Flipper: a New Method of Digital Document Navigation

Liyang Sun
GVU Center
TSRB / 85 5th Street, NW
Atlanta, GA 30332
sunl@cc.gatech.edu

François Guimbretière
Human-Computer Interaction Lab
University of Maryland,
College Park, MD, 20742
francois@cs.umd.edu

ABSTRACT
Page flipping is an important part of paper-based document navigation. However this affordance of paper document has not been fully transferred to digital documents. In this paper we present Flipper, a new digital document navigation technique inspired by paper document flipping. Flipper combines speed-dependent automatic zooming (SDAZ) [6] and rapid serial visual presentation (RSVP) [3], to let users navigate through documents at a wide range of speeds. It is particularly well adapted to rapid visual search. User studies show Flipper is faster than both conventional scrolling and SDAZ and is well received by users.

Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces - Graphical User Interfaces

General Terms: Design

Keywords: Experimentation; scrolling; flipping; visual search; SDAZ; RSVP

INTRODUCTION
Page flipping is an important aspect of the interaction with paper documents. It allows for fast navigations through large bound documents like books. In particular, page flipping is very effective in both perusal and visual search. During a perusal search, users are trying to get a quick grasp of the general structure and content of a document. For example consumers often flip a book in a bookstore to get a "feel" for a book. During a visual search, users use flipping to rapidly identify a passage in a book, using surrounding cues such as an image, a figure, or certain typographical features.

Several systems have been proposed to convey the affordance of books in the digital world, such as 3Book [1] and Zinio. These systems focused on providing a literal representation of page turning in order to make it easier for users to transfer their document navigation skills to digital books. Yet they did not explore flipping per se. Other systems such as speed-dependent automatic zooming (SDAZ) proposed by Igarashi and Hinkley [6] use scrolling combined with an automatic zooming mechanism to provide fast visual search. While SDAZ efficiently controls the visual flow on the screen, it might at times feel overwhelming because of the large number of pages on screen at high scrolling speeds. Furthermore, by using scrolling, it blurs the boundaries between pages, which serve as an important structuring mechanism for large documents.

In this paper, we present Flipper, a new digital document navigation technique combining aspects of SDAZ with Rapid Serial Visual Presentation (RSVP) [3, 4]. De Bruijin and Spence [3] noted that using Rapid Serial Visual Presentation of each page of a book would be a natural candidate rendering the feel of flipping, but Flipper is the first system to propose a practical implementation. By combining SDAZ and RSVP, Flipper offers a smooth transition between the wide range of speeds used during document navigation. With Flipper, users might first scroll a couple of lines at a time as they read a passage, then smoothly transition to RSVP to rapidly locate a passage they have just remembered. As in paper documents, RSVP maintains the page structure of the document, and supports the use of spatial information such as "the image in the upper left corner of a page". Once users have located the target of interest, Flipper's automatic backtracking mechanism helps users to rapidly access the surrounding text.

To evaluate the potential of Flipper for visual search, we ran an experiment comparing Flipper, a SDAZ system, and a system similar to the ubiquitous Adobe Acrobat Reader. Our results show that Flipper allows for faster visual searches than the other two techniques, and is generally well received by users. Furthermore, our results show that when the same rendering technique is used for the conventional scrolling method and SDAZ, there are no significant differences between the two methods. Based on these results, we discuss how our technique could be implemented to serve everyday use such as visual search and perusal exploration of books in online bookstores.

THE FLIPPER TECHNIQUE
When navigating through a document Flipper might be in one of three states: **scrolling** for slow navigation below three pages per second (pps), **page flipping** for fast navigation between 3 and 20 pps, and **backtracking** to help users locate their target when they stop their navigation abruptly. A typical interaction sequence is presented in Figure 1.

In the scrolling mode, Flipper uses a scrolling visualization that is well adapted to reading and small movement. Below

CHI 2005, April 2–7, 2004, Portland, Oregon, USA.
ACM 1-59593-002-7/05/0004.

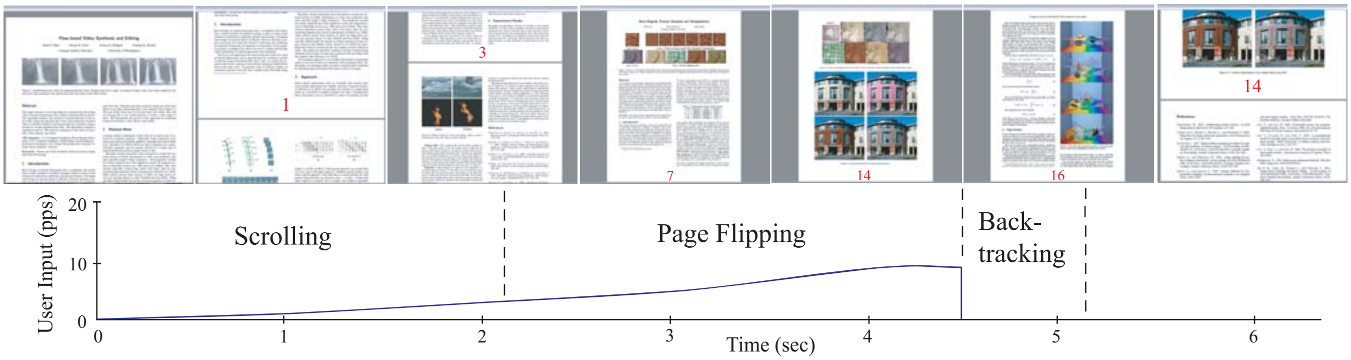

Figure 1: Flipper's Behavior. As speed (shown on the graph) increases, Flipper switches from scrolling to page flipping. An abrupt change in speed at 4.5s causes an automatic backtrack. Each image illustrates the state of the screen at one second interval. Page labels were added for clarity.

2 pps, the text on the screen is scrolled at full scale. Then, between 2 and 3 pps, we used SDAZ to provide a smooth transition between full scale and the scale at which a full page is guaranteed to fit on the screen. This scaling is required for the page flipping mode to work properly.

In the page flipping mode, we used an RSVP approach. Each page, now scaled to fit the screen, is presented one at a time at the rate requested by users up to a rate of 20 pps.

If users request to abruptly stop scrolling, (by setting the requested speed to 0 pps), the system enters the backtracking mode. In this mode, the system flips pages backwards at 8 pps to the document position 400ms before the abrupt speed transition. This takes into account the users' reaction time between the time they see a target and the time they effectively stop scrolling, and help them locate their target.

Since users control Flipper by specifying the rate at which pages are displayed, it seems natural to use an isometric device such as the IBM ScrollPoint [6, 8]. However initial reports from users showed the ScrollPoint was too rigid to be comfortable. Switching to a joystick solved this problem. Note that while our prototype used the joystick on a standard game pad (Saitek P2500 Rumble), joysticks can be easily added to PDA or tablet PC as exemplified by the jog dial navigator on the Sony CLIÉ PEG-N760C. Of course Flipper can also be controlled using a mouse.

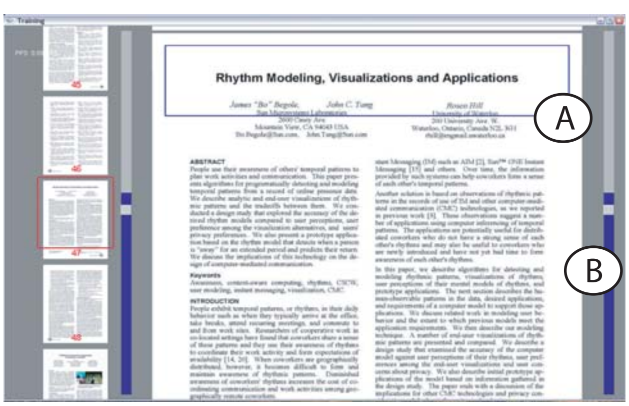

Figure 2: Search screen for Search and Return: A: (circled) target box; B: (circled) the area to search in is highlighted in blue. The Acrobat condition is shown. In Image Search the target box covers most of the screen.

EVALUATION

The goal of our evaluation was to compare Flipper to both SDAZ and a setting simulating the ubiquitous Acrobat Reader for simple visual search tasks. Our goal was to compare these techniques in as realistic a setting as possible. In particular, we wanted to compare not only each technique's raw speed, but also how well each technique lets users transition from high speed navigation to a speed fit for reading.

To do so, we evaluated each technique for two tasks. In the first task (Image Search), participants were asked to locate an image in the current document and place it completely within a target box on screen (Figure 2). Each target image was first presented for at least 4s on both a primary and secondary display. When they were ready, participants pressed the "start" button and began navigating the document presented on the primary display (the secondary display turned grey). When they had correctly placed the image inside the target box, participants pressed the "done" button. During the navigation phase, participants could review the target image on the secondary screen by holding down a button on the game controller. While the target was displayed, scrolling was disabled. The task time was computed from the time users pressed the start button to the time they pressed the done button and subtracting the time spent looking at the target during navigation. Users had to complete one task to move to the next. The Image Search task had 4 trials each for searching distances of 10, 25 and 50 pages to the target image. The scrollbar area highlighted in blue indicates the direction of search (Figure 2).

Our second task, "Search and Return," was similar to the first task, but after finding the image, participants were asked to return to the title of the paper to place it completely in the target box (Figure 2). Due to time constraints, only two searching distances (10 and 25) were used.

Experimental design

We used a within-subjects design, with technique, distance, and task as independent variables and search time as the dependent variable. We fully counter-balanced technique order, and the assignment of technique to documents. Trials were presented at random.

Twelve participants (3 males, 9 females) drawn from a varied background at the University of Maryland, College Park community, were recruited for our experiment. No participant had seen or heard of SDAZ or Flipper. Six participants reported playing games for an average of 3-6 hours per week. Participants completed the experiment in approximately 1.5 hours and were paid $20 for their participation.

Scrolling techniques

The Flipper technique has been described above, so we only present specifics here. The document remained the same size (125% magnification) up to 2 pps, when it started to zoom out until a single page filled the screen (90% magnification at 3 pps).

The configuration of scrolling with thumbnails was aimed at simulating the standard Acrobat setting, in which the current page is presented side by side with a set of thumbnails. The document remained fixed at 125%, and 5 thumbnails were visible at any give time, as shown in Figure 2. The maximum scrolling speed was 20 pps.

Our implementation of SDAZ was derived from the version used by Cockburn and Savage [2]. It zoomed from 125% magnification at rest to 6.25% at 6.73 pps, and stay at this scale up to 20 pps. This setting was similar to Savage's recommendation [7]. Since Savage's study limited zoom to 25% (or 4 pps), we ran a small pilot study to calibrate our system between 25% and 6.25%. Pilots indicated that 6.25% was the maximum comfortable zoom setting. The zoom-in animation of SDAZ was to zoom to the center of the screen using the same zoom/speed relationship as described by Cockburn and Savage [2].

Protocol

To familiarize them with the game pad transfer function, participants first completed a training session, in which they were asked to maintain a scrolling speed of 1, 4, 8, 10, and 20 pps. Only the joystick values and the scrolling speeds were displayed on screen. Next participants completed the training session for the current scrolling technique. The training session was the same format and length as the actual testing, except for the document used. To make participants aware of the range of speed available for each technique, they were first asked to maintain a constant speed of 1, 4, 8, 10, and 20 pps. Participants then completed the actual testing and a NASA Task Load Index (TLX) worksheet [5]. When all three methods were completed participants were asked to complete a questionnaire about the experiment. Throughout the experiment, participants were encouraged to take breaks as needed.

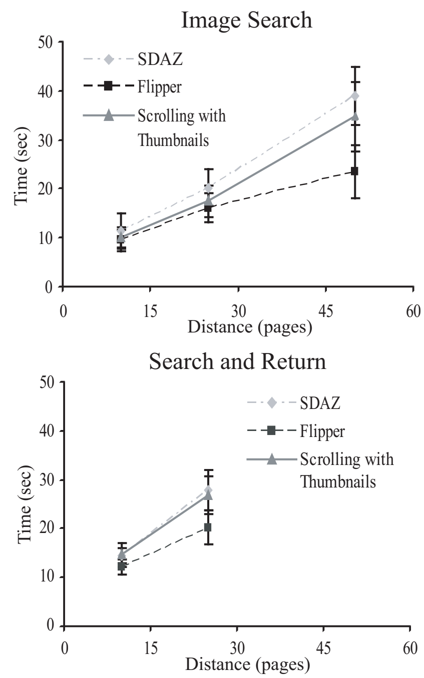

Figure 3: Mean Task Times for Different Techniques and Distances. Error bars show 95% confidence interval.

Apparatus

We ran the experiment on a laptop computer (Sager NP8790) with a 17" LCD display (1440 x 900 pixels). An additional 17" LCD display was used as the secondary display (NEC MultiSync 1700M+, 1280 x 1024). All techniques were implemented in C++ using OpenGL with GLUT 3.7.6. The rendering engine for all techniques was derived from Cockburn and Savage [2]. At initial zoom settings (125%) a page was 27 cm high. We also used a non-linear acceleration function to scale user input. We used 94 consecutive pages from the UIST 2003 proceedings during training and 3 sections of 94 consecutive pages from the SIGGRAPH 2004 proceedings during measurement. All pages were rendered as JPEGs.

RESULTS

To limit the influence of outliers and to ensure the normality assumption required by our statistical test was met, we used each user's median task times. We tested for both skill transfer and a possible influence of game playing but found no significant effect. Bonferrroni corrections were used to correct for multiple comparisons. For Image Search, a one-way ANOVA (technique) on the average time participants spent observing the target showed no significant effect for technique ($F(2,22) = .668$, $p = .523$). For Search and Return a similar ANOVA was performed, and a significant effect for technique was found

(F(2,22) = 3.54, p = .046), but Flipper had the lowest mean observing time (.261s).

A 2-way (technique X distance) ANOVA was performed on the normalized task completion times for both Image Search and Search and Return. For Image search, a significant main effect for technique was observed ($F(2,22) = 8.83$, $p = .002$). Post-hoc comparisons showed Flipper (16.3s) was significantly faster than SDAZ (23.5s, $p = .010$) and scrolling with thumbnails (20.8s, $p = .038$). SDAZ and scrolling with thumbnails were not significantly different from each other ($p = .423$). We also found a significant interaction between technique and distance ($F(4, 44) = 8.67$, $p < .001$), suggesting the advantage of Flipper is increasing as the target distance increases (Figure 3).

For Search and Return, similar results were found. There was a significant main effect for technique ($F(2,22) = 12.2$, $p < .001$), post-hoc comparisons showed Flipper (16.1s) was significantly faster than SDAZ (21.1s, $p = .005$) and scrolling with thumbnails (20.7s, $p = .004$). SDAZ and scrolling with thumbnails were again not significantly different ($p > .950$). A significant interaction between technique and distance ($F(2, 22) = 4.29$, $p = .027$) was again found (Figure 3).

The TLX worksheets did not show a uniformly lower workload rating for any scrolling method. Of the 12 subjects, 5 reported they would use Flipper to scroll through a large document, 5 reported they would use SDAZ, and 2 reported they would use scrolling with thumbnails.

DISCUSSION

The experiment showed that Flipper is significantly faster at tasks representative of visual searches. From our observation during pilots, we believe that this result can be explained by observing that SDAZ [2, 6] has to strike a balance between maintaining a constant flow and visibility of each page. Sometimes it has to scale down pages so much that they become difficult to identify. While this problem might be alleviated with semantic zooming, our pilots also showed for image rich contents, the scope of this approach is limited by the need to avoid image overlap.

While Cockburn and Savage [2] reported that SDAZ was 22% faster when compared to traditional scrolling techniques, our results seem to suggest there are no significant differences between the two. This can be explained by the difference in the experimental setup. They chose to compare an experimental implementation of SDAZ with commercial Acrobat Reader 5.0 to increase ecological validity. They also chose not to include the use of the thumbnail option included in Acrobat Reader. We chose to focus on the scrolling techniques themselves and compared them using the *same* rendering technique. Our results indicate the rendering technique plays a critical role in comparing different scrolling techniques.

Our tests also showed that the backtracking feature was sometime confusing for users. This may be caused by the fact that backtracking was animated which could have disoriented users. Additional research will be needed to understand the role of backtracking in our technique.

Our current prototype could be seen as a reference implementation in the sense that we did not pay attention to resource constraints such as memory requirements and bandwidth. In particular, we used high quality images at all speeds. However, our observations suggest as the flipping speed increases, the quality of the images used for rendering can be significantly reduced. For example, the text can be rendered as simple lines and image quality can be degraded. This suggests progressive rendering might be well adapted to Flipper, and should limit the memory requirements. This should make it possible to use Flipper in the context of perusal exploration at an online bookstore.

CONCLUSION

We believe Flipper effectively carries the affordance of page flipping to the digital world. Flipper not only provides users with a way to perform a rapid visual search, it also offers a smooth transition between fast and slow document navigation. Our tests show Flipper is the fastest method in our testing environment, and that its advantage increases with greater document size. We believe the advantages of Flipper are not restricted to document navigation, and can be applied to many other areas, such as online bookstores.

ACKNOWLEDGEMENTS

We would like to thank Andy Cockburn and Joshua Savage for providing us with the source code of SDAZ, Dave Levin and Corinna Löckenhoff for insightful comments on this work, and Arthur Fisk for his help with SPSS.

REFERENCES

1. Card, S.K., Hong, L., Mackinlay, J.D., & Chi, E.H., *3Book: A 3D Electronic Smart Book*, in *Proc. of AVI'04*. p. 303 - 307.
2. Cockburn, A., & Savage, J., *Comparing speed-dependent automatic zooming with traditional scroll, pan and zoom methods.*, in *Proc. of HCI '03*. p. 87 - 102.
3. De Bruijin, O., & Spence, R, *Rapid serial visual presentation: a space-time trade-off in information presentation*, in *Proc. of AVI'00*. p. 189 - 192.
4. Forster, K.L., *Visual perception of rapidly presented word sequences of varying complexity.* Perception and Psychophysics., 1970. **8**(4): p. 215 - 221.
5. Hart, S.G., & Staveland, L.E., *Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research*, in *Human Mental Workload*. 1988, Elsevier Science. p. 139-183.
6. Igarashi, T., & Hinckley, K., *Speed-dependent Automatic Zooming for Browsing Large Documents*, in *Proc. of UIST'00*. p. 139 - 148.
7. Savage, J., *The Calibration and Evaluation of Speed-Dependent Automatic Zooming Interfaces*, in *Computer Science*. 2004, University of Canterbury: Christchurch, New Zealand. p. 129.
8. Zhai, S., Smith, B.A., & Selker, T., *Improving Browsing Performance: A study of four input devices for scrolling and pointing tasks*, in *Proc. of INTERACT'97*, Chapman & Hall, Ltd. p. 286 - 293.