haptic "views" being used to prioritise those parts of the user's graph schema that assist in answering specific questions. For example, quickly communicating the relative values of each bar in a bar graph to allow for trend information to be determined. Each view would only be useful for a small subset of tasks so many would be required to communicate all of the information required, leading to a question as to how users would navigate these different views. We propose that Lederman and Klatzky's exploratory procedures (EPs) [4] could be exploited to provide access to these different views. Whilst EPs are much talked about in the literature, there has been little discussion as to how these could be exploited with haptic graphs. Wall and Brewster [9] have proposed the use of different haptic properties to represent bars in a bar graph but have not considered this as part of a different view as described here. For example, lateral motion which is used to detect features such as friction and texture could be used to provide a rapid overview of the graph (see Figure 2) by the user moving a haptic stylus quickly across the bars, the relative friction of each bar producing a sensation of graph trend. This view would not provide any detailed information on the exact values of the graph but would more quickly allow a trend to be built up. Additional EPs could be exploited for other questions, e.g. the use of compliance for two bars which are not adjacent to each other.

Using this approach the information that is required to determine the trend of the graph is transferred to the graph schema first, maximising the use of the available bandwidth and thus reducing the demands on working memory. Additionally since the use of EPs is well known, gesture recognition could be used to determine what EP, and thus which view of the graph, the user is trying to access. In doing this the graph could be slightly modified to optimise the presentation of that view. For example, the ability of a user to determine the trend of the graph shown in Figure 2 depends on where users browse the graph. If the stylus is too high, the user may easily miss a bar (e.g. April in Figure 2). With the ability to predict the EP being applied, all bars could be "raised" such that the user could not miss any bar.

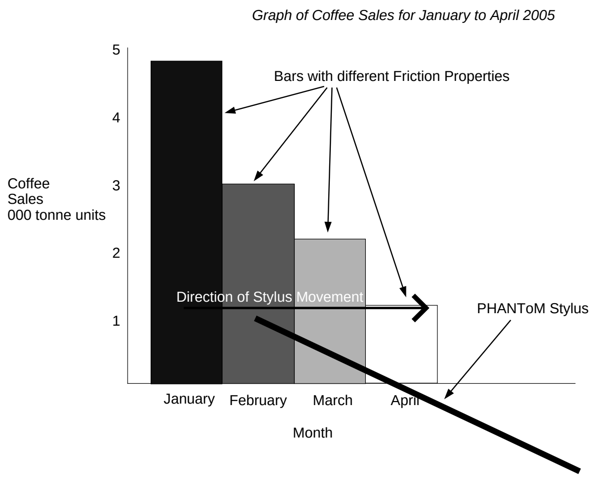

Figure 2: An example of a bar graph showing coffee sales over a four month period, and how the trend of coffee sales may be presented to the user via the exploitation of exploratory haptic procedures.

HAPTIC AND AUDIO INTEGRATION

One of the issues of the use of multiple views is the ability of users to form associations between them, such that the user can understand the underlying data model and as such integrate the different views to form a bigger picture. Card [3] discusses linking and brushing, where data in multiple views are linked, e.g. by using the same colour. However, the user would need to be in the correct part of the haptic visualisation to "observe" this change which may not be possible or convenient. One way in which this association might be made however is through the use of a metaphorical physical model. This would allow the easy integration of another useful feature in the presentation of non-visual graphs: sound. Sound graphs have been in existence for the last twenty years [6], and tend to have been used for line graphs, where the x-axis is mapped to time and the y-axis is mapped to musical pitch. Such graphs have been shown to be effective at communicating trend information [1] and users have been shown to be able to draw approximations of the graphs [2]. It remains a question however, how best to integrate haptic and sound views such that each presents a unique viewpoint on a data set, yet allows these views to be integrated leading to a better understanding of the underlying data set. However, when we consider that haptic sensation arises from the interaction of two objects (one of which is directly or indirectly attached to a human being) and that sound occurs (at least naturally) due to the interaction of two objects, exciting possibilities for multiple view integration exist.

A common childlike activity is to run a stick along a fence of iron railings, producing a rhythmic hitting sound. This same approach can be used to integrate sound and haptics. Consider that the graph as shown in Figure 2 is actually a sequence of metal poles driven into the earth at different heights. The poles can be driven further into the ground or extracted by the user's "stick" (the PHANToM stylus.) Hitting each pole with the stick produces a sound which is in pitch, proportional to the height of the pole. The user can then, in order to get trend information, run their stick across the graph in much the same way as using friction in the example from the previous section. The extensive research on sound graphs could then be exploited to communicate trend information to the user.

RESEARCH QUESTIONS

Whilst the ideas discussed in the previous sections may lead to significantly improved browsing of haptic graphs and may reduce the demands on working memory making it easier for users to gain specific understanding from the graph there are issues that must be considered. In this section we will introduce several research questions which must first be answered if we are to be able to exploit the work discussed above.

Firstly, what are the salient properties of individual graph types that would need to be communicated? Clearly there are a limited number of unique views that can be communicated, and it is therefore important to determine what information is and is not relevant to answering particular questions. It is clear in the example given in the previous section what information is required to identify trend information. However, what other questions do users wish to know about graphs and what features of the graph are relevant to answering these questions?

How can EPs be exploited to communicate different views? What are the most appropriate EPs for different views? There are a limited number of EPs available (especially if we remove temperature and volume which are difficult with the PHANToM), how can these be best exploited, can we reuse EPs with different haptic properties? Friction and texture are related and both are determined using the lateral motion EP, can each haptic attribute be used to encode some different view of the graph?

How accurately can information be communicated using these different haptic views? Wall and Brewster [9] have determined that up to a 40% difference may be required between two presented friction, texture and compliance values for participants to reliably determine those values as being different. Is this enough to allow different views of the graph to be effectively used?

Can users form associations between the different graph views? We have proposed a way in which the integration problem may be overcome when combining haptics and audio, however how can these associations be communicated between different haptic views? In order to avoid the user needing to spend time moving between different views of the graph, we have proposed that the views are collocated and overlapped by encoding each view as a different haptic property. Will this be confusing to users, will they be able to understand the underlying data, or will they fail to understand the graph?

CONCLUSIONS

Whilst there are many questions that must be answered before we can claim the full usefulness of multiple haptic views as discussed in the previous sections, the research described on graph theory perception is clear: simple reproduction of visual graphs in the haptic modality, whilst allowing for simple questions to be answered, are problematic. The time taken by users is greater than using the visual modality, and due to both the time taken and the demands that are placed on working memory; the ability of participants to gain understanding of the graph is reduced. Therefore we must consider new, novel ways to exploit the available haptic bandwidth to communicate information that is salient to the user's current task, and by doing so allow a greater understanding of the information presented to be gained.

ACKNOWLEDGEMENTS

This work is supported by EPSRC grant GR/S86150/01.

REFERENCES

1. L. Brown, S. Brewster, R. Ramloll, W. Yu, and B. Riedel. Browsing modes for exploring sonified line graphs. In *Proceedings of BCS-HCI 2002*, volume 2, pages 6–9, London, UK, 2002. BCS.

2. L. Brown and S. A. Brewster. Drawing by ear: Interpreting sonified line graphs. In *Proceedings of ICAD 2003*, pages 152–156, Boston, Massachusetts, 2003. ICAD.

3. S. Card. Information visualization. In J. A. Jacko and A. Sears, editors, *The Human Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications*, pages 544–582. Lawrence Erlbaum and Associates, Mahwah, New Jersey, 2002.

4. R. J. Klatzky and S. J. Lederman. Touch. In A. F. Healy and R. W. Proctor, editors, *Handbook of Psychology: Volume 4 Experimental Psychology*, volume 4 of *Handbook of Psychology*, pages 147–176. John Wiley & Sons, Hoboken, Hew Jersey, 2003.

5. G. L. Lohse. Models of graphical perception. In M. Helander, T. Landauer, and P. Prabu, editors, *Handbook of Human-Computer Interaction*, pages 107–135. Elsevier, Amsterdam, 2 edition, 1997.

6. D. L. Mansur. *Graphs in Sound: A Numerical Data Analysis Method for the Blind*. M.sc., University of California, 1985.

7. G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 63(1):81–97, 1956.

8. S. Pinker. A theory of graph comprehension. In R. Freedle, editor, *Artificial Intelligence and the Future of Testing*, pages 73–126. Laurence Erlbaum Associates, Hillsdale, New Jersey, 1990.

9. S. A. Wall and S. A. Brewster. Assessing haptic properties for data representation. In *CHI 2003*, volume 2, pages 858–859, Ft. Lauderdale, Florida, 2003. ACM Press.

10. W. Yu and S. Brewster. Comparing two haptic interfaces for multimodal graph rendering. In *IEEE VR2002*, Florida, USA, 2002. IEEE.

11. W. Yu and S. Brewster. Multimodal virtual reality versus printed medium in visualization for blind people. In *ASSETS 2002*, pages 57–64, Edinburgh, Scotland, 2002. ACM Press.