THE EFFECT OF CHANGES IN INFORMATION ACCESS TIMES ON HYPERTEXT CHOICES

D. Scott McCrickard

Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg VA USA 24061-0106
mccricks@cs.vt.edu

ABSTRACT

This research examines ways in which information processing decisions are affected by temporal delays in the acquisition of information. Specifically, we are interested in learning whether the value a user gives a piece of information depends on the amount of time spent obtaining the information, and whether the decision to obtain more information is affected by the time required to obtain it. An experiment examined people's reactions to delays in information propagation in a hypertext environment. The results suggest that rapid propagation times typically lead to better performance, but also they can lead to poorer choices in selecting paths in the hypertext.

1 INTRODUCTION

Everyone familiar with the World Wide Web has encountered delays in downloading a page, particularly when the network is heavily loaded or the page contains significant graphical content. However, these delays may not lead to poorer performance on a browsing task, in which the user must examine pages and select links that may or may not lead to the desired information. Longer delays may encourage more careful reading of the information, which could result in better selection of links.

Propagation delays occur when requests for information are slowed by excessive network traffic, slow connections between machines, and other factors. Depending on the time of the day and site location, a file that typically takes a few seconds to download may take considerably longer.

This paper examines ways in which information processing decisions are affected by temporal delays in the acquisition of information. Specifically, we are interested in learning whether the value a user gives a piece of information depends on the amount of time spent obtaining the information, and whether the decision to obtain more information is affected by the time required to obtain it. We plan to examine this question in a hypertext environment similar to the World Wide Web.

The hypertext model provides the perfect environment for incomplete information processing. A hypertext link is an embedded access point to another document that is activated by clicking on a highlighted word or phrase in the current document. The very nature of this link suggests that a user will move on to the next document before completely processing the information in the current one. If the costs (in time) of the current document as well as of obtaining a new document are low, people might follow a link without completely exploring the current one. On the other hand, if the costs are higher, more time should be spent on evaluating the current information and in choosing more information.

We designed, implemented, and ran an experiment to examine people's reactions to delays in information propagation in a hypertext environment. The pages in the environment are simplified text-only versions of pages found on the Web. The participants were asked to find pieces of information using the environment.

We simulated a delay in the loading time of each page, as if the page were being downloaded from a remote site. The participants were divided into three groups, one group always had a delay of 2 seconds, another 8 seconds, and the third 16 seconds. We measured the amount of time required and the number of links traversed to find the solution.

The remainder of this paper discusses the background, construction, and results for this experiment. Section 2 examines some of the related work in the field. Section 3 outlines our hypotheses. Section 4 discusses the details of the experiment. The results are in Section 5, and some conclusions are in Section 6.

2 RELATED WORK

Effects of response times have been a topic of interest for over thirty years. Understanding effects of long response times is important for graphical systems, particularly virtual reality (VR) systems, where expensive graphical calculations often resulted in a delay in responses. Mackenzie and Ware (MacKenzie & Ware, 1993) looked at the effects of delays between movements and results when users are trying to acquire a target. They found that even a 75 ms lag time resulted in performance degradation.

Work in the area of Web page design mirror some of the issues addressed in this paper. For example, in his collection of online essays, Jacob Nielsen argues that the design of pages that require scrolling should be avoided, as should the design of high graphics sites (Nielsen, 1995). Reasons given are that Web users want their information quickly and compactly, and they are unlikely to spend more than 10 seconds waiting for a page to load.

Some work has been done in the analysis of Web delays on user perception of information worth. Chris Johnson reports that as retrieval delays increase, the value of the information decreases (Johnson, 1997). However, this result was shown for a single download with a fairly long delay that could be over 70 seconds (like those experienced when downloading video clips), while we are interested in much shorter delays of up to only 16 seconds (like those for mostly-text pages). O'Donnell and Draper show that users will adopt coping strategies to avoid downloading unnecessary information (O'Donnell & Draper, 1996). It will be interesting to observe in our experiment the things participants do to cope when presented with a long delay.

3 HYPOTHESIS

We expect that the time required to obtain some piece of information will be in direct proportion to the time used to process the information. That is, the longer it takes to obtain a page of information, the longer a user will spend evaluating that information before obtaining more information.

In other words, participants with small delay times should select one of the first pages that seems at all relevant, while those with larger delay times would spend more time processing the information on each page. There are two reasons we expect this. First, since the participants invested more time in obtaining the information, they should value it more and thus study it more closely. Second, since obtaining more information will be costly, the participants should spend more time choosing the next link from the information that they already have.

4 EXPERIMENT

The experiment asks participants to answer a series of five questions using information found in a hypertext environment. Each question requires the participant to traverse through the environment until a solution is found, then enter the information into an entry widget. We simulated a delay in the loading time of each page, as if the page were being downloaded from a remote site. The time required to load a page differed for the participants: for some, it took only 2 seconds to load each page, for other participants, 8 seconds, and for others, 16 seconds. We will measure the time spent and the number of scrolls performed on each page, and the number of link traversals and the total time required to discover the solution.

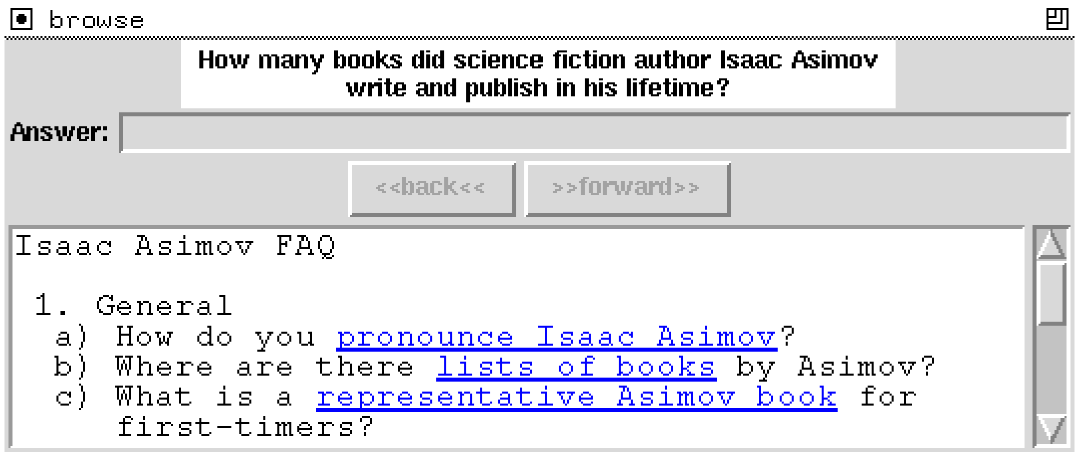

Figure 1: A screenshot from the experiment. The question is in the top box, with the answer entry widget below it. The "back" and "forward" buttons allow a user to view previously visited documents. The scrollable text area at the bottom contains the hypertext information. The links are colored blue and underlined.

Fifteen participants performed the experiment, five for each of the page loading delays. The participants were computer science graduate students with extensive Web experience. By using experienced participants, we hoped to minimize the learning curve effects for our hypertext interface. The nature of the tasks and a description of the interface were explained to them before they began. To minimize the time required to read and process the question, the participants first read the question, then pressed a button to load the first page and begin the timed search task. The participants were told that time is an important factor, and that Krispy Kreme doughnuts would be given to the person with the best combined time on all questions.

The hypertext environment is a simplified text-only model of the one used in most Web browsers like Netscape Navigator and Microsoft Internet Explorer (see Figure 1). Links to other pages are indicated by underlined and colored text. The participant clicks on a link to load the corresponding page and can move between previously viewed pages with "back" and "forward" buttons. In many cases, the solutions to the tasks can be obtained following a number of different paths, and it is not completely clear from the text which path is the best.

Each question has a different set of hypertext pages associated with it. The pages are simplified versions of pages found on the World Wide Web. Each set of pages (and each related question) deals with a different topic that we expect most participants to have some knowledge, though not enough knowledge to answer the question. The hypertext was designed such that the path of links to the answer was not immediately obvious, and none of the participants knew any of the answers before performing the experiment.

Section 5 describes the questions and explains the methods used by the participants to find the solutions. The system used for the experiment was implemented in the Tcl/Tk graphical scripting language. The system recorded the time of every user action (scrollbar clicks and drags, entries in the entry widget, and link traversals). Our hope was that this information, combined with observations of the participants interacting with the environment, would verify on our hypotheses.

5 OBSERVATIONS AND RESULTS

This section lists some general observations and a few results from our experiment. Seventeen participants began the experiment, and fifteen completed it, five for each of the three page loading delays (2 seconds, 8 seconds, and 16 seconds).

The first part of this section examines each question individually, exploring the expected and actual solutions. The second part focuses on the results from several individual participants, noting how their actions corresponded to our predictions.

	2-second delay						8-second delay						16-second delay					
#	A	B	C	D	E	μ	F	G	H	I	J	μ	K	L	M	N	O	μ
0	72	35	18	29	19	34.6	45	96	39	42	20	48.4	47	31	154	88	43	72.6
	6	1	1	1	1	1.4	1	3	1	1	1	1.4	1	1	3	3	1	1.8
1	15	52	40	40	50	39.4	94	80	42	79	126	84.2	100	221	97	94	91	120.6
	1	3	3	3	6	3.2	5	5	1	3	5	3.8	3	9	3	3	3	4.2
2	134	73	31	55	32	65.0	74	139	86	36	102	87.4	187	95	77	68	133	112.0
	13	4	2	2	2	4.6	2	6	2	1	2	2.6	6	2	2	2	2	2.8
3	134	40	15	20	10	43.8	23	17	59	34	49	36.4	64	144	47	36	30	64.2
	18	1	2	2	1	4.8	1	1	2	1	2	1.4	2	6	2	1	1	2.4
4	92	43	72	42	24	56.6	43	46	40	36	65	46.0	52	61	66	58	131	73.6
	8	2	2	2	1	3.0	1	2	1	1	2	1.4	1	2	1	1	3	1.6

Table 1: Summary of times required to answer each question (top line of each pair) and the number of links traversed in answering it (bottom line).

5.1 Per-question results

This section outlines each question in the experiment and analyzes the times required and number of links traversed to answer it. A summary of the times and number of links is provided in Table 1 with graphs of the results provided in Figure 2. The full data from the experiment is available by request (there is over 4000 lines of it).

The first question asks how many runs the Atlanta Braves scored in a baseball game several weeks ago. That game featured the last big offensive output by the Braves, so it is mentioned at least in passing on many of the pages, including Today's feature, Today's preview, and Yesterday's game. There is a Previous games link that we thought would be the most conservative choice, but it is located at the bottom of the page and a participant would have to scroll to find it.

As seen in Table 1, the majority of the participants in each group found the correct answer in the first visited link. Most scrolled down to find the Previous games link, but others used alternate links. We hoped that this task would provide participants with the opportunity to get accustomed to the environment and the questions.

The second question asks how many books author Isaac Asimov wrote and published. The hypertext environment is patterned after a FAQ (Frequently Asked Questions) list for Asimov. None of the questions in the FAQ list exactly matched the one to be answered, and in fact only two links contained the correct answer.

As a result, the participants visited many more links to find the correct solution. Surprisingly, more links were visited on average for larger delays, perhaps because the participants did not yet have a feel for the amount of time it would take to load a page. Correspondingly, the time required to find the correct solution was greater for higher delay times as well.

The third question asks the participant to find the year in which work on Mount Rushmore was begun. The pages for this question are very verbose, requiring a participant to scan a lot of information to find the link that will lead to the answer.

We anticipated that the participants with smaller delays would be more likely to click the first link that seemed at all relevant, since the cost of loading a page is fairly small. The results of this question seem to back up that claim. The participants with a 2-second delay on average chose more links than those with a longer delay. However, they still had a smaller overall time to find the correct answer, perhaps because the participants with a longer delay spent more time analyzing the information.

The fourth question asks for the number of Western films that feature John Wayne. There are a number of promising links in the visible portion of the page, particularly the Filmography link. However, none of these links lead directly to the solution, though the Filmography link contains another link that has the solution. If a participant scrolls down a few lines, a Westerns link that contains the solution will appear.

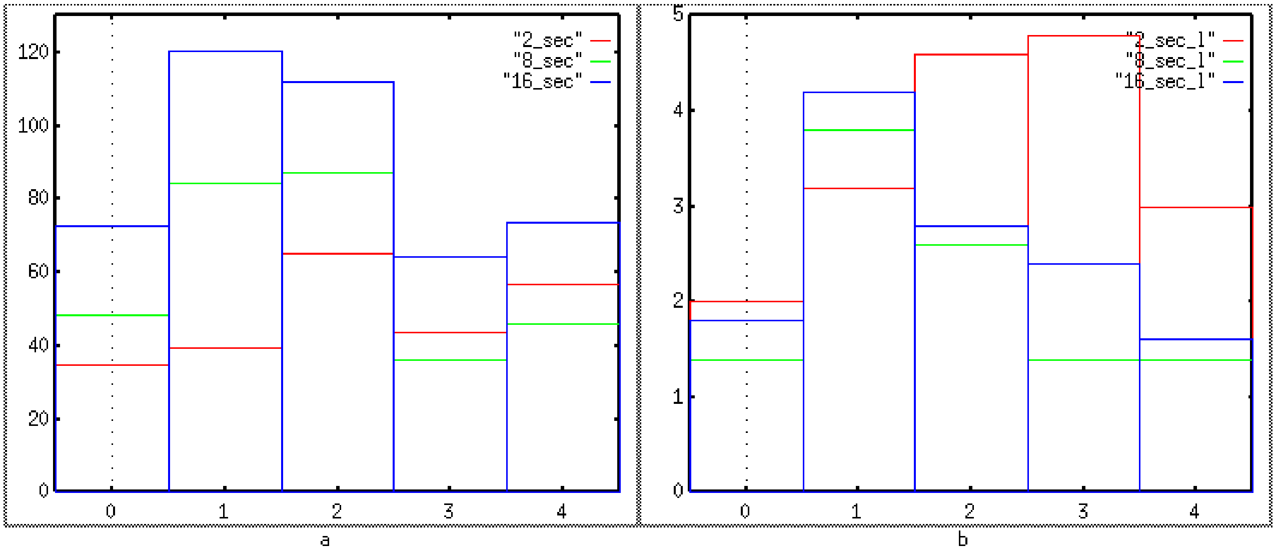

Figure 2: Graphical display of (a) average times and (b) average number of links visited for each of the five tasks (labeled 0-4).

By this point, it seems like the participants with 8- and 16-second page loading delays have learned to scan the entire page to try to identify the best link. Most of the participants with 2-second delays seemed to choose the best link visible without scrolling (the Filmography link), while most with the 8- and 16-second delay scrolled down to find the Westerns link. Surprisingly however, this also resulted in a smaller mean time to find the solution for the participants with an 8-second delay than for those with a 2-second delay!

The fifth question asks for the length of the Mississippi River. Similar to the last set of pages, this set has an encouraging FAQ link that is immediately visible, while a more promising Vital Statistics link can be found by scrolling. The FAQ link does have a link that leads to the correct solution, so it is not a dead end.

Again we saw that the participants with a 8-second delay complete the tasks more rapidly than those with a 2-second delay. While the mean times for the last two tasks were lower for the 8-second delay compared to the 2-second delay, the participants with a 16-second delay had a significantly higher average, probably in part because they were showing signs of frustration and boredom at the large loading delays (see Section 5.2 for details on this phenomenon).

Note from Figure 2, the participants with 2-second delays were finding the answers the quickest in the first task, even though they visited slightly more links. By the later rounds, the gap in links visited for the participants with 2-second delays was much greater, so much so that the average times for those with 8-second delays were slightly smaller than those with only a 2-second delay.

5.2 Per-participant results

This section looks at some of the more interesting reactions from individual participants.

One of the participants with a 2-second delay exhibited behavior we thought would occur more frequently. He clicked on any link that he thought would have any chance of containing the solution. He visited 8 pages for one question, 13 for another, and 18 for a third! Since none of the sets of pages had more than 7 or 8 unique pages, he was actually visiting many of the pages multiple times. Sometimes he would scroll blindly up and down the page, not so much looking for the answer but looking for another link to select. The small delay in page loading time made obtaining more information so cheap that he did not bother to process the information already in his possession. While others with a 2-second delay took non-optimal paths to the solution (settling for the first good link they found), none came close to the behavior of this participant.

Two of the participants quit the experiment before answering all of the questions. They claimed that they

#	2-second delay	8-second delay	16-second delay
0	15.3	26.6	24.3
1	10.3	14.2	12.7
2	12.1	25.6	24.0
3	7.1	18.0	10.8
4	16.9	24.9	29.6

Table 2: Average number of seconds participants spent on each page for each task.

were bored with waiting and suggested that perhaps the delay was longer than we had claimed. This result is in line with one noted by Johnson (Johnson, 1997), who predicted that at some point in time, users of a remote information retrieval system would "run out of time" and quit their download. These participants clearly felt that their own interests were more urgent than the completion of this experiment.

For those participants who did endure the longer delays, many were not completely focused on the task. One participant played with the window while waiting for the page to load, others looked around, several verbally harassed the proctor for building an experiment that is so slow. The participants seem to have an expectation that a hypertext page should load in under 10 seconds, as noted by Nielsen (Nielsen, 1995). Participants approach the problem differently depending on the delay, similar to the strategies observed by O'Donnell and Draper (O'Donnell & Draper, 1996). The participants with a 2-second delay use a rapid-fire selection method, those with an 8-second delay use the delay to think and plan, and those with a 16-second delay become annoyed and distracted.

6 CONCLUSIONS

We hypothesized that people would spend longer examining the data when its cost (in time) was higher, both because they spent more time obtaining the data and because they would have to spend more time obtaining new data. As we see in Table 2, that was the case in certain situations. The data collected for 2- and 8-second delays shows that a greater delay does indeed result in a larger information processing time, while the data for 16-second delays shows that after a point the information processing time becomes smaller. It may also be that (up to a point) the participants use the page loading delay to think about the information they had previously viewed and to plan their future strategy.

One might think that it would take longer to find a piece of data in a hypertext environment when there is a larger delay and a larger information processing time. However, this is not always the case. In this experiment, better decisions were made when more time was spent analyzing the data, resulting in less overall time spent on the search.

References

Johnson, C. (1997). What's the web worth? The impact of retrieval delays on the value of distributed information. In *Workshop on time and the web.* Staffordshire, England. (available at www.soc.staffs.ac.uk/seminars/web97/papers)

MacKenzie, I. S., & Ware, C. (1993). Lag as a determinant of human performance in interactive systems. In *Proceedings of the 1993 joint conference of the acm special interest group on human factors in computing systems and ifip interact (interchi 1993)* (pp. 488–493). Amsterdam, the Netherlands.

Nielsen, J. (1995). *Multimedia and hypertext: The internet and beyond.* Boston, MA: AP Professional. (See www.useit.com for online essays)

O'Donnell, P., & Draper, S. (1996). How machine delays change user strategies. *SIGCHI Bulletin*, *28*(2), 39–42. (Also appeared at July 1995 workshop on Temporal Aspects of Usability at the University of Glasgow)