Interact, Excite, and Feel

Parisa Eslambolchilar
FIT Lab
Computer Science Department
Swansea University of Wales
SA2 8PP, Wales
p.eslambolchilar@swansea.ac.uk

Rod Murray-Smith
Department of Computing Science
Glasgow University, Glasgow
G12 8QQ, Scotland
&
Hamilton Institute
NUI, Maynooth, Ireland
rod@dcs.gla.ac.uk

ABSTRACT

This paper presents a dynamic system approach to the design of multimodal interactive systems. We use an example where we support human behavior in a browsing task, by adapting the dynamics of navigation using speed-dependent automatic zooming (SDAZ), allowing the user to switch smoothly among different modes of control. We show how the user's intention is coupled to the browsing technique via the dynamic model, and how the sdaz method couples the document structure to audio samples using a model-based sonification. We demonstrate that this approach is well suited to mobile and wearable applications, and audio feedback provides valuable information, supporting intermittent interaction, i.e. allowing movement-based interaction techniques to continue while the user is simultaneously involved with real life tasks.

Author Keywords

Multimodality, speed-dependent automatic zooming, mode switching, dynamic continuous interaction, sonification, mobile devices.

ACM Classification Keywords

H.1.2 User/Machine Systems, H.5.2 User Interfaces, H.5.5 Sound and Music Computing.

INTRODUCTION

We use tools to facilitate our interaction with computers and extend human powers to overcome the limitations of the body and controlling information flow. *Continuous control* is at the very heart of tool usage in the interaction between the human and computer as a tool [1, 21]. It differs from discrete interaction in that it occurs over a period of time, in which there is an ongoing relevant exchange of information between user and system at a relatively high rate, somewhat akin to vision/audio/haptic interfaces which we may not model appropriately as a series of discrete events [6]. It is also closely related to the development of dynamic systems since in these systems we can control what we perceive and we are dependent on the display of feedback (either visual, audio or haptic) to help us pursue our potentially constantly changing goals. Furthermore, feedback may influence an uncertain user's actions as more information becomes available [11].

Continuous control systems to overcome limited input states, which is a common problem on handheld devices, offer mode switching and transition and the controller supports the user in completing the task with less effort by changing the interpretation of the inputs to being reference values, rather than control commands [7].

In this paper, as an illustration of how this approach can support multimodal interaction, we use the example of browsing and sensing the structure of texts. Here, speed-dependent automatic zooming (SDAZ) method and the adaptive dynamics are coupled with sonification, which can be linked to a wide range of inputs and feedback/display mechanisms.

CONTINUOUS CONTROL AND TEXT BROWSING ON HANDHELDS

Navigation techniques such as scrolling (or panning) and zooming are essential components of mobile device applications, such as reading text, allowing the user to access to a larger information space than can be viewed on the small screen. Traditional scroll bar can increase the effort required to use the interface as the user must move back and forth between the text and the scroll bar. In addition, in a long document, a small movement of the handle can cause a sudden jump to a distant location, resulting in disorientation and frustration [2].

Speed-dependant automatic zooming (SDAZ), a relatively new navigation technique [4, 5, 17, 26, 28, 29], unifies rate-based scrolling and zooming to overcome the restrictions in screen space for browsing large images and long texts. Using this technique, the user can smoothly locate a distant target in a large document without having to manually interweave zooming and scrolling, and without becoming disoriented by extreme visual flow. Many researches in the past five years have demonstrated that speed-dependant automatic zooming (SDAZ) is well suited to implementation on small screen de-

1.(a) 1.(b)

Figure 1. Pocket PC and accelerometer attached to serial port (1a). 3 screen shots of the document browser (1b) showing a zooming window (red box) moving rapidly over the picture(left), the user finding the picture and landing there(middle), and the zoomed-in picture(right).

vices [3,7,8,19,24,25]; because compact visualisation coupled with zooming makes more effective use of small screen real estate.

A comprehensive study in [3] shows that coupled interaction (SDAZ) performs better than concurrent, but separate, zooming and panning techniques; because mapping limited input events to multiple actions requires the user to switch manually between modes, and it engages their visual attention more than coupled speed-zoom SDAZ.

Additionally, results in [3, 7, 8] demonstrate that users like tilt interaction because the interaction can be comfortably controlled in a single-handed fashion (Figure 1) and one-handed control requires less visual attention than bimanual interaction. However, using motion as input in a handheld device may reduce the quality of the visual display for the duration of the input, due to reflections from the screen and difficulty in concentrating on a rapidly moving screen [22]. One solution is to use audio/haptic feedback in such interaction scenarios. Audio or vibrotactile feedback may be crucial to support tasks or functionality on mobile devices that must continue even while the user is not looking at the display and is engaged in real life tasks [16]. Thus, audio/haptic feedback supports intermittent interaction. In this paper we focus on tilt-controlled SDAZ as an example of continuous interaction. However, the discussions in the following sections can be extended and applied to pressure sensitive or mouse-based interactions. In the next section we show that browsing and targeting can be facilitated by using a model-based multimodal text browser.

MULTIMODAL FEEDBACK IN A TILT-CONTROLLED SDAZ

Interaction models and state-space design presented in [8] focuses on the lower level of sensory-motor phenomenon and in this paper we demonstrate how this method may help in higher level of context of use.

Sonification is defined as the use of non-speech audio to convey information [13]. More specifically, sonification is the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation [14]. There are several stages in sonifying raw and abstract data. A general method which has been used in our application is presented in Figure 3. The starting point in adding multimodality is the raw data. This data, for example, a Microsoft Word document, is in "Data Tables" format. Data tables can describe many types of data including header, subheader, paragraph, table, figure, variables, and so forth in multiple dimensions. Data tables are transformed into visual/audio/haptic structures by applying visual/audio/haptic mappings. In our application visual structures are transferred to an image format, BMP or PNG.

Design

We have used two mechanisms by which we can support tilt-controlled SDAZ with multimodal feedback for rate of scrolling and structural information, to highlight specific information that is currently on the screen: audio and haptic. As an intuitive model of the sonification process, we can imagine the text on the screen to be embossed on the surface. This embossed type excites some resonating object (an elastic band or guitar string, for example) as it is dragged over the text (Figure 2). This physically motivated model is similar in nature to the model-based sonifications described by Hermann and Ritter [15].

Sound Synthesising

We simulate this model in our implementation by drawing an audio sample and placing that in an audio/vibrotactile buffer, as each line in the document"hits" the zooming window (see Figure 1). This technique is a form of granular synthesis; [30] gives other examples of granular synthesis in interaction contexts. A real world example would be the perception of continuous radiation values via discrete pulses from a Geiger counter; in this paper the continuous variable is the text flow rate (Figure 3(d)).

The strength of excitation associated with higher rate-of-scroll changes the acoustic/haptic response of the system, e.g., sampling frequency and volume of the audio sample decreases and provides the sense of distance to the text. As the speed increases, headings, sub-headings, figures and tables become relatively more prominent, to give a better overview of the document structure (an audio equivalent of semantic zooming), because these structures have a higher priority than the standard lines of text (Figure 3(c)). The rate of scroll controls the play rate of audio samples inside the audio buffer, i.e., the sound samples with higher priority are played sooner than sound samples for each line passing. Furthermore, volume and frequency have inverse relation to rate of scroll. At greater "heights" the features are blurred and damped suitably. As frequency and volume of the audio samples decreases when the user zooms out, scrolling slowly gives distinct audio feedback for individual lines of text. In this fashion, the audio texture as we pass over the document gives both an impression of the structure of the text, as well as the speed and zoom level at which we are passing it.

The particular audio cues chosen in our implementation of the text browsing are: scrolling over lines gives the impression of the sound of typing with a teletype keyboard. Passing over headers produces the sound of knocking a heavy object; sub-headers have a less important role than headers thus the sound is the same as header but with lower pitch, and the sound for figures gives the sense of dragging a large object.

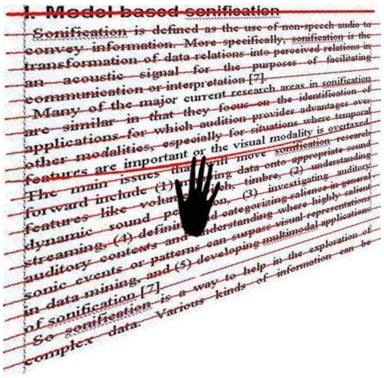

I. Model based sonification

Sonification is defined as the use of non-speech audio to convey information. More specifically, sonification is the transformation of data relations into perceived relations in an acoustic signal for the purposes of facilitating communication or interpretation [7].

Many of the major current research areas in sonification are similar in that they focus on the identification of applications for which audition provides advantages over other modalities, especially for situations where temporal features are important or the visual modality is overtaxed.

The main issues that will move sonification research forward include (1) mapping data onto appropriate sound features like volume, pitch, timbre, (2) understanding dynamic sound perception, (3) investigating auditory streaming, (4) defining and categorizing salience in general auditory contexts and understanding where highly salient sonic events or patterns can surpass visual representations in data mining, and (5) developing multimodal applications of sonification [7].

So sonification is a way to help in the exploration of complex data. Various kinds of information can be

Figure 2. Simulating dragging over the text by exciting an elastic band or guitar string. As an intuitive model of the sonification process for a document, we can imagine each line of the text on the screen as a string of guitar. When the user scrolls over the text he/she excites the string associated to each line and it gives a notion of embossed text.

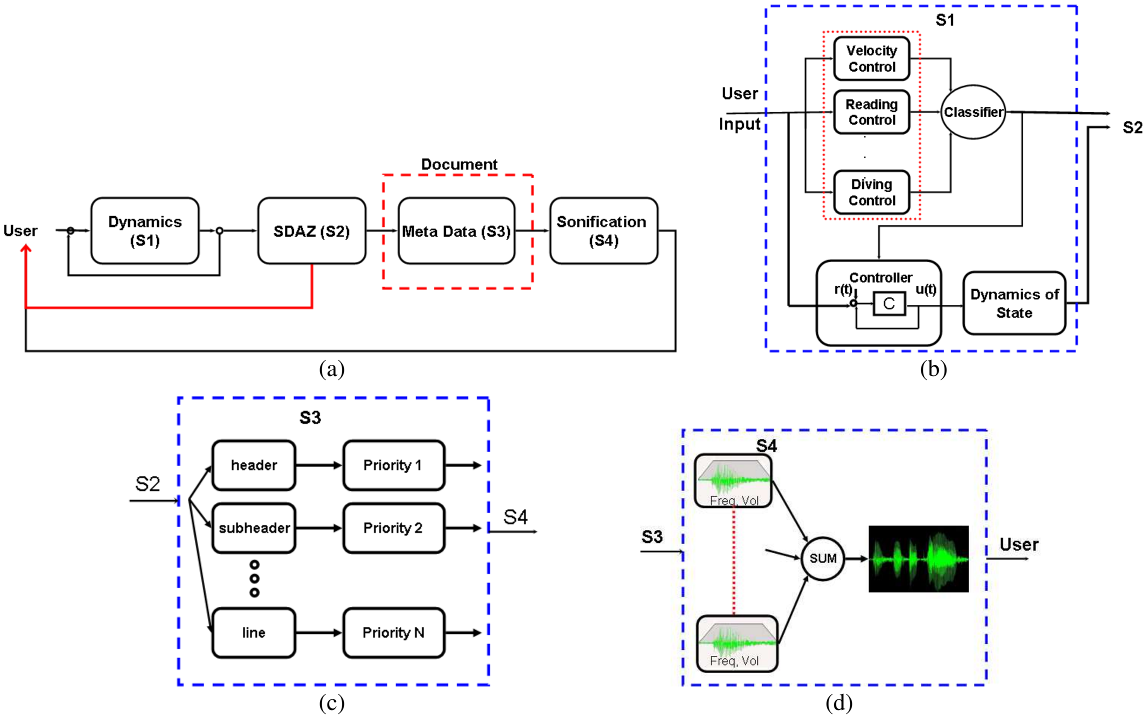

Figure 3. (a) A general framework of the model-based behavior system. (b) A classifier classifies the user's input. Its output and the user's input come to the controller and change the dynamics, e.g., state variables, and consequently speed of scroll and level of zoom changes (see Figure 1). (c),(d) The audio textures as the user passes over the document gives the both the impression of the structure of the text as well as the speed and zoom level at which she is passing it.

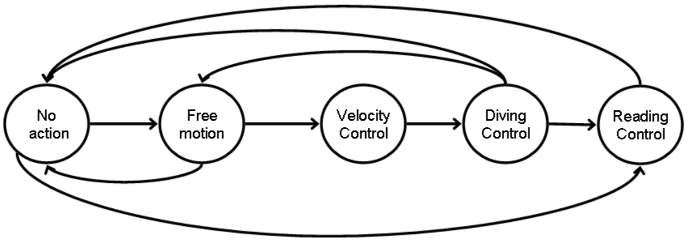

Figure 4. Control states/modes in text-browsing example and transitions among them.

The sound for tables highlights the tabular structure with a "bing" sound. A change in zoom level, as e.g., in target acquisition in the diving control mode, is accompanied by the sound of rushing air. This gives the user an impression of transition which is especially useful when there is no change in document position.

MODEL-BASED MULTIMODAL TEXT BROWSER

In a model-based multimodal text browser the user provides raw tilt input data as action via accelerometer and the user's action controls what she perceives from the display. Therefore, the dynamic system representation couples the user's intention to SDAZ via tilt input. However, in continuous interaction the user input may have different meanings. Therefore, interpreting the input and smoothly transiting the current mode of interaction to the mode the user is interested in is a delicate task [7, 9].

In tilt-controlled text browser task the user might look for a specific piece of information (searching) or target something on the display, or read the text. Thus, we define a number of different control modes: no action, free motion, velocity, diving and reading control (see Figure 4). Figures 3(b) and 5 illustrates how the user behaviour is being used to switch the control mode. The selected mode is then coupled to the SDAZ controlled variables, e.g., speed of scroll and zoom. For instance, in reading mode the controller adjusts the zoom level to stay in the maximum level (100% zoom), while if the user "breaks out" into free motion, the zoom level is decreased smoothly to a lower level.

Detecting state transitions

As we saw in Figure 4, the user can go to different control modes depending on the input behaviour. As an example, at the beginning the user is in the no-action state because there is no input from the user to the system. By tilting the device, therefore increasing the speed, the user goes to the free motion control mode. Free motion control is a transient state and the user may go back to no-action or velocity control mode afterward. In the velocity control mode the user is usually looking for some piece of information and she may spend a long time browsing at a steady speed in the document searching for a certain data. In this mode, the user controls the desired velocity manually. Then, the controller maintains this velocity automatically and completes the scrolling task with the desired speed for the user.

A general technique described in [9], Bayesian classification, can be also applied for model-based tilt-controlled SDAZ to classify the likelihood of being in one of control modes according to the joint probability of the input and output time-series [20]:

$$P(Mode \mid X) = \frac{P(Mode)P(X \mid Mode)}{P(X)} \tag{1}$$

where X is an array of previous inputs and possibly also outputs. $P(X \mid Mode)$ can be identified from experimental data collected from test users using standard density estimation models.[1]

Reference Signals as Inputs

The reference signal is a goal or a target the user is aiming to reach. In control engineering, engineers design systems to control variables with respect to reference signals, and they plan to be able to manipulate the references (inputs) when they want to get the system to control a variable at a different level (as when we change the setting on the thermostat) [23].

In control systems the controller supports the user completing the task with less effort by changing the interpretation of the inputs to being reference values, rather than control commands. In modern aircraft controllers there are different interpretations of aircraft controls depending on flight mode (e.g., take off, altitude-hold and so forth.) and blend seamlessly between modes [27].

Figure 5 presents a user's trajectory and position, velocity, and zoom-level data when he is looking for the 5th header in the document when there is only visual feedback from the screen. While the user tilts the device at a constant angle, the controller maintains the desired velocity automatically and completes the scrolling task with the velocity the user wants to achieve, rather than the user having to do this. Any change in the tilt angle for the controller means the user wants to change the reference signal. When the user finds the target and it is located inside the zooming window, she may return the PDA back to the equilibrium point thus the controller scrolls to the zooming window with a constant magnification level. This means that as the user performs the various tasks they switch between control modes automatically, and their inputs have different meanings, but that the transitions are always smooth and natural, and the user is often not even aware that their movements are having a different effect in the different modes.

Experiment

This study is carried out with an HP 5500 PDA with the Xsens P3C accelerometer and a headset (shown in Figure 1). In this study we map the zooming window's position and its speed to the audio or vibrotactile space. In three tasks, i.e., audio only, visual only and vibrotactile only, we asked 8 users to find three different targets, the 4th and 5th headers and the 7th figure in the document in a counterbalanced order. Figures 6, 7 and 8 present the users' trajectories, time series of the zooming window position, velocity, zoom, tilt input and phase plots (position vs. velocity). The users were given time to familiarise themselves with the system, and the specific document. They browsed the document in three versions of the application: (a) tilt-based SDAZ with visual display only, no audio or vibrotactile feedback, (b) with audio feedback only, no visual/haptic display (the user did the experiment blindfold), and (c) with vibrotactile feedback, no audio/visual feedback. The results show that the audio/haptic cues were sufficient for the users to distinguish when they passed the target, slowed down and returned to the target. The phase plots, trajectory and tilt angles are as smooth as the visual-only SDAZ. The difference here is that the users

---

[1] Density estimation is the construction of an estimate, based on observed data, of an unobservable underlying probability density function. The unobservable density function is thought of as the density according to which a large population is distributed; the data are usually thought of as a random sample from that population [20].

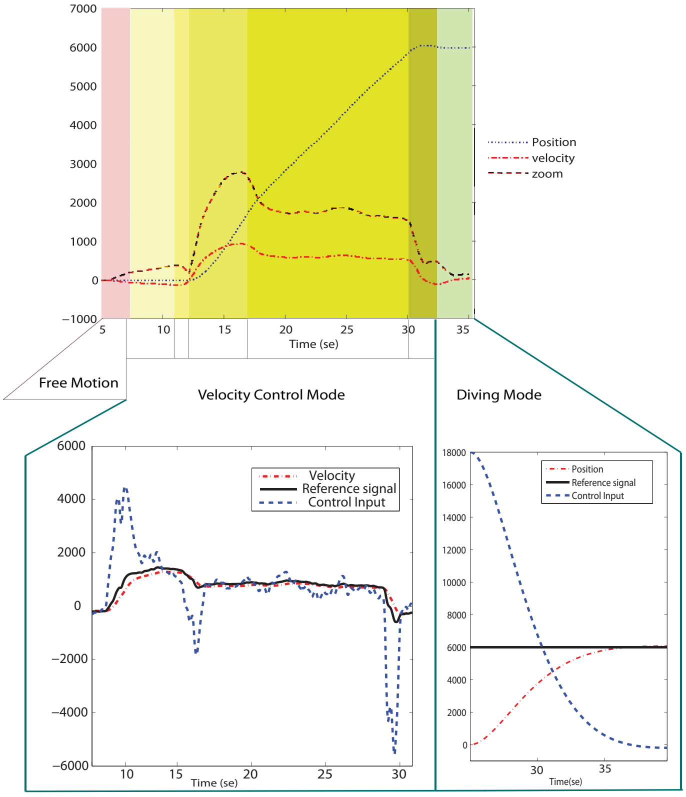

Figure 5. A user's trajectory and position, velocity, and zoom-level data when he is looking for the 5th header in the document and targeting that. Different colours highlight different modes of control or changes in the reference signal. Bottom figures highlight the changes in reference signals. In the velocity control mode, the system is moving velocity state variable toward the reference input. In the position control mode, the user is over the target and dives toward it. The system is moving position state variable toward the reference signal.

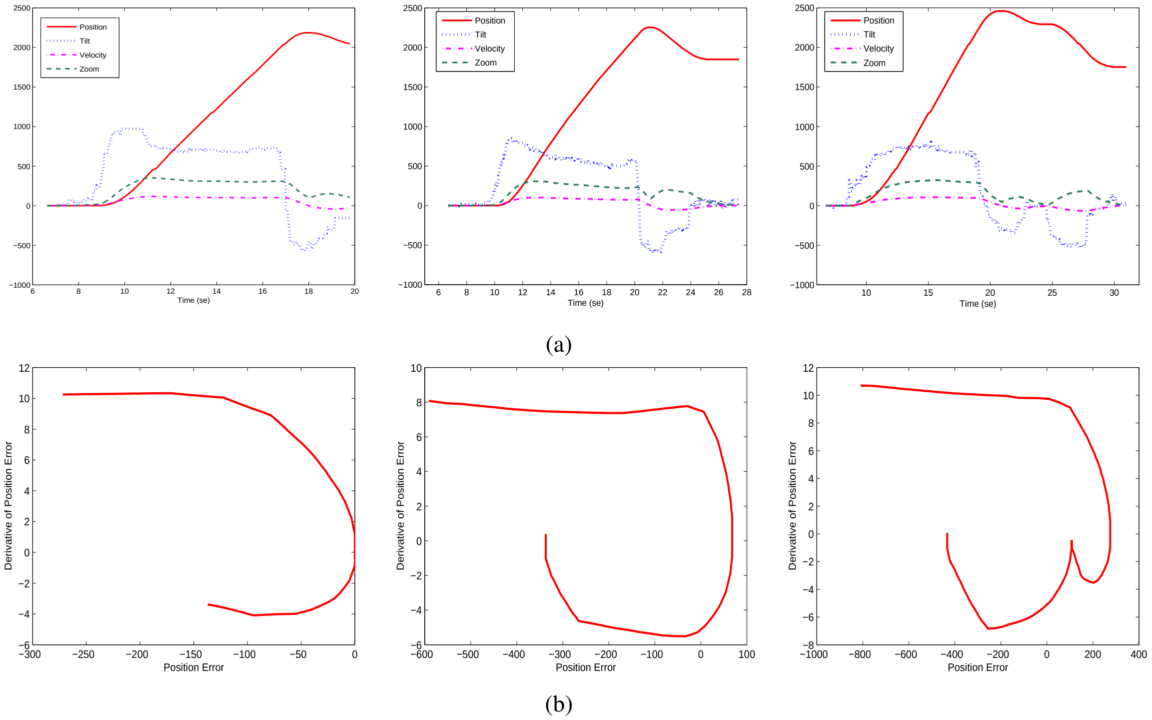

(a)

(b)

Figure 6. Searching for the 4th header in the document. (Left-to-right) Visual only, audio only and vibrotactile only tasks. (a) Time series of the zooming window position, velocity, zoom and tilt input. (b) Position vs. velocity.

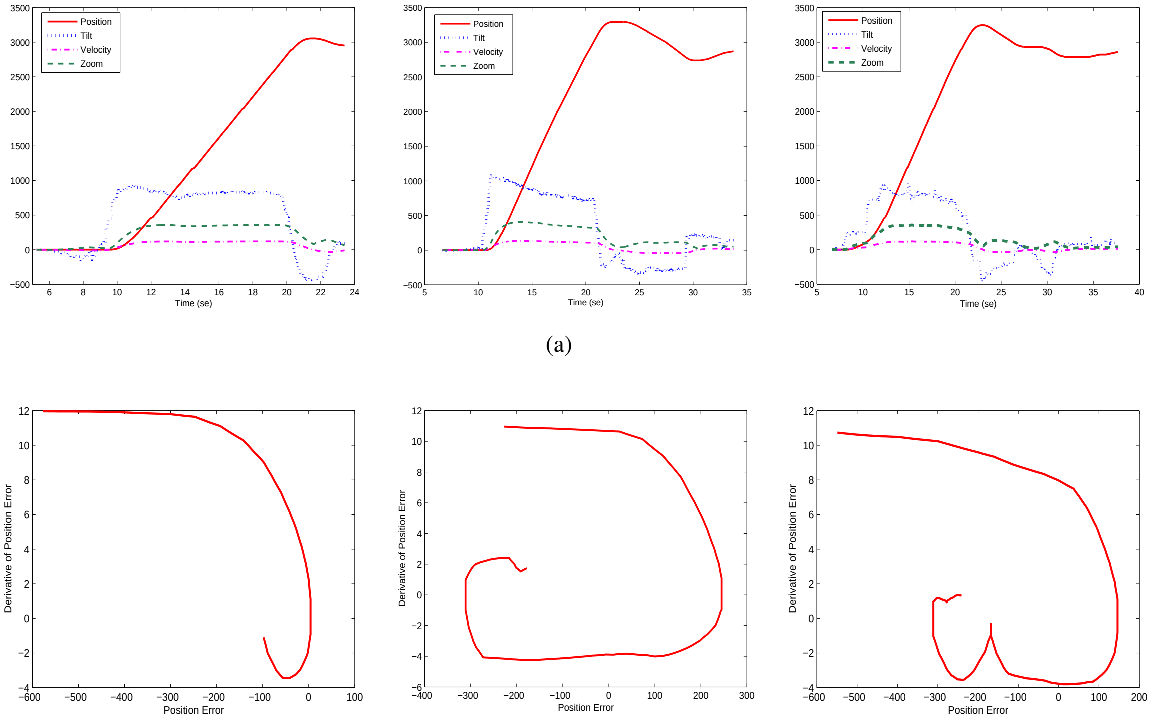

(a)

(b)

Figure 7. Searching for the 5th header in the document. (Left-to-right) Visual only, audio only and vibrotactile only tasks. (a) Time series of the zooming window position, velocity, zoom and tilt input. (b) Position vs. velocity.

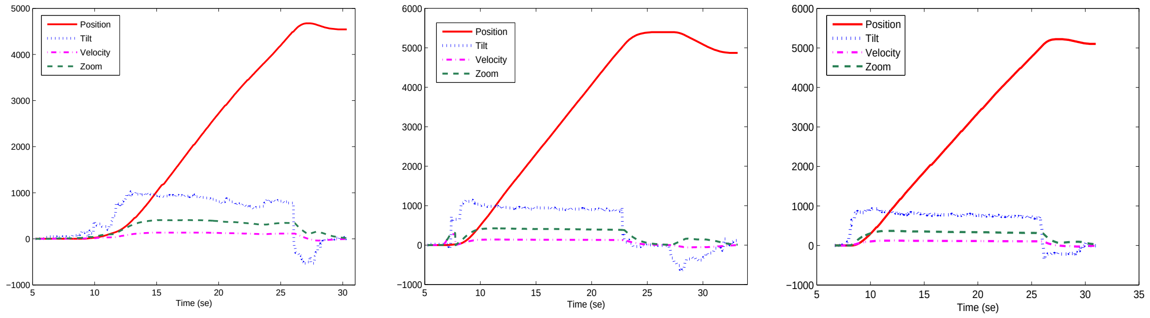

(a) Visual only task.

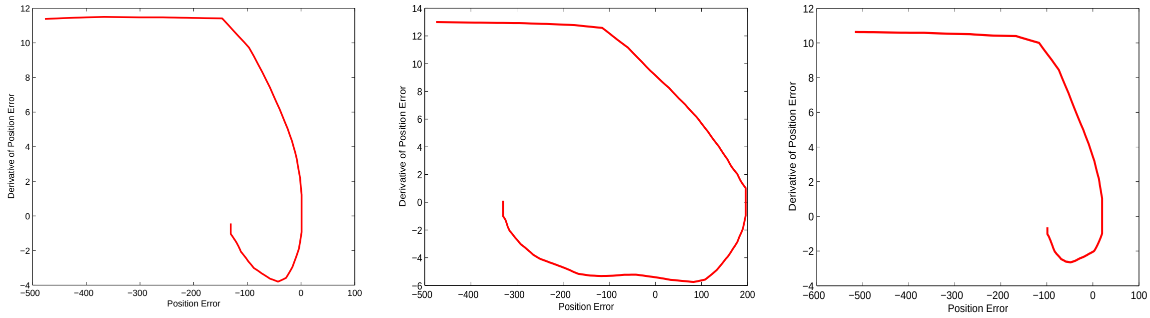

(b) Audio only task.

Figure 8. Searching for the 7th figure in the document. (Left-to-right) Visual only, audio only and vibrotactile only tasks. (a) Time series of the zooming window position, velocity, zoom and tilt input. (b) Position vs. velocity.

who worked with visual display only were faster in locating the target and landing on it. In both audio and vibrotactile feedback, we see overshoots on the trajectories, whereas the visual one allowed an over-damped target acquisition because of the extra predictive power of the visual display. Phase plots of velocity vs. position close to landing (See Figures 6 to 8) show that the narrower the target, more the overshoot happens.

Performance Measures

The "*performance criteria*" provide a value system for identifying an optimal path. "*The performance criteria are typically expressed as a function to be minimised* [18]." The differential and algebraic equations in the state-space mode can be used to determine the path through state space (or the control law) that minimises (or maximises) the performance criterion. Typical criteria include time, distance, or resource consumption.

In Fitts' law experiments [12], the goal is to capture the target in minimum time. Thus, the performance in this discrete positioning can be described as minimising the function:

$$J = \|t_f - t_0\| \tag{2}$$

t_0 and t_f are initial and final time in Fitts' Law task respectively. In our continuous tracking task, the function to be minimised would be:

$$\begin{aligned} uf &= filter(u) \\ J_t &= \|t_f - t_0\| && (3) \\ J_s &= \sum_{t=t_0}^{t_f} \|uf_{t+1} - uf_t\| && (4) \\ J_a &= \frac{\sum_{t=t_0}^{t_f} \|uf_{t+1} - uf_t\|}{t_f - t_0} && (5) \end{aligned}$$

t_0 and t_f are initial and final time and uf_t is the user's tilt input to the system, which has been filtered. Equation (3) presents the total time of the completing the task as a performance measure. Thus, when the use is comfortable with the system while interacting she should complete the task in the minimum time. Similarly, in equations (4) and (5) the total sum of changes or mean sum of changes in the tilt input should be minimised if the interaction is smooth.

We calculated the users' performance in these tasks, shown in Table 1. This provides strong evidence that audio and haptic feedback needed more effort from the user for targeting and landing and as the target gets bigger, performance difference between visual task and audio/haptic task grows smaller (compare targeting the figure with the header in Table 1). The system was also informally evaluated by a blind user. The user commented that

> ... [the system] has potential as a scrolling interaction for non-visual interfaces such as speech, but it has yet to be integrated with speech-based content. The sounds used are well-chosen from the point of view of controlling the speed of scrolling and drawing attention to key features in the document.

Table 1. Users performance in different targeting tasks using vision only, audio only and haptic only

Task	sum of changes (unit)
Visual– 4th header	36113
Audio– 4th header	41712
Haptic– 4th header	47293
Visual– 5th header	38501
Audio– 5th header	47118
Haptic– 5th header	55063
Visual– 7th Figure	51969
Audio– 7th Figure	55826
Haptic– 7th Figure	56577

An interesting observation relating to the vibrotactile work in [22] was that most users thought that the purely vibrotactile system also had audio feedback, and similarly in our work, although it was only audio feedback, users had a strong sense of vibrotactile feedback. A combination of the approaches seems a promising research direction.

CONCLUSIONS AND FUTURE WORK

In this paper we presented a model-based interactive method for browsing texts based on a sonification method, SDAZ technique and continuous interaction interface. The dynamic system representation coupled the user's intention to the browsing technique via tilt input, allowing the user to switch smoothly among different control modes.

We presented an audio/haptic feedback representation of the speed-zoom coupling involved in SDAZ. We demonstrated the applicability of the approach by implementing the SDAZ interface for a text browser system on a PDA instrumented with an accelerometer. Initial informal user evaluation of this implementation of multimodal SDAZ on a Pocket PC was positive. Sonifying each piece of structural information in the document was easy to understand by users and they could feel different textures in the document even the blind people.

Audio/vibrotactile feedback supports tasks running on mobile devices even while the user is not looking at the display and is engaged in tasks in real world. Thus, audio/haptic feedback supports intermittent interaction. The results showed that the audio/haptic cue is sufficient for the blindfolded users to distinguish where the target is but when they start to land they have already passed the target, so they have to adjust the tilt and go back to the target. Additionally, the users' activities in the audio and haptic only interface were higher than the visual only, most likely because the vision is the leading sense. Providing audio feedback about the user's predicted position instead of their current position may solve this problem [10].

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of IRCSET BRG SC/2003/271 Continuous Gestural Interaction with Mobile devices, HEA project Body Space, and SFI grant 00/PI.1/C067, the IST Programme of the European Commission, under PASCAL Network of Excellence, IST 2002-

506778. This publication only reflects the views of the authors.

REFERENCES

1. M. Beaudouin-Lafon. Designing Interaction, not Interfaces. In *AVI'04: Proceedings of the working conference on Advanced visual interfaces*, pages 15–22, Gallipoli, Italy, 2004. ACM Press.
2. S. Brewster, P. Wright, and A. Edwards. The Design and Evaluation of an Auditory-Enhanced Scrollbar. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems: Celebrating Interdependence, CHI'94*, pages 173–179, Boston, Massachusetts, USA, 1994.
3. T. Büring. *Zoomable user interfaces on small screens – Presentation and interaction design for pen-operated mobile devices*. PhD thesis, University of Konstanz, Germany, 2007.
4. A. Cockburn, J. Looser, and J. Savage. Around the World in Seconds with Speed-Dependent Automatic Zooming. In *ACM Conference on Human Factors in Computing Systems, Vancouver, Canada*, pages 35–36, 2003.
5. A. Cockburn and J. Savage. Comparing Speed-Dependent Automatic Zooming with Traditional Scroll, Pan, and Zoom Methods. In *People and Computers XVII: British Computer Society Conference on Human Computer Interaction, Bath, England*, pages 87–102, 2003.
6. G. Doherty and M. Massink. Continuous Interaction and Human Control. In J. Alty, editor, *Proceedings of the XVIII European Annual Conference on Human Decision Making and Manual Control*, pages 80–96, Group-D Publications, Loughborough, 1999.
7. P. Eslambolchilar. *Making Sense of Interaction Using a Model-Based Approach*. PhD thesis, Hamilton Institute, National University of Ireland, NUIM, Maynooth, Ireland, October 2006.
8. P. Eslambolchilar and R. Murray-Smith. Tilt-based Automatic Zooming and Scaling in mobile devices-a state-space implementation. In S. A. Brewster and M. D. Dunlop, editors, *Mobile Human-Computer Interaction–MobileHCI 2004: 6th International Symposium*, volume 3160 of *Lecture Notes in Computing Science*, pages 120–131, Glasgow, Scotland, September 2004. Springer-Verlag.
9. P. Eslambolchilar and R. Murray-Smith. Model-Based, Multimodal Interaction in Document Browsing. In *Machine Learning for Multimodal Interaction, Invited paper, 3nd Joint Workshop on Multimodal Interaction and Related Machine Learning Algorithms, MLMI*, volume 4299 of *Lecture Notes in Computer Science*, pages 1–12, Washington DC., USA, May 2006. Springer Berlin.
10. P. Eslambolchilar, R. Murray-Smith, A. Crossan, S.Dalzel-Job, and F. Pollick. *Handbook of Research on User Interface Design and Evaluation for Mobile Technology*, chapter Model-based Target Sonification in Small Screen Devices: Perception and Action. Idea Group Reference, 2007.
11. G. Faconti and M. Massink. Continuous interaction with computers: Issues and Requirements. In C.Stephanidis, editor, *Proceedings of Universal Access in HCI, Universal Access in HCI - HCI International 2001*, volume 3, pages 301–304, New Orleans, 2001. Lawrence Erlbaum Associates, Inc.
12. P. M. Fitts. The Information Capacity of the Human Motor system in Controlling the Amplitude of Movement. In *Journal of Experimental Psychology*, volume 47, pages 381–391, 1954.
13. W. W. Gaver. *Auditory Interfaces*. Handbook of Human-Computer Interaction, 2nd edition, 1997.
14. T. Hermann and A. Hunt. The discipline of interactive sonification. In T. Hermann and A. Hunt, editors, *Proceedings of the International workshop on interactive sonification*, Bielefeld, Germany, January 2004.
15. T. Hermann and H. Ritter. Listen to your data: Model-based sonification for data analysis. In G. E. Lasker, editor, *intelligent computing and multimedia systems*, pages 189–194, Baden-Baden, Germany, September 1999.
16. K. Hinckley, J. Pierce, M. Sinclair, and E. Horvitz. Sensing Techniques for Mobile Interaction. In *UIST'00: Proceedings of the 13th annual ACM symposium on User interface software and technology*, pages 91–100, San Diego, CA, USA, 2000. ACM Press.
17. T. Igarashi and K. Hinckely. Automatic Speed-Dependent Zooming for Browsing Large Documents. In *UIST'00: 13th Annual Symposium on User Interface Software and Technology*, pages 139–148, San Diego, CA, USA, 2000. ACM Press.
18. R. J. Jagacinski and J. M. Flach. *Control Theory for Humans: Quantitative approaches to modeling performance*. Lawrence Erlbaum Associates, Inc., Mahwah, New Jersey, 2003.
19. S. Jones, M. Jones, G. Marsden, D. Patel, and A. Cockburn. An evaluation of integrated zooming and scrolling on small-screens. In *International Journal of Human-Computer Studies*, volume 63, pages 271–303. ACM Press, September 2005.
20. D. J. C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
21. M. McCullough. *Abstract Craft: Practical Digital Hand*. The MIT Press, 1998.
22. I. Oakley, J. Ängeslevä, S. Hughes, and S. O'Modhrain. Tilt and feel: Scrolling with vibrotactile display. In *Proceedings of Eurohaptics 2004*, Munich Germany, 2004.
23. K. Ogata. *Modern Control Engineering*. Englewood Cliffs, NJ: Prentice Hall, 1990.
24. D. Patel, G. Marsden, M. Jones, and S. Jones. Improving Photo Searching Interfaces for Small-Screen Mobile Computers. In *Proceeding of the 8th International Conference on Mobile Human-Computer Interaction, Mobile HCI'06*, pages 149–156, Espoo, Finland, September 2006. ACM.
25. D. Patel, G. Marsden, S. Jones, and M. Jones. An evaluation of techniques for browsing photograph collections on small displays. In *Mobile Human-Computer Interaction–Mobile HCI 2004, Lecture Notes in Computer Science*, pages 271–303, Glasgow, Scotland, September 2004. Springer.
26. J. Savage. The calibration and evaluation of speed-dependent automatic zooming interfaces. Master's thesis, Computer Science Department, University of Canterbury, New Zealand, 2004.
27. M. B. Tischler. *Advances in Aircraft flight Control*. Taylor & Francis, 1994.
28. A. Wallace. The Calibration and Optimization of Speed-Dependent Automatic Zooming. Honours Report, University of Canterbury, Christchurch, New Zealand, November 2003.
29. A. Wallace, J. Savage, and A. Cockburn. Rapid Visual Flow: How Fast is Too Fast? In *Proceedings of the Fifth Australasian User Interface Conference (AUIC2004)*, pages 117–122, Dunedin, New Zealand, 2004. Australian Computer Society, Inc., Darlinghurst, Australia.
30. J. Williamson and R. Murray-Smith. Granular synthesis for display of time-varying probability densities. In A. Hunt and Th. Hermann, editors, *International Workshop on Interactive Sonification (Human Interaction with Auditory Displays)*. Bielefeld University, Germany, January 2004.