Document Resizing for Visually Impaired Students

Michael Connolly
Dept of Software Engineering
University of Auckland
Private Bag 92119
Auckland, New Zealand
mcon059@aucklanduni.ac.nz

Christof Lutteroth
Dept of Computer Science
University of Auckland
Private Bag 92119
Auckland, New Zealand
lutteroth@cs.auckland.ac.nz

Beryl Plimmer
Dept of Computer Science
University of Auckland
Private Bag 92119
Auckland, New Zealand
beryl@cs.auckland.ac.nz

Abstract
The ability to read documents and notes is a crucial part of the education system, but for over 1200 visually impaired students in New Zealand and many more worldwide, large and clearly printed documents remain elusive. Resizing documents for visually impaired readers currently requires a mixture of time, patience and experience with word processors such as Microsoft Word. This paper describes the design and construction of an add-in to simplify the process of resizing documents so that they become more readable to the visually impaired. This paper discusses common problems with the resizing of documents, and the tools produced to help reduce or eliminate these problems. The tools were evaluated in the resizing of workbooks by staff at a visual resource centre with promising results.

Keywords
Document Resizing, Visually Impaired

INTRODUCTION
In 2005 there were approximately 1200 visually impaired students in the New Zealand education system that required extra support to help them at school (Petty, 2005). Many of these students are unable to, or struggle with, reading worksheets, workbooks and other forms of documentation. Such material is a core educational resource in many New Zealand schools, and generally teachers share resources to assemble the workbooks for their courses.

The simplistic solution for this problem is to enlarge the document by photocopying it onto a larger piece of paper - often A3. However this has 3 major disadvantages, first it makes reading uncomfortable, wide columns require the user to turn their head to read each line – people with impaired vision often have narrow fields of focus, exacerbating this problem. Secondly, each element is enlarged including blank space; this fails to take advantage of the blank space and therefore uses large amounts of paper for each document. The final disadvantage is portability: documents larger than A4 are difficult to carry; this is particularly a problem for students who are expected to carry large amounts of handouts from class to class.

A better solution is to resize the document content. This is done by increasing the size of the elements contained within the document while maintaining the logical relationship between the elements. At the same time fonts that are difficult to read can be replaced. Today most documents are written and formatted within a word processor. These word processors provide the user powerful functionality for formatting but provide very limited information about the logical relationship between each component in the document. Often documents are considered 'well formatted' if every component appears in the appropriate position. However it is rarely the case that the logical relationships are also so well organised, and as a result changes to the document font size can result in the user having to reformat the entire document. Our analysis of the documents that the visual resource centre of a secondary school has given us to reformat shows that they use a very diverse range of logical relationships to create similar visual layouts.

In order to automatically resize a document the logical relationships must be established between each of the different components, and then each component must be resized and repositioned to ensure the graphical representation reflects the logical representation. There is currently no software that is able to generate and manage logical relationships. Furthermore, in order to establish logical relationships the document must be initially well formed, and constructed using a known template. The diversity of logical representations we found in documents can be explained by the process by which these documents are typically created: teachers cut-and-paste a variety of resources from different documents together to suit their teaching goals, hence is unlikely for the documents to be consistently formed or well constructed. We have found that it is unlikely that a complete set of logical relationships can be generated from documents at a level where it is suitable for mainstream use. Instead the best idea is to rely on the teacher or the person resizing the document to work out these relationships.

OZCHI 2010, November 22-26, 2010, Brisbane, Australia.

OZCHI 2010 Proceedings ISBN: 978-1-4503-0502-0

Section 2 discusses related work in the field of document resizing and various printing standards amongst visual impairment organizations around the world. Section 3 discusses the requirements for software to simplify the process of resizing documents. Sections 4 and 5 describe the interaction design and functionality of the software. Finally, Section 6 describes how the software was evaluated.

RELATED WORK

Since research into document resizing is rarely undertaken, concepts from different areas of visual impairment studies need to be combined to produce a set of standards for resized documents. There are many assistance organizations for the visually impaired, which publish professional printing standards for visually impaired readers. Although a 1993 study in Australia showed that 0.4% of the population of under 15-year-olds are visually impaired (Castles, 1993) there are no global standards for large print documents. This section reviews selected printing standards and discusses the differences between them. It also discusses research from the related fields of document layout analysis and layout management.

The Association of Directors of Social Services (Association of Directors of Social Services, 2002) suggested minimum font size of 14. A simple and clear layout was recommended – complex structure increase the reading difficulties for visually impaired readers. A sans-serif typeface, e.g. Arial and Helvetica, is recommended.

A specialist from the California Transcribers & Educators of the Visually Handicapped organization suggests that the minimum font size should be 18 (Hudson-Miller, 2007). The suggestion on images is interesting: relevant images may not be separated into several smaller ones, but less important images can be divided. In the process, the loss of print quality is forbidden. It is also suggested retaining the pagination of the document, as it is helpful for readers to find content. MacNeill suggests that contrasting colours are preferred by visually impaired readers (MacNeill, 2006). This author also recommends a minimum font size of 16 points.

From those papers, it is clear the main directions of print standards from different originations are similar: large print and less confusion. However, they have different approaches and the details are very different. One example is the minimum font size, varying between 14 and 18 points, which might be due to the fact that the definition of 'visually impaired' is different in different organizations. Since a print standard needs to be created for this project, these standards need to be combined and aligned with the NZRFB (New Zealand Royal Federation for the Blind) suggestions, which are applied widely in New Zealand.

The PDF document format is widely used. Structural analysis of PDF documents as some similarities to document resizing. There are many approaches to extract logical structures from PDF files and convert the PDF file to other formats. For example, Déjean and Meunier (2006) presented a system to convert a PDF document to a structured XML document. Converting a PDF file to the XML format requires categorizing each document element and capturing the recognition order among the elements (Hadjar, Rigamonti, Lalanne, & Ingold, 2004). By converting a PDF document into XML, the logical structure of the PDF is recovered, since the tags and hierachy of XML reflect the logical structure (Mao & Kangungo, 2001; Vanderdonckt, Bouillon, & Souchon, 2001). Both the logical structure and the recognition order are important information in the resizing process, hence similar approaches may be used to support document resizing.

Document layout analysis is used to guide the process of converting a document from an image to a textual electronic version (Breuel, 2003; Mao & Kangungo, 2001; Déjean & Meunier, 2007). When working with a bitmap image of a document, the computer stores and processes the document as a series of pixels, with no distinction between the paragraphs, figures and tables inside the image. In order to convert those pixels into an electronic document, different parts such as paragraphs must be identified. This is accomplished by document layout analysis (Breuel, 2002), which analyses the layout structure of a document, considering only its visual appearance. This technique is similar to the process of document resizing, where the document is resized following a set of guidelines based on the relationships of the elements in the document.

Document layout management focuses on using page space efficiently. Document layout management systems are able to automatically reformat, resize and paginate electronic text and graphics for different media, so that a document looks as good as when diplayed on its orginal medium (Jacobs, Li, Schrier, Bargeron, & Salesin, 2004). Document layout managment provides information and ideas on managaging the document elements when changing the size of their medium. These systems typically require that the logical structure of a document is given as an input. However, this is typically not the case for document resizing.

GUI layout management focuses on managing the layout of widgets in a changeable graphical user interface (GUI). Modern GUIs allow the user to resize the GUI to accommodate a variety of different needs and screen resolutions. When a GUI is resized, the positions and sizes of the widgets need to be adjusted. This adjustment is typically done with the help of a constraint-based layout specification (Lutteroth, Strandh, & Weber, 2008). For document resizing, we are typically not given such a specification, therefore it is much harder to automate. However, there are approaches for recovering a layout specification from a GUI that does not use a layout manager, e.g. (Lutteroth, 2008). Similarly, reverse engineering could also be applied to recover a logical structure for a document. Since the reverse engineering has to rely on heuristics, there cannot be a guarantee that it produces accurate results.

REQUIREMENTS

Due to the hundreds of different document styles and layouts as well as the differences in the visual representation of the documents internal data model, we have found that full automation of document resizing is impractical and would not be reliable. Previous fully automated solutions that we considered are slow and not reliable enough for variety of different document types (Zhao, 2009). Through testing we found using a semi-automated approach, which can rely on the user's judgment of the appearance, is a more practical and reliable solution.

After working closely with staff at a visual resource centre of a secondary school, the following document resizing functions were agreed upon:

- Break down large documents into smaller documents
- Insert page numbers automatically
- Change the orientation of pages easily
- Change the margin settings easily
- Aid in the selection of related objects
- Increase the font size of all text and change its formatting
- Replace certain types of formatting with others
- Automatic removal of columns

Each of these points is discussed in detail below. Note that all of the documents we encountered were created and resized using Microsoft Word, hence the Word user interface serves as our benchmark when trying to improve the document resizing process.

Document Splitting and Joining

Resizing Documents can, at times, be extremely difficult. The major problem is the displacement of some of the document elements when the resizing seemingly unrelated sections. As a document is resized, it is logical for certain areas to take up extra space. This extra space causes following parts of the document to move down. If graphics are not anchored to the correct paragraph, this can cause significant displacement of graphic elements, leading to strong distortions of the original layout. This is best illustrated in Figure 1 where the rectangle and the highlighted paragraph are closely related logically and physically. However, when the page is resized, the rectangle and the paragraph are on separate pages because the rectangle was anchored to the page rather than the paragraph. The default anchor positions set by Word, although often logically incorrect, are usually not changed by users because they are concerned only with the visual appearance.

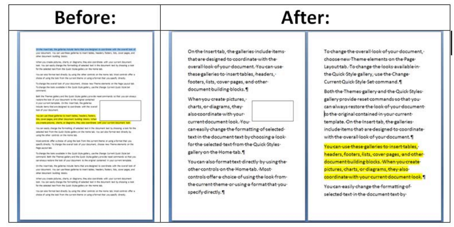

Figure 1 - Incorrect Anchoring

In small documents these issues do not cause any major time delays. But in large workbooks, finding the correct paragraph to relocate the diagram back to could require the reading of dozens of pages. Often these small inconveniences build up, cascading through the document until later sections of the document are heavily distorted. There is no simple solution for this problem. We have thoroughly explored automatic and manual approaches to building a correct semantic document structure - the complexity of the word object model, including its capability to nest objects (which may be anchored to different things) make it almost impossible to cover all eventualities.

The visual resource centre staff have two approaches to resizing a difficult section of a document: recreate this section of the document from the paper version, or extract the section from the original, deal with it separately, and then merge it back in once enlarged. The latter is a practical solution to limit the aforementioned problem of cascading distortions during resizing. It can be done by splitting the document up into many smaller documents, and resizing each section independently. This limits the amount of distortion each section's resize can cause since each section is isolated. Once each section has been enlarged, the document can be rejoined. The rejoined version is the composite of all the resized sections, and is not affected by the distortions of the resizing process.

Page Numbers

It is common for documents to hold references to certain pages. However, when a document is resized these page references are no longer correct. Replacing these page references with new ones is not useful in the education system, as teachers and students in a class will make references to the original version. Instead the original page number should be kept. Currently, staff resizing a document need to go through the document page by page, writing the original page number on the top of each page so that students can navigate through a resized booklet using the references of the original booklet.

Page Rotation

When resizing documents in order to maximize the space available for elements such as tables and some images, it is common practice to rotate pages. In Microsoft Word 2007, rotating specific pages can be complicated for non-experienced users. Figure 2 shows the 6-step process required to rotate a single, specific page. As the diagram shows, the user must first select the text, navigate to the Page Layout Tab, access the Extra Page setup options, change the orientation to Landscape and change the apply to section to selected text. This process is slow, not very intuitive, and requires experience with Word. Simplifying this process will save time and allow Word novices to rotate pages.

Figure 2 - Rotate a single page in Word 2007

Page Margins

Although the process of changing margins in Microsoft Word has been simplified in the 2007 edition, it still offers very little support for those who want to mix and match custom margin settings. When workbooks are resized, the margins should be changed to accommodate the binding on the side while optimizing the space available to the content. Although changing margins to one of the preset settings is quick and easy, setting exact or custom document margins requires manual input of measurements, as shown in Figure 3. Providing a set of margins that are optimized for document resizing will save time compared to inputting margin settings manually.

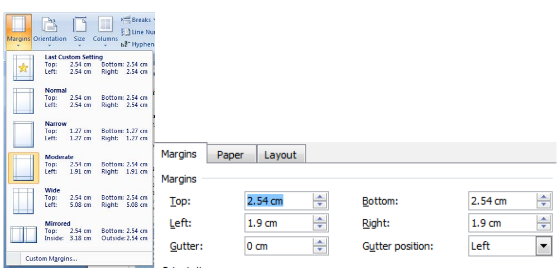

Figure 3 - Margin Settings in Word 2007

Related Object Selection

Most modern word processors allow graphics to be inserted in one form or another. Word supports many different forms of graphics, and when resizing documents it is important the graphics are resized correctly as well. Graphics should be moved above the paragraph they are related to, and if the graphic contains fine details then it should be moved onto a separate page and enlarged. One type of graphic that can be particularly difficult to resize is diagrams. Diagrams can be assembled from multiple components and are difficult to resize, especially if the components are not grouped. Figure 4 shows a diagram that needs to be resized, and all the resize handles in the user interface. In order to enlarge this shape, all components must be selected and then uniformly increased. Failing to select even one component results in a distorted diagram, as shown in Figure 5. Simplifying the process of selecting complex structures and grouping them so only one selection point is required will reduce the likelihood of incorrect resizing of diagrams and speed up the resizing process.

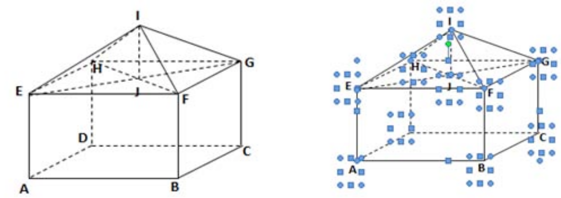

Figure 4 - A Diagram and all its resizing handles

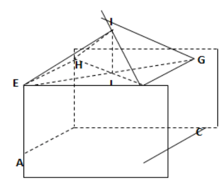

Figure 5 - Failed resizing of complex shape

Text Resizing

Font size, shape and colour are all important for the visually impaired. Some fonts and types of formatting can be difficult to read. Figure 6 shows four different fonts with the characters 1 L ! and I. It shows that some fonts make it incredibly difficult to distinguish between certain characters. Currently, when resizing a document the text must be selected and the user must navigate between various menus to select fonts and resize the text. This repetition of a simple task is time consuming, and also requires the user to know where advanced font options are located. Automating text resizing will free up the user's time and reduce the knowledge required of font styles that are difficult for the visually impaired to read.

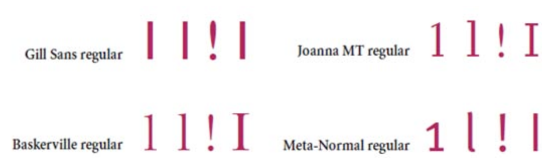

Figure 6 - Differences in fonts

Font Effects Conversion

In many workbooks and exams it is common for specific text effects to be used to emphasis parts of the text (bold, italic, etc). This can be problematic for document resizing, as often these effects are difficult for visually impaired students to read. Currently, when resizing a document, the user must manually find and replace all instances of effects, and then check if the text still makes sense. Producing an automatic method of finding and replacing effects will simplify the process and save time.

Column Removal

Columns can be great for allowing large quantities of text to be visible on one page. However, when columns are used the text size used is often very small. When resizing documents, columns should be removed because they can make the resized document difficult to read. This is illustrated in Figure 7 where the word "photosynthesis" is split between multiple lines. It can also be difficult for the visually impaired to distinguish the different columns, causing sentences to appear unintelligible. A specific tool for removing columns of text would simplify this relatively frequent operation.

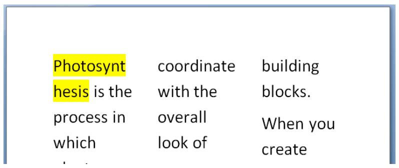

Figure 7 – Columns

DESIGN AND IMPLEMENTATION

Microsoft Office is the most widely used word processor in the world, and the visual resource centres of educational institutions are no exception. Due to this popularity and the many functions already built into Word, the tools need to be accessible from within Word (preferably a newer version such as 2007) to speed up the process.

Since the main requirement is to simplify the process of resizing documents, all functionality should require minimal configuration and be usable without the need to navigate between multiple dialogues or option menus. All functionality should be quickly and easily accessible from within the Word user interface. In order to ensure that the tools are simple and easy to use, the controls for the user Once the document has been split into multiple document parts and each part has been successfully resized, it is necessary to join the parts back together in the correct order. Joining the parts together requires several steps. First, a new empty document needs to be created. All the content from the document parts will be moved to this new document. Each document part is then opened, and the final line of the new document is checked to see whether it contains a section break or a page break. If no section or page break exists, then a section break is added to the final line. By adding this section break and copying the section break along with the content, the entire section and the page need to be kept to a familiar style. Word 2007 and 2010 use a ribbon and tab interface, where controls are grouped within tabs and then displayed on ribbons. Each set of tools is given its own group within a new 'Document Resizing" ribbon, as shown in **Error! Reference source not found.**8, thus organizing the interface so that related parts are close together.

Document Splitting and Joining

Word is a word processor and not a desktop publisher. Because of this, Word logically organizes documents in sections and not pages; page layout information is stored within the section, not within a page. This needs to be taken into consideration when splitting documents into smaller parts. Copying each page does not always copy the page layout: Figure 9 shows a page from a math's exam that has been copied and pasted into a new document. As shown, the content is copied over correctly and the relationships between each section is maintained, however, the page margins on the new document are different from those of the original, and this has resulted in the document overflowing onto a second page. This occurs because copying a page's content does not automatically copy its section information; hence the page layout is not transferred. Because of this, the page's original section needs to be located and the page layout settings for this section reapplied to the new document after a page's content has been copied.

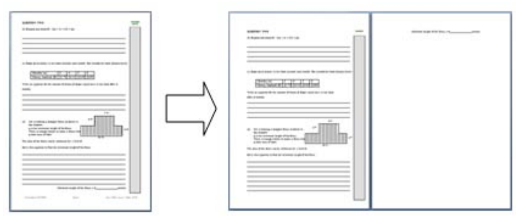

Figure 9 – A page copied to a new document loses page layout information

formatting are copied. The copied content is then pasted into the awaiting new document file. Due to the variety of different page layouts and content types, several different techniques need to be applied to deal with special cases.

Page Number Insertion

Inserting page numbers into a document can be done with a single click. But when a document has been split into multiple documents for enlargement, it would be unable to insert the page numbers without opening each section. Therefore, the following approach was decided upon: when a document is first split the user is queried with a dialog box asking whether or not page numbers are desired; if the

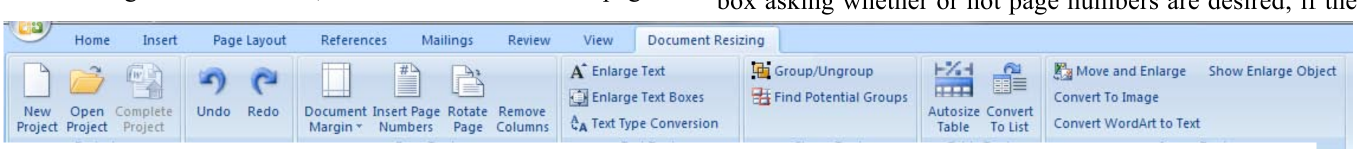

Figure 8 - Plug-in Interface

user selects yes then page numbers are inserted as the document is split. If the user chooses not to add them during the splitting process, by using the insert page number button on each part of the document, it inserts the appropriate page numbers in the section. Finally, if the user does not split the document, the insert page number button adds them to the entire document.

Page Rotation

Page rotation is a one-button-click operation. But what is the best way for users to select the document parts to rotate? Using the particular page the user is viewing is not accurate enough as the user may view multiple pages at once. We decided that the user's selection and caret position is the most reliable and logical way, because this is consistent with existing manipulation operations work, e.g. for changing text formatting. When the user has selected an area of the document, then this area is placed on rotated pages. When the user has not made any selection then the page the user's caret is on is rotated.

Related Object Selection

Word has two separate selection types: the default is 'select text' the other is 'select objects'. 'Select objects' allows users to drag a rectangle around an area and select all objects within this area. Using the select object technique on documents with lots of diagrams can be time consuming as it requires: 1) manually locating diagrams that are not grouped, 2) selecting the area of the diagram and then 3) grouping all the items together. A single button approach to locate any ungrouped diagrams, and group them together is a simpler technique. When the user chooses to locate ungrouped objects, the first ungrouped object is selected on screen. The user is then able to choose whether or not to group the selected object and add or remove elements from the group. This allows more flexibility as the user is able to navigate between different groups through the user interface at the top of the screen as shown in Figure 9.

This was achieved by indexing all the shapes in the document and sorting them by page and the logical position on the page. For each shape, we determine whether the shape's coordinates overlap (or are within a fixed distance from) another shape. If this is the case, the shapes are considered related and the two shapes are added to a candidate list representing a potential group. This is done until all shapes are either placed in a candidate list with other related shapes, or each shape has at least been checked for relations to other shapes.

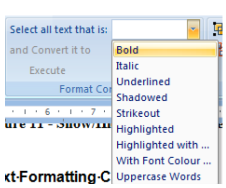

Figure 12 - Font effects conversion

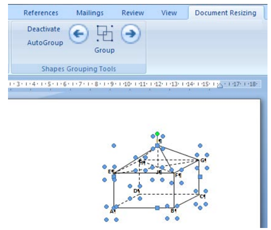

Figure 9 - Shape grouping toolbar

Text Resizing

For the resizing of text, a simple one-button-click approach is important as this is such as frequent task. Furthermore, the user needs to be able to change resizing settings quickly, without a cluttered user interface with too many options. Therefore, the function for showing and hiding of parts of the ribbon was used. Thus the user interface remains clean and tidy, while also allowing the access to advanced user controls without the need to navigate complex menu structures. Parts of the "Document Resizing" ribbon, with extended options shown and hidden respectively, are shown in Figure 10.

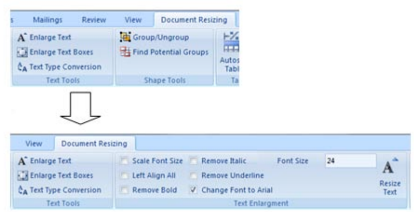

Figure 10 - Show/hide of UI components

It was decided to split the text resizing tools into two groups: Enlarge Text and Enlarge Text Boxes. This decision was made to allow different options for each. When text is resized in a text box, it is common for the text to become larger than the size of its container. In this scenario, the container needs to be resized. Each shape object in Word has the ability to resize itself to suit the size of the contents, this ability can be enabled via a tick box, if this is not enabled the user is notified of any shapes whose contains exceed their bounds.

Font Effects Conversion

Replacing the text effects in a document is similar to a standard find-and-replace in any text processor. Just as with find–and-replace, the user must specify what they are looking for and what they wish to replace it with. In the case of effects conversion, the user selects the effect to convert from (the source effect) as well as the effect type to convert to (the target effect), using two drop down menus. Since there is a limited number of effect options in Word, these lists are relatively short as shown in Figure 12. Analogously to the text find-and-replace, the effect conversion function finds all the text portions that match the selected source effect, and reformats them by applying the target effect.

EVALUATION

The requirements for the document resizing tools were elicited by interviewing the staff of a visual resource centre of an Auckland secondary school, and acquiring a representative sample set of 5 documents with 64 A4 pages in total, and resizing these documents manually. One of the documents was resized twice, as it was to be used by two visually impaired students with significantly differing requirements (one of the students needed a much bigger font size). This exercise demonstrated the time consuming nature of the process, and the importance of being able to resize documents for different people using different parameters.

An important insight from this initial phase was to shift the focus from resizing a document as a whole to splitting the document into smaller parts and resizing them separately using specialized tools. The former approach turned out to be too complex and practically infeasible. During the resizing it became apparent which operations are the most frequent and time consuming. This data forms the basis of our requirements and our implementation, which provides simplified tools for the most frequent operations.

In order to understand and evaluate our implementation of the resizing tools, the 5 documents that we resized manually were resized again, but this time using our tools. Then, another 8 documents were acquired from the visual resource centre and enlarged. These documents ranged in complexity and style, and consisted in total of 219 A4 pages. Documents made up mainly of text with limited graphics were very easy to enlarge using our tools. The tools provided for text significantly reduced the amount of interaction with Word in order to achieve the desired resizing tasks. Documents consisting mainly of tables and text were also resized easily and without the need for any major interaction with the document.

Documents with limited amounts of text and diagrams were easily resized with the text tools. The page rotation tools worked efficiently and effectively on all documents, reducing the amount of interaction required with the Word 2007 interface significantly. The table tools, which link directly to Word 2007 functions, were also very useful and reducing the need to switch between tabs and menus.

Documents with complex diagrams and large quantities of graphical components on one page required a large amount of manual work. The tool for selection of related objects often selected appropriate groups, but failed to be of practical use due to the lack of responsiveness and the lack of any real need for grouping of small items in any of the documents that were resized.

Word is technically very complex, which leads to some rare special cases that cause our tools to fail. For example, a failure occurs when a table spans multiple pages and the function for inserting page numbers is used. In this case, a page number is inserted inside the table instead of above it. Such problems are mostly due to the Word API, which is poorly documented and only provides limited control over the document.

For the documents that have been resized with both the manual and the semi-automatic approach, we found that depending on the complexity of the documents, a time saving between 20% and 50% was achieved when using our tools. That is, in the best case only half the time was needed to do the resizing. Note that the person doing this resizing work is a very experienced Word user. Furthermore, when applying the semi-automatic approach, he already knew the documents as he had resized them manually. As a consequence, it is reasonable to assume that for inexperienced Word users who have no previous experience with the documents they have to resize, a higher time saving can be achieved.

The tools were trialled at the visual resource centre by one of the most experienced staff members. The methodology used was: observation and talk-aloud protocols while the tools were being used, and a post-task an interview. Using the tools she resized a 25 page workbook in roughly 45 minutes. The staff member believed that resizing the workbook without the aid of the tools would have taken her around 5-6 hours.

The staff member was particularly impressed by the document splitting and joining functions, as would allow her to work a document in sections. She thought this would make it easier for her to work on the documents in small amounts of spare time over several days, without worrying about getting lost in the document. The tools for page numbering, page rotation, text enlargement and document margin adjustment were also highly valued, but the tools for grouping of related objects were found to be too complicated and not very valuable.

It is apparent that by providing a simpler interface to perform common tasks, our tools can significantly reduce the time needed for document resizing. A lot of the functionality made available in the tools can also be accomplished from within Word. However, it would require more advanced knowledge of Word, and is therefore very hard for novice users. Even for expert users, using Word without our tools would require a significant additional amount of time because it involves the navigation of complex menu hierarchies and dialogues.

During our observations it became clear that the current set of tools does not yet support all the common tasks that are necessary during resizing. By producing further tools and with further refinements to the user interface, the time required to resize a work book could be reduced further.

CONCLUSIONS

This work has explored the requirements of a system to assist resizing documents for the visually impaired. Requirements were garnered from standards, by analysing a representative sample of resource documents and advise from school staff experienced in document resizing. The goal was to create a tool to assist novice and expert users of Word to resize documents more easily. It became clear that it is not possible to fully automate the document resizing process, and that a semi-automatic approach would be preferable. In particular, the system to support resizing cannot be a single tool, but rather a set of tools to improve common resizing tasks within the document processing environment, which in the majority of cases is Microsoft Word.

Based on the requirements, a set of tools was implemented. The evaluation included: an empirical comparison of the resizing time with and without the tools; observations of a real user using the tool with talk-aloud protocol; and an interview. The results indicate that simplifying the steps required to resize a document with specialized tools provide a significant time saving. We found that the tools are most effective on text-heavy documents with a limited amount of graphics. Graphics-heavy documents were still easier to resize with the tools compared to manual resizing, but required the user to perform more resizing tasks manually. Initial results indicate a significant reduction in the interaction steps, knowledge and time required to resize a document. Based on document resizing trials, we estimate that the tools provide a time saving between 20% and 50%.

Our current focus is on ensuring the tools are compatible with all the different styles, components and formats that Word contains. This will help to reduce the rare cases where the tools do not perform as expected. Also, further functionality to assist in the resizing of graphics is desired and will be developed in the future, as this is one of the most daunting resizing tasks.

We plan to deploy the tools at the largest visual resource centre in New Zealand later this year. This will give us the opportunity to perform a long term trial, collect more feedback from staff, and determine the efficiency gain offered by the tools more accurately.

REFERENCES

Association of Directors of Social Services (2002). *Progress in sight: national standards of social care for visually impaired adults.* London.

Breuel, T. M. (2003). "High Performance Document Layout Analysis". *Proceedings of the Symposium on Document Image Understanding Technology*, pp. 209-218.

Breuel, T. M. (2002). "Two Geometric Algorithms for Layout Analysis". *Proceedings of the 5th International Workshop on Document Analysis Systems*, pp. 188-199.

Castles, I. (1993). *Disability, Ageing and Carers Australia: Visual Impairment*. Australian Bureau of Statistics.

Déjean, H., & Meunier, J.-L. (2007). Logical Document Conversion: Combining Functional and Formal Knowledge. *DocEng '07: Proceedings of the 2007 ACM Symposium on Document Engineering*, pp. 135-143.

Hadjar, K., Rigamonti, M., Lalanne, D., & Ingold, R. (2004). Xed: a new tool for extracting hidden structures from electronic documents. *Proceedings of the 1st International Workshop on Document Image Analysis for Libraries (DIAL'04)*, IEEE, pp. 212-224.

Hudson-Miller, J. (2007). *"LARGE PRINT for Mainstreamed Visually Impaired Students"*. Online at: http://www.ncaer.net/documents/3LargePrintStandardsArticlePartII08.doc

Jacobs, C., Li, W., Schrier, E., Bargeron, D., & Salesin, D. (2004). Adaptive document Layout. *Communications of the ACM 47(8)*, pp 60-66.

Lutteroth, C. (2008). Automatic Reverse Engineering of Hard-Coded GUI Layout. *Proceedings of the 9th Austalasian User Interface Conference*, pp 65-73.

Lutteroth, C., Strandh, R., and Weber, G. (2008). Domain specific high-level constraints for user interface layout. *Constraints*, 13(3): pp 307-342.

MacNeill, A. (2005). Overcoming Barriers to Creativity with the older Visually Impaired Person. *B.F.A. Eductional*.

Mao, S. and Kanungo, T. (2001). Empirical performance evaluation methodology and its application to page segmentation algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(3): pp 242-256.

Moh, C.-H., Lim, E.-P., & Ng, W.-K. (2000). Re-engineering Structures from Web Documents. *Proceedings of the 5th ACM Conference on Digital Libraries*, pp 67-76.

Petty, N. M. (2005). *Using student perceptions to evaluate the effectiveness of education for high school students with vision impairment.* PhD Thesis, University of Cantebury.

Vanderdonckt, J., Bouillon, L., & Souchon, N. (2001). Flexible Reverse Engineering of Web Pages with VAQUISTA. *Proceedings of the 8th Working Conference on Reverse Engineering,* p. 241.

Zhao, J. (2009). *Document resizing for visually impaired readers.* MSc Thesis, University of Auckland.