Exploring Small Screen Digital Library Access with the Greenstone Digital Library

George Buchanan[1], Matt Jones[2] and Gary Marsden[3]

[1]Interaction Design Centre, Middlesex University
g.buchanan@mdx.ac.uk
[2]Department of Computer Science, Waikato University
mattj@cs.waikato.ac.nz
[3]Department of Computer Science, University of Cape Town
gaz@cs.uct.ac.za

Abstract. In recent years, the use of small screen devices has multiplied rapidly. This paper covers a number of different issues which arise when digital libraries are used in combination with such displays. Known limitations of small screens are presented to the Digital Library community. Two evaluations of pilot small-screen DL systems are presented, with some unexpected cultural and socio-technical concerns which arose. The pilot systems also demonstrate the delivery of small-screen access using an existing popular DL system.

1 Introduction

The usual tool for interacting with digital libraries is a generic web browser on a standard desktop display. However, there are many circumstances where this norm is not going to hold. Many uses of small-screen devices as access tools to DLs are emerging. For instance, handheld computers, such as Palm and Pocket-PC devices, are increasingly being used for casual browsing alongside extended work co-ordination tasks by knowledge workers – typically, frequent users of digital libraries and web-published information. Secondly, the desk-based stereotype of the information worker is both historically biased and under attack from changing work practices. For example, in many sciences key work is performed out of doors, in areas where broadband network access or physically large displays are impracticable, and mains power supplies an expensive luxury. Thirdly, in much of the world, particularly where population and/or wealth is sparse, the absence of current fixed infrastructures for services such as power and communications, is leading to a leapfrog to wireless-based technologies.

These differing pictures of the use and users impact all forms of information services, from executive information systems in businesses to open, public, digital libraries. It is unsurprising, therefore, that the combination of small screen technology and digital libraries is of increasing significance; for instance, Cathy Marshall and Christine Ruotolo have recently presented a paper on small-screen reading at JCDL 2002 [20], and other mobile use of DLs was observed in [19].

This paper presents two separate pilot small-screen DL systems implemented on Greenstone. We include the evaluation of a novel presentation of search results, and a brief evaluation of the provision of DL access through small displays using the WAP protocol, which has different limitations to HTML-based access.

Our findings are not only of interest in the context of small screens – cultural and technical expectations of users significantly impacted the outcome of our evaluations.

2 Background and Overview

The New Zealand Digital Library project's main work has been on the open-source Greenstone Digital Library software [29]. Greenstone's adoption in a wide variety of organisations has lead to an exposure to a wide variety of uses. Recent related DL projects include information resources for field workers for Non-Governmental Organisations in isolated areas; and in Cape Town, South Africa, one of the authors has observed the widespread adoption of handheld devices by students - offering an opportunity for providing information more pervasively than with desktop-based access. Circumstances such as this have alerted us to the need to support small-screen devices.

As well as being involved in digital library research, we have also studied the issues of small-screen web use [2, 7, 8], particularly the improvement of browsing between documents. In the context of small-screen DL access, as in our earlier work, there is a need to separate the different factors involved in wireless and/or palmtop systems. Our interest is focussed upon the peculiarities of small screens alone. The wide variation of factors such as network communications (wireless, wired or none), input devices (pen, keyboard, keypad, touch-pad etc.) and form factor (handheld or embedded) cannot be successfully confounded with display size if generalisable understanding is required.

Our earlier research [15] demonstrates that techniques which give benefits to users of small screens can also benefit users of normal, desktop-sized displays. Therefore, discovering small-screen access techniques for digital libraries may result in novel and effective means of accessing digital libraries generally, which particularly benefit those working with less display capacity.

3 Small Screen Usability

With the recent growth in popularity of a variety of small-screen consumer devices, increasing attention has been given to the challenges of providing effective interaction through small displays. However, it would be inaccurate to believe that small screens are a recent novelty – they have long been used in Automatic Telling Machines at banks, visitor information systems at museums and controlling machine tools. Unsurprisingly, related research dates back to these earlier uses.

The recent mushrooming of mobile phone usage has rather clouded matters by synonymising "small screen" with handheld consumer technologies and very small displays.

It is worth clarifying what size of display we intend when we say “small screen”. There are two primary types: micro-displays, more synonymous with mobile phones and containing only perhaps 80 letters in approximately 5 lines of text; and small-displays, common with Pocket-PC, Palm etc. handheld computers, containing approximately 500 or more characters in up to twenty lines of text depending on the screen orientation and text size. Compared to the capacity of the average desktop display, typically 7 to 12 thousand characters, either represents a significant loss of display capacity. Our recent work [4] indicates that some of the problems expected of micro-displays is not misplaced; however, we and others have demonstrated that small displays can be used with less penalty in terms of effectiveness (task time and success rate) than expected [7,8].

3.1 Reading on the Small Screen

A primary form, some may say the principal form, of activity in a digital library is the reading of documents. Studies of the reading of texts using small screens date back to their earliest uses. In various studies, including [9, 10, 23], researchers tested the reading rate and comprehension performance of users using various screen sizes and proportions (covering both small- and micro- displays). Consistently, researchers found no difference in comprehension rates across screen sizes [9]. The number of displayed lines of text little affected reading speed, with smaller displays reducing performance by 15% [24]. The only point at which the effectiveness of users dropped significantly was when the number of lines of text displayed was very small (4 lines of text or fewer), and particularly when only one line of text appeared. Line width had a bigger impact – a 25% drop in reading rate observed when the display was reduced to 1/3rd [10] – but, again, the speed of reading remained substantially similar to that on desktop-sized displays.

3.2 Access: Searching and Browsing

A second key component in digital libraries is the access to documents. This is traditionally achieved through two methods: firstly, the discovery of documents using a query; and secondly through the browsing of an organised hierarchy or list, such as topic or author indexes.

To take the latter first, the activity of browsing across and within documents on small displays has also been assessed. Swierenga [27] tested the performance of users choosing commands from a hierarchical list of selections, again with differing screen topologies. When small and large screens were compared, users achieved similar accuracy and speed on each display type.

However, menu and hypertext navigation are not the same. We wondered what the impact of screen size would be on hypertext browsing. We tested the difference when browsing a hypertext system with index, and discovered that task completion times and outcomes were poorer on small screens than on conventional desktop displays [14].

In response to this, we developed an outline-style interaction technique in which the topic tree could be expanded or contracted interactively by the user, before a final document selection was made. On small screens, the overall task completion times fell by 35% and success rates rose [15]. However, there were identifiable performance issues surrounding the structure of the hierarchy – large, 'flat' structures proving poorer than deeper, more balanced ones.

The study of interactive querying on small screens is, unlike the question of reading and browsing, considerably less researched. Some work has been done on adapting desktop visualisation methods (e.g. [26]) to small displays [12], and certain systems have been evaluated in isolation [7]. However, comparative studies of search tools are unavailable.

3.3 Summary

From existing studies of reading, we can be assured that the use of small screens does not mean per se that reading becomes ineffective. However, as Marshall and Ruotolo's recent studies [20] find, screen size may impact the purpose and form of reading done by users – small displays for casual, opportunistic reading, larger displays for more intensive study.

Our knowledge of information seeking on small displays is more varied: browsing on PDA-sized displays is covered, and the results are encouraging, though the use of very small displays requires further study. In the case of search, current scientific understanding is poorly developed.

In this paper, we will present an experiment to assess different presentation methods on small displays, and the outcome of initial studies of browsing and searching on micro-displays, widening the scope of available data.

4 Experimental Systems

In order to gain grounded data on small-screen use in DLs, we have developed two pilot systems built upon the open-source Greenstone DL software:

1) Providing outline-based searching on small screens §4.1

2) Browsing and Searching using WAP devices §4.2

These systems highlight distinct, yet complementary, aspects of small-screen use and usability. In creating them, we have observed some difficulties which may impact on the use of categories and outline presentations in DLs, both on small screens and generally.

The first system evaluates the use of category hierarchies and outliner interaction, successfully used in browsing, in the context of search and highlights some interesting learning effects.

The second implements a general DL system in a context more restricted than the internet, and unveils some cultural difficulties with classification and hierarchies.

4.1 LibTwig – An Outliner-Style DL Browse and Search Tool

LibTwig is a Greenstone-based DL implementation of our WebTwig browsing tool [15]. LibTwig, and an updated WebTwig, have been created to start to systematically evaluate a number of different alternative interaction styles for search tasks on small screens.

The interaction method for browsing is the same, outliner, style as WebTwig: the user browses over a hierarchical index, and can expand a category to reveal its component documents and sub-categories, or close it to leave just it's own title visible. Within a category, documents are listed in Greenstone's default, alphabetic, ordering. LibTwig also supports searching, with results given in one of two presentations: a reduced outline hierarchy containing only items and branches which match the search; or alternatively a traditional ranked list. In the case of the outline mode, the ordering of documents within categories is the same, alphabetical, order as when browsing.

If you refer to figure 1, the outline presentation, with categories partly (left) and fully (centre) expanded, can be compared to the ranked list display (right), with documents ranked by relevance, the most relevant at the top.

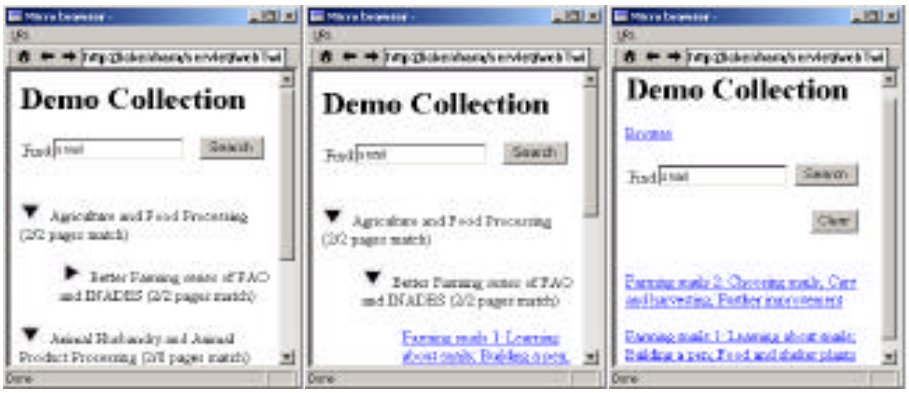

Fig. 1. LibTwig in use on the Greenstone Demo Collection. Left – Outline search; Centre – Outline search with category expanded and a document link visible; Right – Ranked List. The browser window is fixed to Pocket-PC display capacity.

Our initial study of patterns of information seeking by users of small-display devices [14], demonstrated that interactive search was the initial information seeking method for 80% of users – twice the rate on desktop displays. Thus, search is particularly important on the small-screen [4].

The use of an outline-style presentation for browsing tasks significantly benefited those using small displays, reducing task completion times and increasing success rates [15]. These benefits also occurred in a smaller degree when using desktop displays. Later alternative implementations of outline browsing on small screens, e.g. by Buyukkokten [6], have demonstrated similar results.

Given the benefits of an outline access method for browsing tasks, it is reasonable to hypothesise that such a presentation may also benefit interacting with search

results. The Cluster hypothesis [23] predicts that most "real" matches for a search will be in a common classifier, so a corollary is that matches not in that category are probably less relevant. Therefore, the thematic division of results by subject category can improve the effective selectivity of the document review task [25].

With LibTwig, we have carried out a pilot study to evaluate the relative merits of a naïvely implemented version of the outline presentation of search results compared to the traditional ranked-list. Experimental systems which use some form of the outline search presentation exist. Some have been presented at DL conferences (e.g. [21]). However, usability evaluations of these have not been widely available.

Design and Implementation

LibTwig communicates with Greenstone through Greenstone's internal CORBA protocol [1], and utilises Greenstone's search facility. Greenstone is entirely a 'full text' digital library, and does not use 'stop words'. This affected our implementation decisions, as we shall see.

There are a number of alternative implementations of the outline search presentation. A significant problem which could be expected may be the loss of relevance information through the loss of ranking. It is well proven that ranked lists are more effective for users than other orders, such as alphabetic, e.g. [2]. Whilst preserving the structure of the hierarchy may give additional information to the user, it is unclear whether the ordering of documents within a category in the hierarchy should be as normal (alphabetic in the case of LibTwig and WebTwig), or ranked instead. Furthermore, how to give feedback as to the number of matching documents in a node or its children is an open question, with a number of sub-problems and alternative solutions. Finally, how to match documents (requiring all or merely some terms in a document) is also unclear, or whether using some heuristic cut on the score of a normal ranked search would be appropriate.

Our 'naïve' approach was to provide as similar a selection of documents as possible as would be the case in ranked list. We anticipated that the use of very common terms, e.g. "the", would retrieve many documents, as Greenstone's search engine does not use 'stop words'. The absence of stop words has little or no effect on the important head of relevance-ranked document lists, usually used by Greenstone, and so normally the lack of stop words has no effect. However, in the case of the 'outline' mode of presentation, relevance ranking is not used, and thus the negative effects of high retrieval would not be palliated. Therefore, to remedy this, search results were refined with a 'cut' on documents with extremely low scores, effectively reintroducing a basic level of 'stop word' provision.

The total number of matching documents within a category was given beside its name in the hierarchy. Documents were presented in the default, alphabetic, order, rather than ranked within each category by relevance score. Whilst the combined effect of these decisions was expected to be sub-optimal, the intent was to secure data towards the lower bound of performance when compared to ranked search.

Method

A panel of 12 subjects was recruited to perform an initial evaluation. All the subjects were regular library users, using a physical library once a month or more. Similarly,

all were computer-literate, either owned or regularly used a mobile phone, and had some degree of experience of using the Web. A range of levels of web experience and library use were recruited. Ages ranged from 19 to 47, and the subjects were a mixture of students and staff in the department of Computing Science, Middlesex University. The material to be searched was selected as being of interest or potential interest, all subjects either studying of teaching the theme of the material (human-computer interaction); again, a range of expertise was recruited. Subjects were screened in advance to ensure a balance of subject expertise across sets of questions.

The collections used were built by end-users, not by expert librarians. Greenstone is in widespread use and many collections are not created by professional indexers. It is possible that using professionally indexed material would suggest benefits which were more related to the expertise of the librarian than the subject of our focus: the access method itself.

Subjects were given an initial training, using both the search presentations, and were then permitted an open-ended familiarisation period with each presentation style. Pre- and post- experiment questionnaires were taken.

A fixed panel of 10 questions was answered by each subject, five with each interface, and with the combination of question and interface balanced across subjects. The order of questions was also randomised to reduce any ordering effect. As with our previous experiments, some questions were more directive, others open ended. All questions had a set of appropriate target documents selected by a subject expert, of which the subjects were unaware. Subjects were invited to select appropriate documents, and had an open-ended time period to complete each question. The subject could choose to move onto another question at any time, though they could not later return to a question. Subjects were asked to give their opinion as to their degree of success on each question. Timings were taken of their performance, as was the number of target documents which they viewed, using a bespoke browser based on Internet Explorer which we have used in previous published experiments.

Results

Our first comparisons were of the time performance of users on each result presentation. Given the small sample size, and the pilot purpose of our study, results are generally indicative rather than inferential, and are seldom statistically significant.

Overall average times were within 10% of each other. Considering only those cases where the subject believed that they had succeeded, average times are virtually inseparable. As we have found previously, failure cases took significantly longer than successful tasks [17].

	Presentation Mode	
Question Set	Ranked List	Outline
All questions	201	220
Answered questions	170	168
Failed questions	380	306

Table 1. Time comparisons for Ranked List and Outline presentations

The small sample size and natural variation in timings between subjects and presentations does not account for the lack of statistical significance, however. Differences between two apparent sets of users contributed substantially to this, and identifying these two groups assists in identifying problems in the outline presentation.

Differing User Behaviours

Early in our analysis, a significant affect was observed from the frequency of web use rather than the length of web experience. Refer to Figure 2 below.

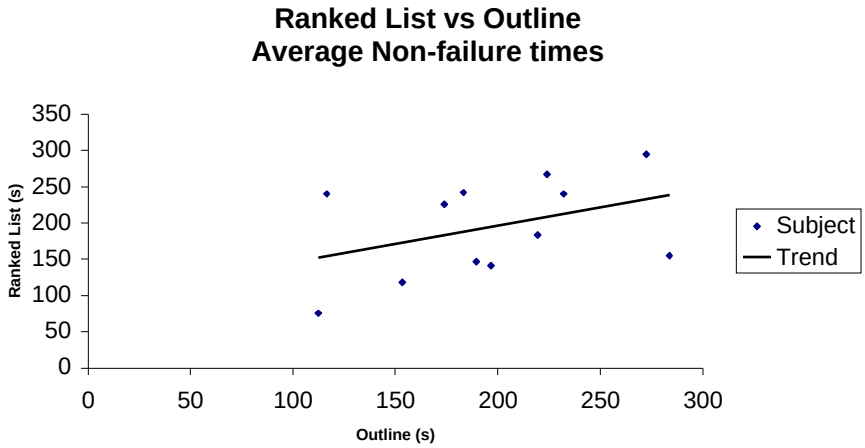

Fig. 2. Ranked List versus Outline time performance – successful and uncertain outcomes.

The average overall trend is plotted as a line – those who took longer with one method, also took longer with another. Six users, below the line, performed much better with ranked list than outline presentation – they all used the web for 10 or more hours per week (self-reported). Those above the line had comparable times in each mode, but average web use was 5 hours/week. The length of web experience was, however, similar – averaging four years (given the self-reporting here, further accuracy is misleading).

The intensive web users averaged 170 secs/question in ranked list mode. Their superior performance in this mode is perhaps unsurprising, as they may be more experienced with the presentation. However, they also experienced fewer failures – just four over all modes and questions. In outline mode, their task times were close to the non-intensive users, averaging 233 secs/question. Whilst completing the post-test questionnaire, three intensive users reported confusion over the difference between the two modes. Even permitting these users to use the tools again after the experiment did not clarify their understanding of the new outline presentation. When asked which presentation they preferred, most of the intensive users selected ranked list mode (4 versus 1, plus 1 abstention).

The non-intensive web users were slower in ranked list mode than outline mode – 244s against 218 secs/question. Furthermore, when successful in the outline presentation, they were faster than the intensive users – taking 188 seconds on average against 255 seconds for successful intensive users. Failures were much more common – 19 in all (nearly 33%) and the better speed achieved in outline mode was bought at a cost – 13 failures occurring in that presentation, twice as many as in ranked list mode. However, inspection of search terms used by the subjects illustrated that these users were much more prone to using stop words or very common words, resulting in larger sets of matching documents. As we shall see, common terms prove a problem even for the intensive users.

In the post-test questionnaire, four of the non-intensive group reported that the outline mode helped them find documents and that they got a "better picture" of the collection organisation and found related documents. The non-intensive web users unanimously preferred outline mode, though three reported problems understanding category headings.

Across these two groups, rates of reformulation (1.4 searches/question) and page viewing (2.8 documents/question) were virtually identical. Excluding the higher use of very common words by the non-intensive user group, query terms were also broadly similar. Commonplace words little affected the ordering of documents in ranked list mode. So, behavioural artefacts did not reflect differences in ranked list performance – perhaps indicating that intensive users took better advantage of the presentation.

General Presentation Differences

Considering the two presentations, failure rates differed from seven (12%) in ranked list mode to sixteen (27%) in outline mode. However, five failures in outline mode occurred on one question (two subjects being intensive users, and three non-intensive) – only one, intensive, user managed to successfully answer the question. All users used "web" or "website" in this search, on a collection about web usability, so these terms were poor distinguishers. In ranked list mode, common words had a minimal affect on the ordering of the top documents, and it is notable that of those using ranked list mode, all (three each of the intensive and non-intensive groups) successfully answered the question. On this one question in outline mode, 90% of documents matched, and one large category had over 200 matches. This poor selectivity degraded the task to browsing.

For every other question the rate of success varied by only one subject between the presentations. Overall, though significance can be gained on the failure rate at the 5% level, removing the outlier also results in a loss of significance, so the global measure cannot be unreservedly trusted.

In the post-test questionnaire, six subjects identified the (alphabetic) ordering of titles in Outline method as a problem, and asked for the documents within a category to be "sorted". Four subjects reported problems with (original author) document titles.

Evaluation

Given the naïve implementation of Outline mode, we expected it to be considerably poorer than ranked list mode. The widespread use of the ranked list presentation also

meant that any novel method would be disadvantaged at first comparison, and a brief training was unlikely to counterbalance this. However, overall the outline presentation compared well, particularly for the non-intensive web users, with whom it was popular.

Clearly the current design can be improved – ranking within categories when searching as suggested by the subjects – and focussing attention on fewer, more relevant documents. The good success rate of experienced users with outline mode showed that outline mode need not impede successful searching.

In the case of the problem question noted above, normally modestly selective words ("web" and "website") were common in the context of the collection being searched. This clearly impacted intensive users (half of all their failures occurring on this one question). Our non-intensive users, using more common terms generally, no doubt suffered their higher failure rate in outline mode as a consequence of the same underlying problem.

Thus, improved precision should improve performance of the outline mode for all groups.

However, in the problem question, the hierarchical structure of the collection was also a problem – one category, in which a number of the best matches occurred, had two hundred documents in it (excluding sub-categories). We have previously observed a drop in performance of the use of hierarchies in browsing tasks when such broad, flat structures occur [15], so that the same phenomenon reappearing in browsing must clearly also be a factor. Though in this case, an improved search method would have reduced the number of items in that category considerably, this may not always be the case. Thus, using Dumais and Chen's approach of only displaying the top five hits (or some other reduced selection) initially in a category, and permitting the user to request to see all the matches, may provide improved performance [8].

Previous experimental work with outline-style presentation of search results has given variable outcomes; for instance Chen and Dumais [8] report an improvement when compared to ranked list displays, whereas the evaluation of the use of clustering systems for accessing search results, e.g. [13], is much less encouraging. In the case of Dumais and Chen's study, the subjects in that case are described as "intermediate" users – and their profile is similar to our subjects whose web use was moderate. Dumais and Chen's system also contains features, such as ranking by relevance within each category, which may well improve on our current performance.

The benefit of topical browsing noted by four of our subjects echoes similar comments by users in a study of a system for clustering search results by Hearst [13].

Thus, there is clearly reason to develop the outline presentation further so that a more certain and precise comparison of an improved method can made with ranked list results. Comparative studies using standard desktop-sized displays also need to be undertaken.

Key questions surround the intensive web users. Clearly, their existing skills did not translate well to the outline method, and the clear confusion of half of this group is perplexing. Can their lack of comprehension be improved? If not, then will performance with the outline mode be permanently affected? If such effects are emerging, there are significant impacts on the future direction of information seeking research – alternative presentations finding adoption harder.

4.2 DL Access via WAP

To work, an outliner based interface into a digital library obviously requires understanding of both libraries and hierarchical document structures. Our research in South Africa has shown that this understanding may not be taken for granted.

Providing people in developing countries with information is difficult – books are relatively expensive and distribution is problematic. However, many of these nations have highly developed cellular networks – e.g. 22% of South Africa's population have a cellular handset, yet only 11% have sufficiently high wages to pay income tax [18]. Within SA, we therefore undertook a study to see if we could provide DL access using WAP technology.

A second challenging issue is the general lack of library experience in South Africa. Studies of other groups in which library exposure is low, such as the Maori of New Zealand [11], observe negative impacts on DL usability.

The system used the same outline-style access as described in the section on LibTwig above. Nodes in the hierarchy could be expanded and contracted just as in an outliner. Indentation was used to emphasise the tree structure visually. The hierarchies of the library and documents were used as one continuous hierarchy, so within a category, expanding a document would list its chapters, expanding a chapter would give its sections, etc., until actual body text was revealed.

The system's usability was assessed through a series of evaluations involving typical end-users as study subjects with handheld (small- rather than micro-display) devices. See Fig. 3. below for an overall impression of the system in use.

Fig.3. Greenstone's WAP interface in use through a Palm-OS Wireless Simulator – used here for image clarity.

As expected, there were some trivial usability problems which would be expected on almost any WAP-based system. One particular difficulty was users not identifying when scrolling was possible. As result, often only the first part of a longer document would be read. When questioned about why they had not scrolled, subjects reported being unaware that any more information was available. In part this appeared to be due to the small size of the scroll bars presented by WAP systems; the level of

visibility is too low, more or less eliminating feedback to the user. However, this also correlates to difficulties observed when subjects were using large-screen devices with a variety of DL systems [3] – even subjects with extensive web and application experience repeatedly failed to scroll down, and repeatedly failed to observe that scrolling was possible. It would seem that a key component of DL usability, particularly on smaller screen devices, is the requirement for the browser to better support reader awareness that scrolling is possible.

On a more profound level, subjects also seemed to struggle with the concept of hierarchical access, contrary to our previous experience of the use of the outline-browsing method of access. The concept of a strict hierarchy seems alien to some cultures as shown by further studies wherein our target user group were unable to draw simple hierarchies such as a family tree [28]. This echoes findings of cultural difficulties in DLs amongst the Maori [11].

Besides the problems in understanding general hierarchical organisations, we conducted a further set of studies with the full Greenstone system on a desktop computer to see if there were problems unique to Greenstone.

Again, browsing was a problem, there appeared to be problems with even basic metaphors and structures such as indexes, chapters, sections. . Many of our subjects had never been to a library before attending university and the distinction between sections and chapters was lost on them. Consequently, they were confused by the behaviour of the interface as sometimes clicking on an icon (e.g.) chapter would give them text to read, and at other times (e.g. collection) would give them more icons.

Finally, when clicking on a heading that contained only body text, users clearly expected the hierarchy to disappear, and the content text only to appear, an expectation perhaps related to their experience on web sites.

This same phenomenon occurred in a slightly different form when searching. An individual result of a search, partially due to Greenstone's full text search facility alluded to earlier, could be anything from an entire document (many individual sections matching), down to a single section (no other matches occurring in another section of the same document). Clicking on a result could result in a variety of responses, from a list of chapter headings (where the result is an entire document), to body text (the result was a section). Again, users expected body text only.

Clearly, there are significant questions as to which issues here are related to small screens, or are particularly acute on small screens, or those which are, on the other hand, cultural issues. Overall, there was surprising symmetry between usability issues on the small and large screens.

Further results of this work are reported in [16].

5 Conclusion and Future Work

In implementing and evaluating these two tools, some common themes have emerged, particularly in regard to the use of hierarchies and outliners. Outliner interactions have provably improved small-screen browsing, and are candidate forms of search access too. However, cultural incomprehensibility and (in the context of searching) learned expectations, may limit or eliminate the benefits of hierarchies.

In the case of WAP access, work first needs to be done to address the cultural issues uncovered first before re-evaluation at Cape Town, whilst the existing system is evaluated in a context where users are more familiar with library metaphors.

The outline presentation of search needs improvements which should make comparison to ranked presentations clearer. Knowledge of the possible training effects of web use can now better inform our experimental methods. The impact of hierarchical search result presentation on desktop displays also needs pursuing, within its own merits and as a comparison against small-screen use. However, the strong preconceptions of search held by advanced users, and their strongly developed skills with ranked list presentation, may result in the outline presentation being less effective for them.

References

1. Bainbridge, D., Buchanan, G., McPherson, J., Jones, S., Mahoui, A., Witten, I. : Greenstone as a Platform for Distributed Digital Library Publications. Proc. European Conference on Digital Libraries, 137–148. Springer-Verlag (2001).
2. Belkin, N.J. et al. Using relevance feedback and ranking in interactive searching. In Harman, D. (ed.) TREC-4 Proceedings of Fourth Text Retrieval Conference. Washington, D.C., 181-209 (1996).
3. Blandford, A., Stelmaszewska, H., Bryan-Kinns, N.: Use of multiple digital libraries: a case study. JCDL 2001: 179-188
4. Buchanan G. & Jones, M. (2000). Search interfaces for handheld mobile devices. Poster proceedings 9th International Conference on the World Wide Web.
5. Buchanan, G., Farrant, S., Jones, M., Thimbleby, H., Marsden, G., Pazzani, M.J.: Improving mobile internet usability. WWW 2001:673-680
6. Buyukkokten, O., Garcia-Molina, H., Paepcke, A. & Winograd, T. (2000). Power browser: efficient Web browsing for PDAs. Proc. ACM CHI 2000, pp 430 – 437.
7. Buyukkokten, O., Garcia-Molina, H. & Paepcke, A., Focused web searching with PDAs. Proc. of WWW 9 (2000). pp213-230.
8. Chen, H., Dumais, S.T.: Bringing order to the Web: automatically categorizing search results. Proc. Of ACM CHI 2000: pp 145-152.
9. Dillon, A., Richardson, J. and McKnight, C. (1990). The effect of display size and text splitting on reading lengthy text from the screen, Behaviour and Information Technology, 9(3):215–227.
10. Duchnicky, R. L. and Kolers, P. A., (1983). Readability of text scrolled on visual display terminals as a function of window size, Human Factors, 25:683–692.
11. Duncker, E., Cross-cultural Usability of the Library Metaphor, Joint Conference on Digital Libraries 2002, in press.
12. Dunlop, M. D. & Davidson, N. (2000) Visual information seeking on palmtop devices. Vol. II Proc. BCS HCI 2000, 19-20.
13. Hearst, M. A. and Pedersen J. O. Reexamining the Cluster Hypothesis: Scatter/Gather on Retrieval Results. Proc. 19th ACM SIGIR Conference on Research and Development in Information Retrieval. 1996. 76-84
14. Jones, M, Marsden, G., Mohd-Nasir, N, Boone, K, & Buchanan, G. (1999) Improving web interaction in small screen displays. Proceedings of Web 8 conference, 51–59
15. Jones, M, Mohd-Nasir, N & Buchanan, G (1999). Evaluation of WebTwig — a site outliner for handheld Web access. International Symposium on Handheld and Ubiquitous

Computing, Karlsrhue, Germany. Gellerson, H-W (Ed.), Lecture Notes in Computer Science 1707:343–345, Springer-Verlag.
16. Jones, M., Buchanan, G., Thimbleby, H., Sorting out Searching on Small Screen Devices, Conference on Mobile HCI, Sept. 2002, In press.
17. Marsden, G., Cherry, R. & Hafele, A. Small Screen Access to Digital Libraries. 2nd South African Conference on Human-Computer Interaction (CHI-SA2001). South Africa. (Electronic proceedings).
18. Marsden, G.: Subverting Technology: Meeting User Needs in a Developing Economy, Social Issues Column, SIGCHI Bulletin, March/April 2002, pp 8.
19. Marshall, C.C., Price, M.N., Golovchinsky, G., and Schilit, B.N. (1999) Introducing a digital library reading appliance into a reading group. Proc. ACM DL'99 (Berkeley, CA, August 11-14), pp. 77-84.
20. Marshall, C.C. and Ruotolo, C. Reading-in-the-Small: a study of reading on small form factor devices. To appear in Proceedings of the Joint IEEE and ACM Conference on Digital Libraries (JCDL02), Portland, Oregon, July 14-18, 2002.
21. Palmer, C.R., et al,: Hierarchical Document Clustering of Digital Library Retrieval Results, Proceedings of the Joint Conference on Digital Libraries, 2001, p.451.
22. Paynter, G.W., Witten, I.H., Cunningham, S.J. and Buchanan, G. (2000): Scalable browsing for large collections: a case study. Proc Fifth ACM Conference on Digital Libraries, San Antonio, TX, pp. 215—223.
23. C.J. van Rijsbergen: Information Retrieval. Butterworths, London, second edition, 1979.
24. Shneiderman, B.,: User Interface Design and Evaluation for an Electronic Encyclopaedia: In Salvendy, G. (Ed.), Cognitive Engineering in the Design of Human-Computer Interaction and Expert Systems, Elsevier Science, 1987, pp 207-223.
25. Shneiderman, B., Byrd, D. & Croft, B (1998). Sorting out searching. Communications of the ACM, 41(4):95–98.
26. Shneiderman, B., Feldman, D., Rose, A., and Grau, X.F. Visualizing Digital Library Search Results with Categorical and Hierarchical Axes. in Proceedings of Digital Libraries 2000 (San Antonio TX, June 2000), ACM, 57-65.
27. Swierenga, S. J. (1990) Menuing and scrolling as alternative information access techniques for computer systems: interfacing with the user, Proc. Human Factors Society, 34th Annual Meeting, Vol. 1:356–359.
28. Walton, M., Marsden, G., Vukovic, V.: Visual Literacy as Challenge to the Internationalisation of Interfaces: A Study of South African Student web users, ACM CHI 2002 Extended Abstracts pp350-351
29. Witten, I. H., Nevill-Manning, C.G., McNab, R., & Cunningham, S.J.: (1998) A public digital library based on full-text retrieval: collections and experience, Communications of the ACM, 41(4). 71–75.