Information Visualization on Small Display Devices

Tomi Heimonen

University of Tampere

Department of Computer and Information Sciences

Master's Thesis

September 2002

University of Tampere

Department of Computer and Information Sciences

Tomi Heimonen: Information Visualization on Small Display Devices

Master's Thesis, 79 pages

September 2002

The amount of information is increasing rapidly all the while the devices with which to access the information are shrinking in size. This leads to a problem of presentation: how to display the growing amounts of information effectively on small displays.

This thesis introduces a number of information visualization techniques as solutions to the presentation problem, as well as devices on which the techniques could be used. Further, an evaluation framework is presented, which can be used to assess the suitability of the information visualization techniques for use on the devices in question. The framework draws together the visualization techniques and devices by comparing their corresponding properties. The proposed solutions are evaluated against the framework in the context of two small display devices.

As a result, it is found that only one third of the proposed solutions are suitable for use as such on low-end display devices, such as mobile phones. The situation is reversed in the context of high-end display devices, for example handheld computers. In their case only one third of the techniques are found unsuitable for use.

Keywords: information visualization, presentation problem, small displays, handheld devices

Index

1. Introduction

In the current information society that we live in the amount of information is increasing rapidly. Take, for example, the World Wide Web. From its origins in the 1970's and the 1980's the number of pages has grown from few hundreds to the range of several billion. Also, the advances in storage and microprocessor technology have made it possible to generate and publish further information much more rapidly than before.

However, these advances are also driving development into another direction. Personal computing devices are becoming smaller and smaller from the monolithic mainframes of the early decades in computing. Nowadays mobile phones and portable computers are a mainstay of information access in the business world where information needs to be available regardless of location or device.

These contradictory developments have created the need for bringing about ways of finding the right information from the vast mass of available information, and most importantly, ways of presenting that information in the most effective way possible. It might be trivial to view the organization chart of a large corporation on a high-resolution desktop computer display, but quite problematic on the small display of a mobile phone. There is thus a need to develop solutions that can facilitate the access to information on devices with small displays. Information visualization is one field of research that in part attempts to build solutions that address this presentation problem.

In this thesis the issues related to the presentation problem are addressed, some ways of resolving the problem in the context of small displays are introduced and evaluated.

First, the main principles and concepts of information visualization are introduced. This is followed by a detailed discussion on the presentation problem and its solutions as proposed in the literature.

After that, a context for the evaluation of the solutions is established by introducing the concept of small displays and actual devices with small displays.

This is followed by the construction of an evaluation framework that introduces the different aspects of both the possible information visualization techniques that could solve the presentation problem and the devices on which information is to be presented.

Finally, the suitability of the solutions for use on small display devices is analyzed and evaluated using the proposed framework. This will result in a relative ranking of the proposed solutions based on the type of device in question.

In conclusion, the presentation problem, its solutions and the suitability of the solutions will be discussed, along with the issues that emerged during the evaluation.

2. Information visualization

In the following chapter the main principles of information visualization and related terms are introduced along with the reference model for visualization that describes the information visualization process in detail.

2.1 Definitions

Information visualization and the terms relating to it have been defined in several different ways. In the context of this thesis it is therefore important to first establish certain definitions. Card *et al.* [1999, p. 7] define *information visualization* as "the use of computer-supported, interactive, visual representations of abstract data to amplify cognition." Information visualization is thus a tool that aids the human in gaining insight into data. The purposes of this insight are discovery, decision-making or explanation. A related activity to information visualization is scientific visualization, which is typically used to gain insight into scientific, usually physical, data.

Both information visualization and scientific visualization fall under the field of external cognition. Matlin [2001, p. 492] defines *cognition* itself as "the mental activities involving the acquisition, storage, transformation, and use of knowledge." Following that, Card et al. [1999, p. 7] specify *external cognition* as the "use of the external world to accomplish cognition."

The term v*isualization*, on the other hand, falls under the domain of internal cognition. Ware [2000, p. 1] defines the act of visualization as "constructing a visual image in the mind." The purpose of visualization is to form a mental model of the data that is being studied. Mental models are thus a human being's internal representations of the external visual world.

2.2 Reference model for visualization

There are several stages in the process that produces visual representations out of raw, abstract data. Card *et al.* [1999, p. 17] and Spence [2001, p. 13] both present similar reference models for visualization. The model presented by Card *et al.* is comprised of the following stages: *data transformations*, *visual mappings*, *view transformations,* and *human interaction*. In Spence's model the stages are named *selection*, *encoding*, *presentation*, and *interaction*. Figure 1 illustrates the reference mode for visualization as presented by Card *et al.*

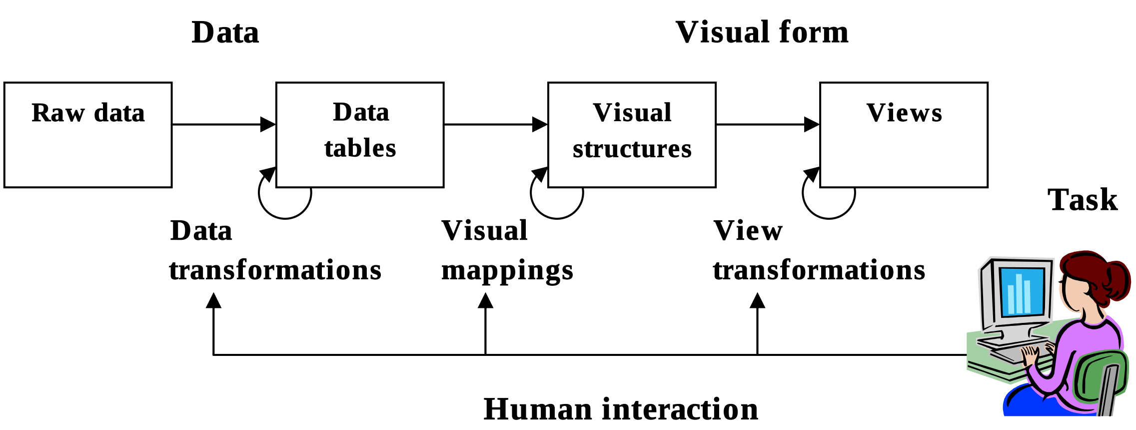

Figure 1. The reference model for visualization by Card *et al.*

2.2.1 *Data transformations*

The starting point of the visualization process is the raw data. This data, for example network traffic data, is in its original form difficult to understand. Therefore raw data is transformed into *data tables* that are relational descriptions of the data, including the metadata describing the relationships between the different points of data. Data tables can describe many types of data, nominal, ordinal or quantitative variables, in multiple dimensions. Information is typically lost or gained when transforming raw data into data tables, as new structures or values are derived from the original structures and values, for example through statistical calculations. A depiction of a data table is presented in Table 1.

	Case$_i$	**Case$_j$**	**Case$_k$**	**...**
Variable$_x$	Value$_{ix}$	Value$_{jx}$	Value$_{kx}$	…
Variable$_y$	Value$_{iy}$	Value$_{jy}$	Value$_{ky}$	…
...	…	…	…	…

Table 1. A depiction of a data table. After Card *et al.* [1999, p. 18].

2.2.2 Visual mappings

Data tables are transformed into visual structures by applying visual mappings. Visual structures combine *spatial substrates*, *marks* and *graphical properties* into a visual representation.

Spatial substrates include different ways of distributing the content of the data tables into space on different axes. Marks are the visible things that appear in space, such as points, lines, areas or volumes.

Graphical properties include attributes such as size, colour, texture, shape, and orientation, which indicate the order of marks and help distinguish the different marks from one another.

Lohse [1994] has developed a higher level classification of visual representations that includes 11 categories, such as graphs, tables, maps, and cartograms.

2.2.3 View transformations

Visual transformations are used to modify and augment the static visual structures in order to create views into the visual structures. These views exist in space-time and visual transformations are used to extract more information from the visualizations through interaction than would be possible statically.

According to Card *et al.* [1999, pp. 31-32] the three most common view transformations are *location probes*, *viewpoint controls*, and *distortions*.

Location probes are transformations that reveal the underlying data table information using location in the visual structure as the focus. For example *magic lenses* by Fishkin and Stone [1995] give an alternate view of the region under the magic lens view. An example of magic lenses is presented in Figure 2.

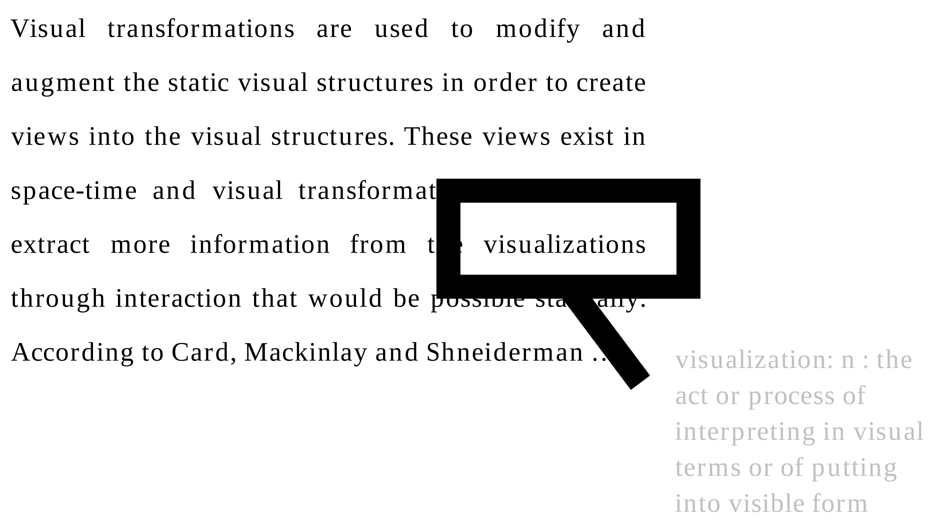

Figure 2. An example of magic lenses. After Spence [2001, p. 125].

Viewpoint controls are used to move the viewpoint into the visualization, for example, by zooming, panning or clipping. One example of viewpoint controls is for example the continuous zooming technique used in the *Pad++* system by Bederson and Hollan [1994].

Another viewpoint control technique is *overview + detail*, in which two windows are used together. An overview of the visual structure is shown in one window and the other window is used to present a detailed view of the focus area. An illustration of the principle of overview + detail is presented in Figure 3.

Figure 3. An example of the overview + detail technique.

Distortion techniques modify the visual structure, creating *focus* + *context* views. The focus area is used to show the underlying details of the visual structure in an undistorted view, whereas in the context area the details are distorted in order to save space. An example of a focus + context view is the *Perspective Wall* by Mackinlay *et al.* [1991], illustrated in Figure 4.

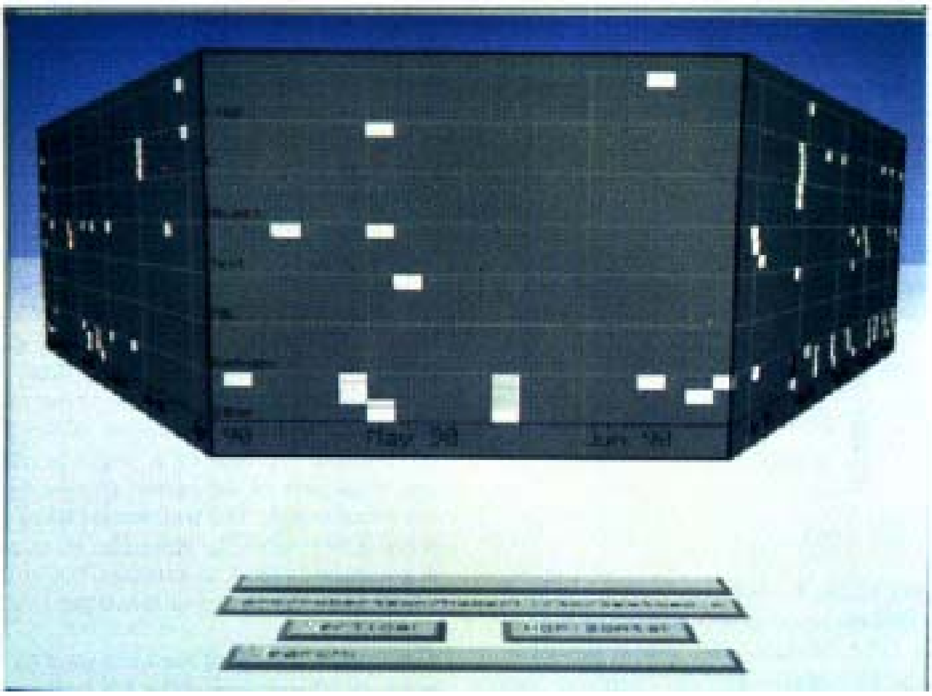

Figure 4. The Perspective Wall by Mackinlay *et al.* [1991].

There is, however, some ambiguity to the definitions of different view transformation techniques and it is therefore sometimes difficult to establish exact boundaries between different techniques. For example, focus + context and overview + detail are in many ways quite similar techniques, as both use the principles of a focus view to present the details and a less detailed overview to present the context.

Further, some focus + context techniques, such as the non-graphical fisheye views, do not employ distortion at all, but rely rather on other means, such as data suppression, to create a difference between the focus and the context. One example of the non-graphical fisheye views technique is the *Power Browser* by Buyukkokten *et al.* [2000]. In Power Browser suppression is used to present hierarchical structures of hyperlinks on web pages based on their importance and proximity to the currently selected page.

2.2.4 Human interaction

Human interaction completes the reference model suggested by Card *et al.* [1999, p. 33]. Human interaction can take place at any stage of the visualization process. In the data transformation stage the user can control which parts of the raw data are included in the visualization. In the visual mapping stage the user selects, for example, the mappings between the axes of the visual structure and the variables in the data table. In the visual transformation stage the user can control the viewpoint and its properties.

The reference model for visualization gives an approximation of the steps that the information visualization process is composed of. In real life, the visual reasoning tasks that users carry out usually combine these steps into more elaborate loops of actions. Thus, the different steps in the information visualization process are not one-off events, but rather loops unto themselves, in which the data tables, visual mappings and visual transformations change based on the human interaction with the system.

3. The presentation problem

In this thesis the emphasis is on discussing the issues related to the view transformation stage and the human interaction stages of the visualization process, especially in the context of small display devices. According to Spence [2001, pp. 111-112], with the recent advent of small and mobile handheld devices such as PDAs (Personal Digital Assistants) and mobile phones, the role of the *presentation problem* has become increasingly crucial.

The problem is threefold: on one hand there is a need to present large amounts of information on a limited display area, on the other hand there is also a need to support the human activity of forming a mental model of the data. In addition, the interaction capabilities found in handheld devices limit the possibilities to interact with the data. Spence presents a number of possible solutions to overcome the presentation problem that will be discussed in the following.

The most obvious solution is to use scrolling in order to bring data in and out of the visible area. For example, when editing a long document in a word processor application the user can move between chapters and sections by scrolling the underlying text via some scrolling mechanism. The problem with this approach is that most of the document is always hidden from view, thus making it difficult to assess the context to which the currently displayed data belongs to. Figure 5 gives an example illustration of scrolling.

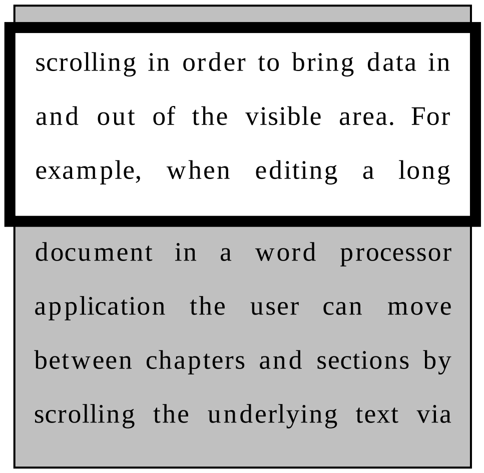

Figure 5. An example of scrolling. After Spence [2001, p. 113].

A way to remedy the problems associated with scrolling is to provide context. This can be done, for example, by providing a less detailed context view near the detailed focus view, adopting a manner of an overview + detail technique. This solution, however, poses problems of its own. It is often difficult to make the mental transfer from the edge of the detailed view to the corresponding area in the context view.

A study by Hornbæk and Frøkjær [2001], in which the usability of linear scrolling, overview + detail, and fisheye views was compared, suggests that overview + detail and fisheye view techniques are superior to linear scrolling in reading tasks. The problems associated with overview + detail techniques are demonstrated by the fact that the fisheye view technique achieved better performance, although it gained higher subjective ratings from the test subjects.

Ware and Lewis [1995] have proposed a solution, the *DragMag Image Magnifier*, to this continuity problem that uses a magnification approach to show both detail and context in the same view. Lieberman [1997] proposes the use of transparency to show the detailed view superimposed over the context view.

Figure 6 illustrates this idea as implemented by Lieberman in an earlier prototype [1994].

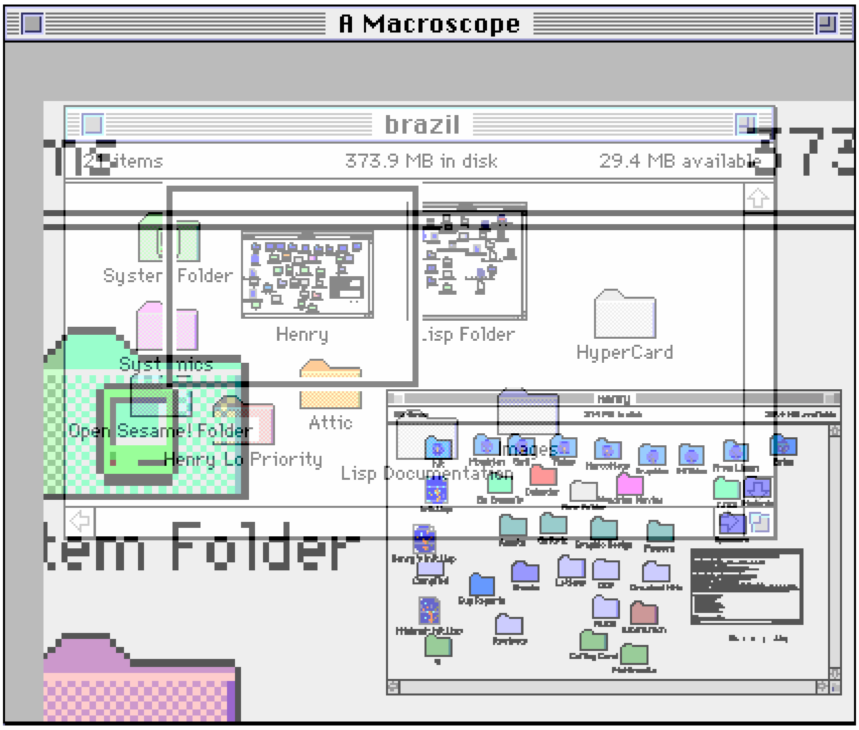

Figure 6. The Macroscope by Lieberman [1994].

The presentation problem can thus be defined as a problem of providing both detailed and contextual information in the same, often constrained, view. In the area of small displays it becomes even more important to be able to do this effectively.

During the recent two decades a lot of attention in research has been put into solving the problems dealing with the successful union of focus and context information. Spence [2001, pp. 116-133] breaks these solutions down into six categories, which are: (1) focus + context, (2) suppression of information, (3) magic lenses, (4) browsing, (5) effective views, and (6) zoom and pan. This categorization, excluding effective views due to their similarity with other proposed techniques, will be used as a basis in this thesis to discuss the issues related to information visualization on small displays. The general principles and prominent techniques behind these solutions will be presented in the following section.

3.1 Approaches for solving the presentation problem

3.1.1 Focus + context techniques

Focus + context techniques have three premises, as outlined by Card *et al.* [1999, p 307). First of all, both the context and the detailed information must be available simultaneously. Second, the type of information shown in the detailed view and context may be different from one another. Third, these two types of information should be combined into one display.

According to Card *et al.* there are several motivations for using a single view. The use of two separate displays incurs a penalty on the performance of visual search and working memory. Also, there seems to be a systematic way in which the users' attention in detail decreases when moving away from the object of attention. This leads to the hypothesis that it may be possible to build visualization techniques in which the peripheral context information is shown in reduced detail while showing the information in the focus in more detail, and the amount of detail is dynamically varied based on the changes in the user's attention.

Focus + context techniques typically use the principle of distortion through a *visual transformation function* to accomplish the focus + context demands. Visual transformation functions distort the visualizations through stretching and compression in order to give the focal portion of the visualization more visual detail. The shape of the transformation, that is how the magnification factor changes across the visualization, is expressed by a magnification function that is a derivative of the visual transformation function.

A thorough review of distortion-oriented focus + context presentation techniques has been presented by Leung and Apperley [1994], in which the transformation functions are divided into two categories: *step functions* and *continuous transformation functions*. Step functions change the values discontinuously from one constant to the next, resulting in separate discrete levels of magnification in the visualization. Continuous transformation functions, on the other hand, scale the values in a continuous manner, resulting in smooth magnification. The difference between step functions and continuous magnification functions is illustrated in Figure 7.

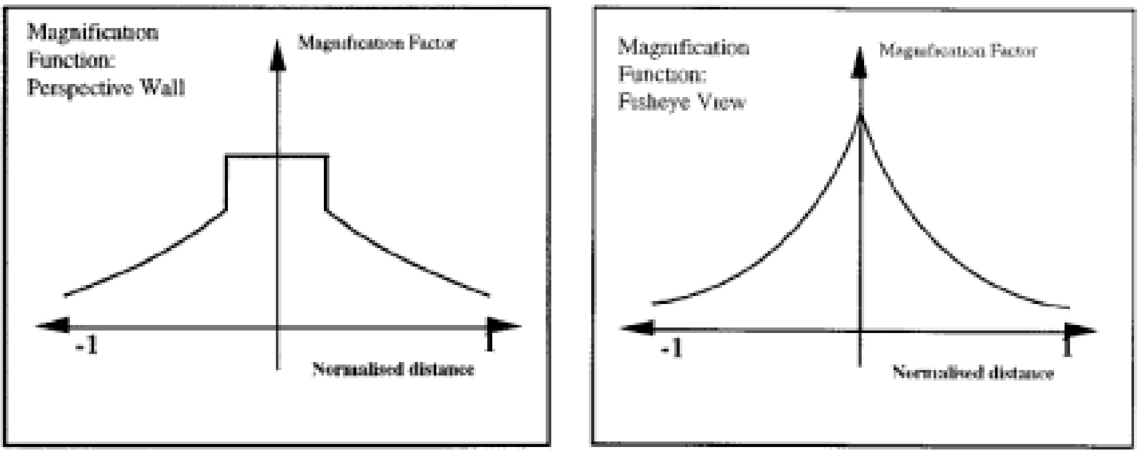

Figure 7. Comparison of step vs. continuous magnification functions [Leung and Apperley, 1994].

The *Bifocal Display* by Spence and Apperley [1977] is one of the first focus + context techniques. With the Bifocal Display a step function is used, and the resulting distortion is analogous to a long strip of paper being shown to the user over two posts in such a way that all the contents of the paper are visible. Only the items in the centre are readable to the user, but the user can still be aware of the items away from the centre. The principle behind the Bifocal Display is presented in Figure 8.

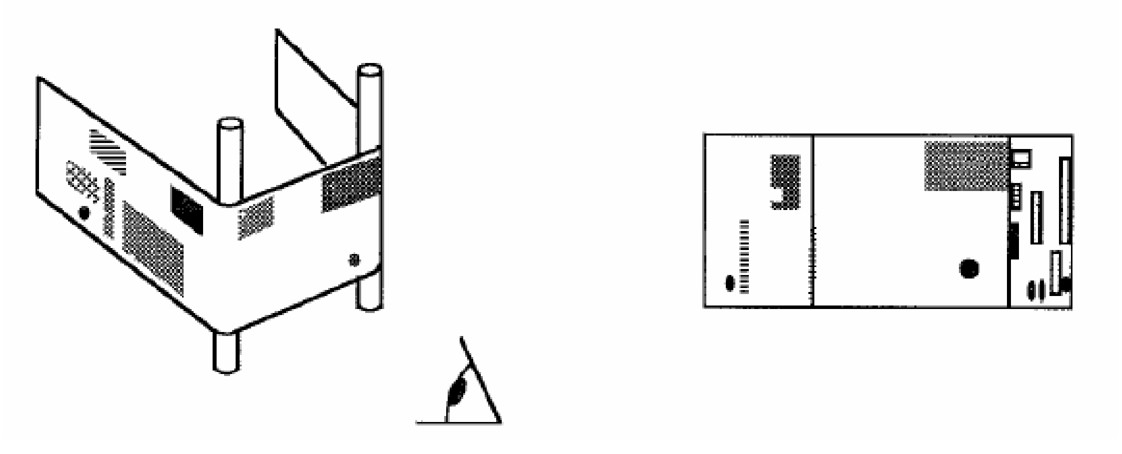

Figure 8. The Bifocal Display principle [Leung and Apperley, 1994].

An extension of the Bifocal Display principle, the *Perspective Wall*, was developed by Mackinlay *et al.* [1991]. It simulates a three-dimensional view of a file structure view, building on the distortion principles advocated in the Bifocal Display. In the Bifocal Display and Perspective Wall the context view is distorted in x dimension only, but the distortion could also be applied in the y dimension, or in both dimensions simultaneously. These distortion principles have been extended in the *Document Lens* by Robertson and Mackinlay [1993] and the *Table Lens* by Rao and Card [1994]. An illustration of the Document Lens is presented in Figure 9.

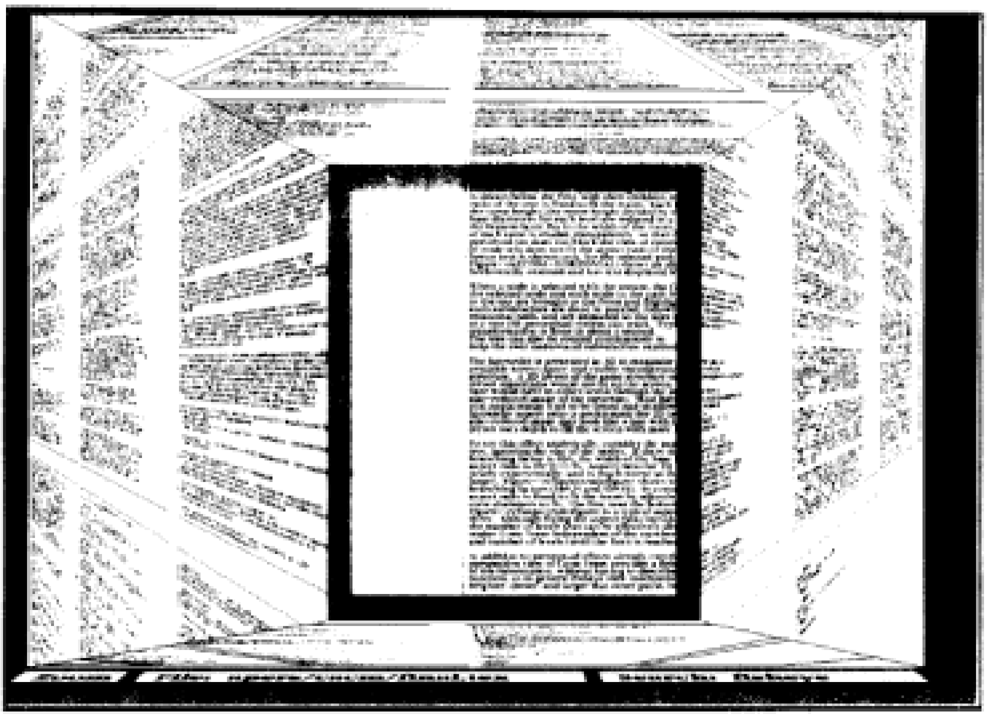

Figure 9. The Document Lens by Robertson and Mackinlay [1993].

Kadmon and Shlomi [1978] proposed a polyfocal projection for presenting statistical data on cartographic maps and its application, the *Polyfocal Display*. The basic principle behind the polyfocal display is similar to that of the Bifocal Display, however Polyfocal Display uses multiple focus areas and the view is distorted in both x and y dimensions using a continuous magnification function.

The previous techniques presented in this thesis have used focus + context techniques to present information in two-dimensional space. However, space itself can be distorted to produce implicit focus + context views by using non-Euclidean geometry. Lamping and Rao [1996] have proposed the use of hyperbolic geometry to achieve focus + context.

Lamping and Rao's application, the Hyperbolic Browser, can be used to present large hierarchies efficiently. The Hyperbolic Browser first maps the hierarchy on a hyperbolic plane and the resulting structure is then mapped to normal

(Euclidean) two-dimensional plane. An illustration of the Hyberbolic Browser is presented in Figure 10.

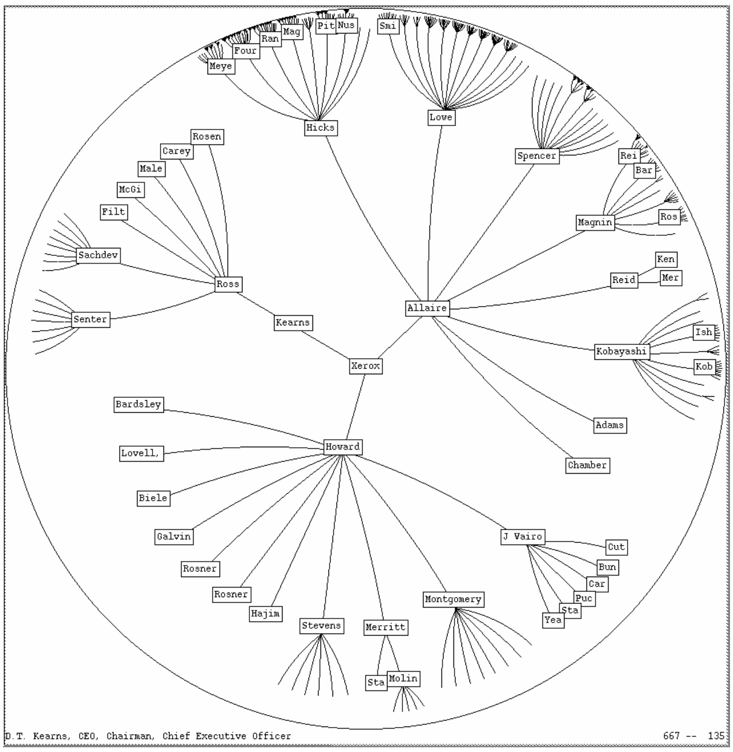

Figure 10. The Hyperbolic Browser by Lamping and Rao [1996].

3.1.2 Suppression techniques

Suppression techniques build upon the notion that not all information is always relevant to the task the user is attempting to accomplish and can thus be suppressed from the view. The most prominent technique to take advantage of this notion is the *fisheye* technique suggested by Furnas [1986]. According to Furnas it is often practical to show data only if its perceived value exceeds some threshold set by the user. The perceived value is calculated by a *degree of interest* (DOI) function. For example, in the most simplistic case the perceived value could be calculated as a factor of the distance of the item from the current viewpoint. As consequence of applying the DOI function, items that have perceived values above the threshold are shown in the view while items with values below the threshold are suppressed from the view.

However, in many cases the situation is more complex. In these cases the DOI function would be composed of a presupposed (*a priori*) component, the general interest of the object, and an observed (*a posteriori*) component, the distance of the object from the current viewpoint. For example, in map applications this would mean that larger cities with high general interest that are far away from the current viewpoint would more likely be shown than small cities with low general interest.

Figure 11 gives an example of the operation of the DOI function. The numbers indicate the value calculated by using a simple proximity-based DOI function. Initially, the focus is on the root node, and then it is shifted to the indicated child node. The tree structure on the right shows which nodes are shown when the threshold is set at -4.

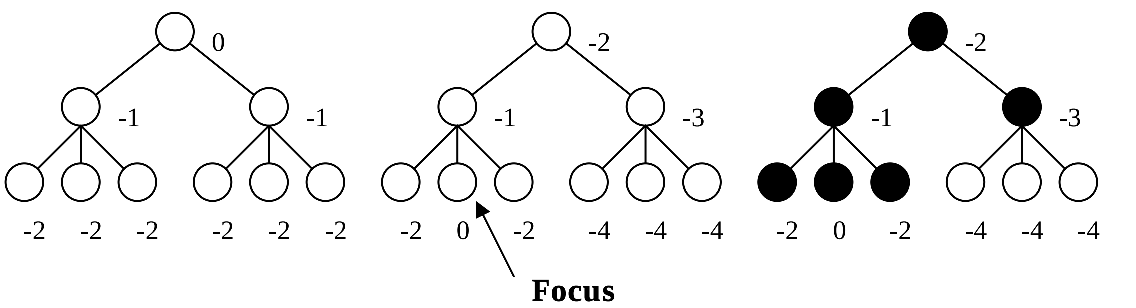

Figure 11. The operation of the DOI function in Furnas' tree-view. After Spence [2001, p. 122].

Card *et al.* [1999, p. 307] suggest that the selective reduction of information for the contextual area can be obtained by several other means as well, for example through highlighting or distortion. Instead of removing the objects from the view, only the objects of interest could be highlighted or the view could be distorted so that only the objects near the points of interest are visible in detail.

In Furnas' work the fisheye principle is used in the domain of non-graphical tree structures, but it has been extended, for example, by Sarkar and Brown [1994] with their graphical fisheye views, and by Bederson [2000] with his application of the fisheye technique to application menus (*fisheye menus*). An implementation of Sarkar and Brown's graphical fisheye view used to visualize the structure of network graph is illustrated in Figure 12.

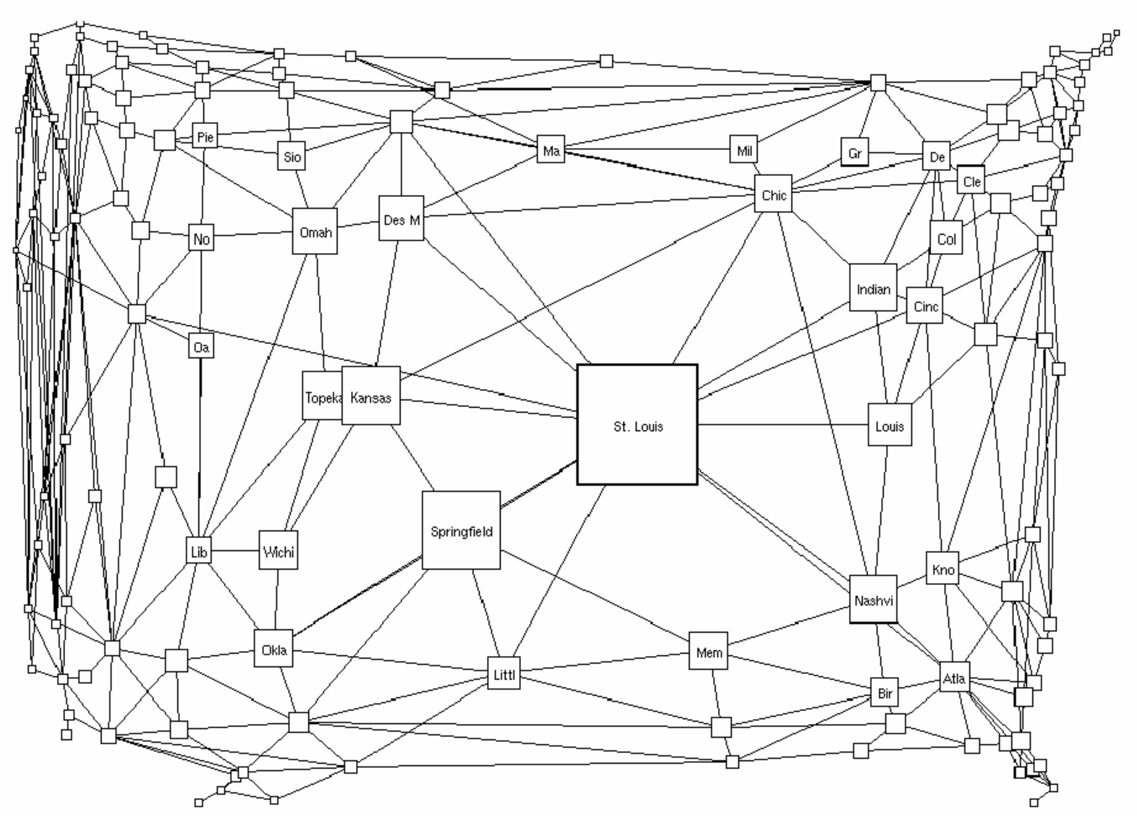

Figure 12. Graphical fisheye view by Sarkar and Brown [1994].

While in this thesis fisheye views are presented as a data suppression technique, they could also be discussed under the heading of focus + context techniques. The division used here is, however, based on the categorization provided by Spence [2001, pp. 116-133], in which the suppression of detail is presented as a sufficient basis by which to separate fisheye views from the larger family of focus + context visualization techniques.

Z-thru mapping by Colby and Scholl [1991] is another technique that builds on the idea of selective display of information. Z-thru mapping views consist of information layers associated with a certain property of the data being displayed. The transparency and focus of these layers can be controlled by the user so that the relevant information can be viewed according to the user's needs, for example the transparency of certain layers can be decreased to make the information on that layer stand out better.

The z-thru mapping concept has been further developed for example by Silvers [1995]. Figure 13 illustrates Silvers' weather data visualization in which different layers are used to present geographical, location, and weather data.

Figure 13. Z-thru mapping by Silvers [1995].

Suppression and distortion can also be combined to create effective visualizations. One hypothetical example by Spence [2001, p. 124] uses the polyfocal display distortion technique by Kadmon and Shlomi [1978] to provide sufficient detail for navigation inside two cities while also employing suppression to filter away unnecessary detail on the trail between the cities. An illustration of Spence's example combining distortion and suppression techniques is presented in Figure 14.

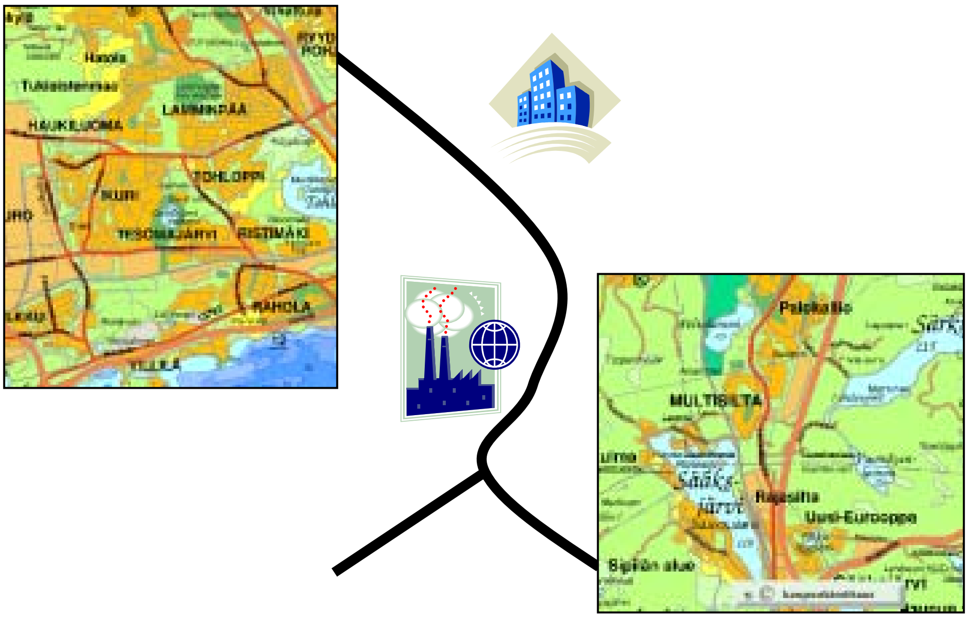

Figure 14. An application of the "A Really Useful Map" example. After Spence [2001, p. 124].

3.1.3 *Magic lenses*

The fundamental metaphor of the *magic lenses* concept is that of using a magnifying glass to reveal more detail about the underlying information. The display of increased detail is limited to a small region in order to prevent visual clutter over the whole display area. The selection of items is done according to the x, y position of the lens and the parameters assigned to the lens. For example, the lens could be parameterised to show only such items that have a property the value of which exceeds a certain threshold value assigned by the user.

Fishkin and Stone [1995] have proposed an application of the magic lenses concept. In their implementation magic lenses are used to perform data base queries. Each lens is assigned a Boolean query that filters the underlying values, and items that match the query are indicated with colour coding. Combined queries can be formed by moving the lenses to overlap one another. The magic lenses implementation by Fishkin and Stone is presented in Figure 15.

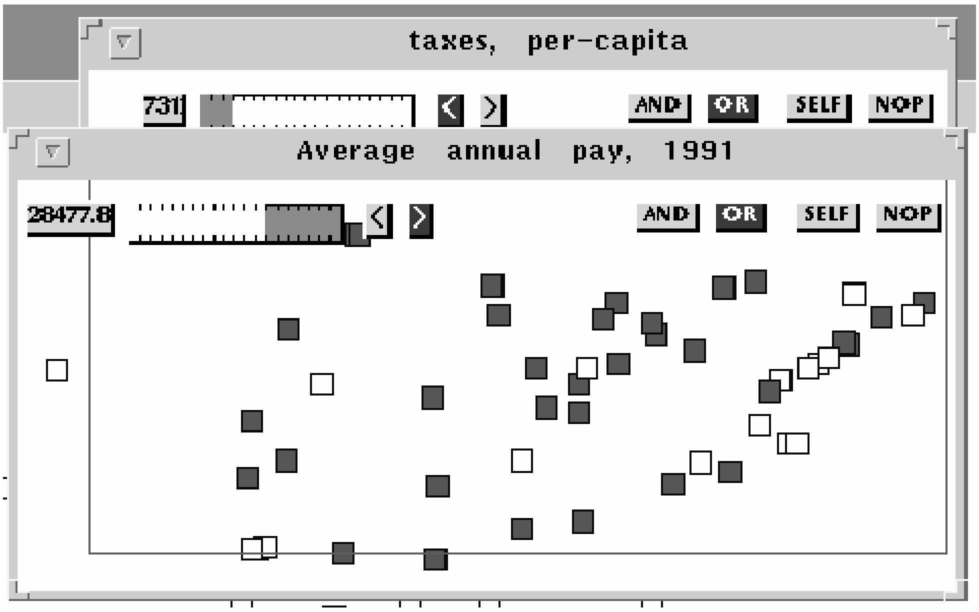

Figure 15. Magic lenses by Fishkin and Stone [1995].

3.1.4 *Riffling techniques*

Riffling allows users to quickly browse through large quantities of data with the goal of gaining some kind of idea of its contents. The concept of riffling is based on a trade-off between space and time. A simultaneous full view of the data might require a much larger display area that can be used – with riffling the same physical space is used to present the user with rapid snapshots of the data in a linear and serial manner. In this sense riffling is analogous to flipping through the pages of a book printed on paper. Spence [2001, p. 126] suggests that at display rates of five to ten images per second users are still able to recognize the images that are being displayed.

A number of interesting applications based on the concept of riffling have been developed. Wittenburg *et al.* [1998] developed the PolyNav™ previewer, a prototype that supports image-based previewing of information spaces. Their implementation is used to present the goods and services of a World Wide Web service to potential customers. In the PolyNav™ previewer the hyperlinks on a page are shown in a sequence as a stream of images. The prototype is presented in Figure 16.

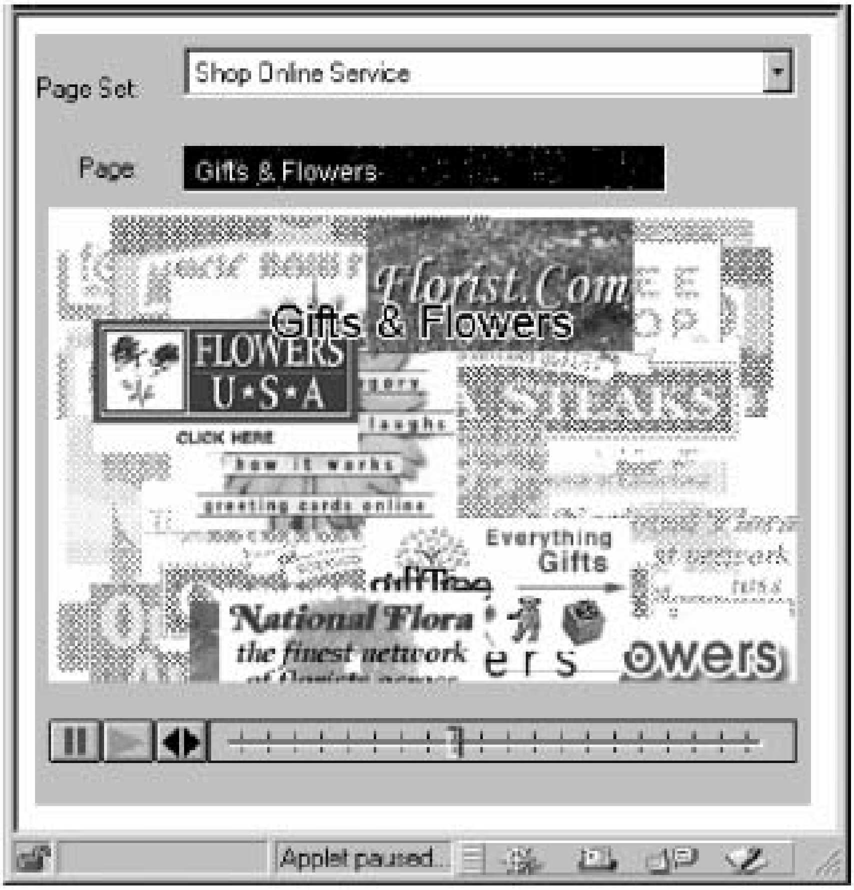

Figure 16. A Java™ applet version of the previewer by PolyNav™ Wittenburg *et al.* [1998].

De Bruijn and Spence [2000] have developed the RSVP (*Rapid Serial Visual Presentation*) technique based on the principle of riffling. In their approach riffling is used to browse through the contents of file-system folders in a sequential manner. Figure 17 illustrates the operation of the RSVP technique.

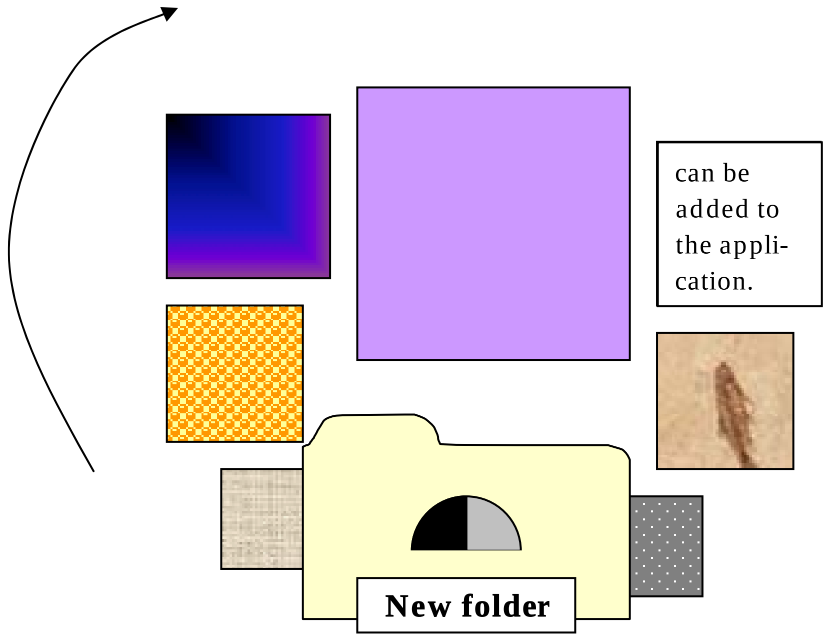

Figure 17. An example of the operation of the RSVP technique. After Spence [2001, p. 128].

3.1.5 Zoom and pan techniques

Zoom and pan are techniques for moving and magnifying the viewpoint in the visual structures. Zooming is defined by Spence [2001, p. 130] as the "increasing magnification of a decreasing/increasing fraction of a two-dimensional image under the constraint of a viewing frame of constant size." Panning is in turn defined as "the smooth movement of a viewing frame over a two-dimensional image of a greater size."

Spence [2001, p. 131] presents an example of zoom and pan by examining two points in an image. In order to investigate the points in detail, the user needs to zoom in on the points. To move from point *A* to point *B* could be accomplished by using conventional panning, however this might result in a loss of context. A better solution is to first examine point *A*, then zoom back to the overview level and then pan the viewpoint to point *B* and zoom in to examine the details. An example of a zoom and pan process is illustrated in Figure 18.

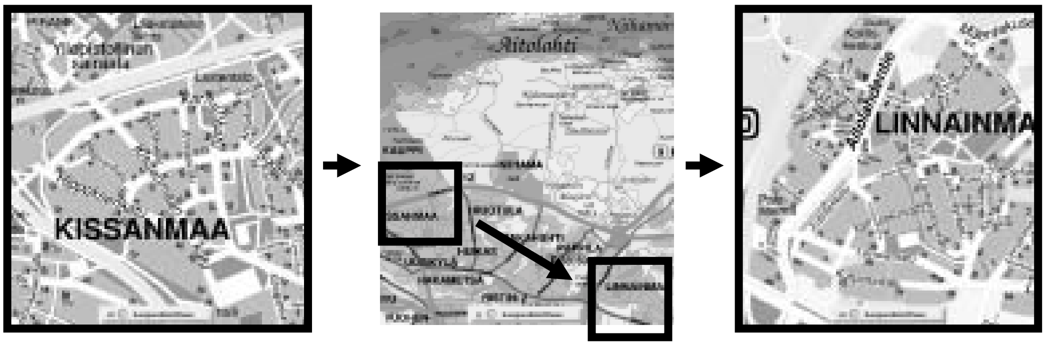

Figure 18. A zoom and pan example. After Spence [2001, p. 130].

The use of zoom and pan alone is not necessarily an adequate solution to overcome the problems presented by small display area. The applications of conventional geometric zooming are also limited. Therefore an extension to conventional geometric zooming called *semantic zooming* has been proposed. In semantic zooming interfaces the objects can also change their appearance and content when zoomed into.

The *Spatial Data Management System* by Herot *et al.* [1980] is one example of semantic zooming. In the overview the user is presented with a collection of available ships. Semantic zooming, combined with geometric zooming, reveals the overall details of the selected ship and further semantic zooming might reveal additional information. An example of the operation of spatial zooming is illustrated in Figure 19.

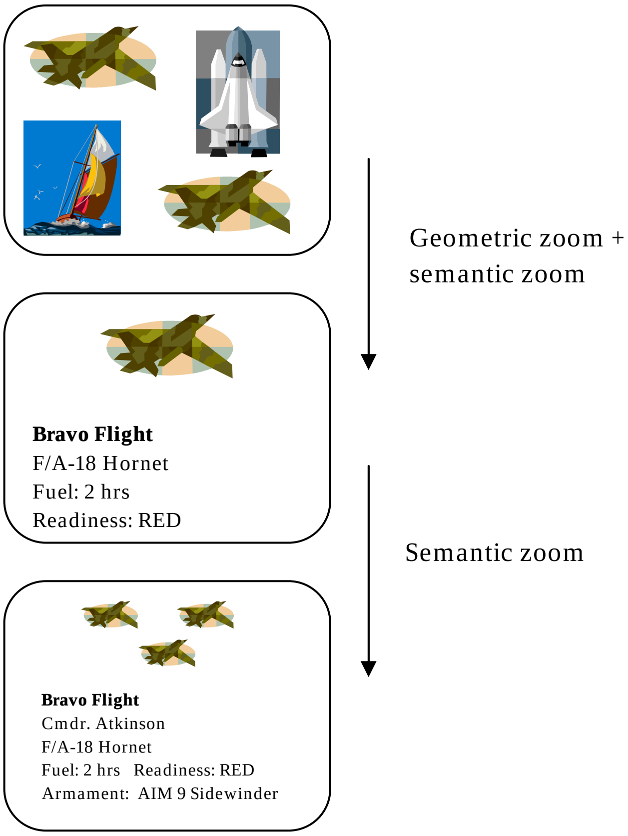

Figure 19. An illustration of spatial zooming. After Spence [2001, p. 133].

3.2 Summary

The visualization techniques presented in the previous sections form the pool of techniques that will be evaluated as possible solutions to the presentation problem on small display devices.

The above techniques produce primarily two-dimensional visualizations. Techniques that produce three-dimensional visualizations, such as the *Cone Tree* technique by Robertson *et al*. [1991], will not be evaluated in this thesis because the current generation of handheld devices does not include devices that would support the use of three-dimensional graphics in a satisfactory fashion, for example through hardware acceleration. Devices with support for three-dimensional will, however, enter the marketplace in a matter of years.

4. Small display devices

The definition of small displays is presented and discussed in the following, establishing a backdrop against which information visualization techniques can be evaluated. A number of existing devices acting as the actual context for evaluation are also introduced in detail.

4.1 Definitions

When considering information visualization in the context of small display devices, one must first define what is meant by small displays. In this thesis small displays are discussed in terms of actual products available to the public. This approach was chosen in order to be able to compare the solutions presented in section 3.1 against real world constraints found in actual devices.

A definition for the size of a display is based on the *resolution* of the display element. Resolution is the number of individual pixels (points of colour) incorporated on a display element. Resolution is expressed by reporting the number of pixels on the horizontal axis and the number of pixels on the vertical axis. The resolution of a typical laptop computer display is, for example, 1024 pixels by 768 pixels (1024 × 768).

A related feature, *sharpness*, depends on the resolution and the physical dimensions of the display element. The same resolution will look sharper on a smaller display element and progressively lose sharpness on larger displays since the same amount of pixels is being spread out over a larger area.

A given display element will have a maximum resolution that depends on its physical ability to focus light. On maximum resolution the physical dot size (*dot pitch)* matches the pixel size. Display elements can typically present several lesser resolutions below the maximum resolution. For example, a display system that supports a maximum resolution of 1280 × 1024 pixels may also support 1024 × 768, 800 × 600, and 640 × 480 resolutions. On a display of a given size, the maximum resolution may offer the sharpest image possible but be spread across a space too small to readable. [1]

Typical desktop CRT (Cathode Ray Tube) displays are capable of presenting resolutions upwards to 1600 × 1200 pixels or more. In contrast, typical handheld devices, including mobile phones, usually employ LCD (liquid crystal display) technology. The resolution of handheld device displays ranges typically from 96 × 65 pixels to 640 × 240 pixels. Figure 20 (after Siirtola [2000]) illustrates the relative proportions of the different display resolutions.

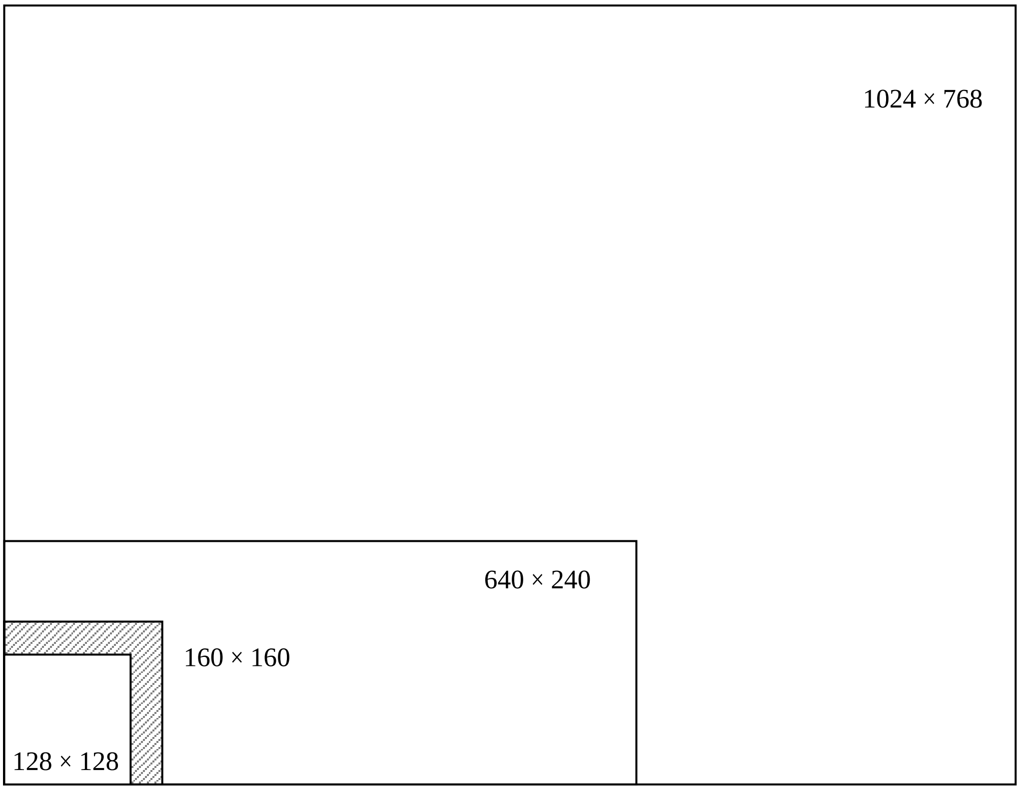

Figure 20. The relative proportions of different display resolutions.

[1] whatis.com on-line dictionary of IT-related words: *http://whatis.techtarget.com/*

However, also the aspect ratio of the display should be considered along with the resolution when defining the effective size of a display. The resolution of the display measures how many individual pieces of information the display can present, whereas aspect ratio indicates the shape of the display.

Studies in the entertainment industry have established that a display with an aspect ratio of 16:9 is the most appealing to users, based on the characteristics of human vision. This would suggest that the aspect ratio of the display could be used to determine the effective size of the display by making the assumption that a display with an aspect ratio of 16:9 is optimal when compared to other displays having the same resolution.

This assumption seems reasonable, if one considers a real world example. A display with a resolution of 75 × 133 pixels is far more useful for visualizing a street map than a display with a resolution of 10 × 1000 pixels, although both displays feature the same amount of pixels. Naturally the situation is not so simple, as the type of data being visualized has also a profound impact on the optimal shape of the display. Additionally, there are other factors that contribute to the effective size of the display, such as the amount of colours the display element is capable of reproducing.

However, as formal methods for unequivocally establishing the size have yet to be presented, the resolution of the display has been used to establish a boundary between small, medium-sized, and large displays. In this study eleven display types of varying size from existing devices were used to calculate the data presented in Table 2.

Device	Display width	Display height	Pixel count
Sony Ericsson t39m [2]	103	33	3 399
Nokia 3310 [3]	84	48	4 032
Nokia 6310i [4]	96	65	6 240
Siemens ME 45 [5]	101	80	8 080
Nokia 6610 [6]	128	128	16 384
Palm m515 [7]	160	160	25 600
Nokia 7650 [8]	176	208	36 608
Compaq iPAQ [9]	240	320	76 800
HP Jornada 720 [10]	640	240	153 600
VGA display	640	480	307 200
XGA display	1024	768	786 432

Table 2. Device display resolution and pixel count.

Based on these figures it can be concluded that a device belongs to the small display category if it has a resolution of approximately 150 000 pixels. This seems to be the “natural” cut-off point between the smallest desktop monitors and handheld devices. In addition, two distinct categories can be discerned

[2] http:// http://www.sonyericsson.com/T39/

[3] http:// http://www.forum.nokia.com/main/1,35452,1_1_75_03,00.html

[4] http:// http://www.forum.nokia.com/main/1,35452,1_1_75_03,00.html

[5] http:// http://www.my-siemens.com/

[6] http:// http://www.forum.nokia.com/main/1,35452,1_1_75_16,00.html

[7] http:// http://www.palm.com/products/palmm515/

[8] http:// http://www.forum.nokia.com/main/1,35452,1_1_75_24,00.html

[9] http:// http://www.compaq.com/products/handhelds/pocketpc/

[10] http:// http://products.hp-at-home.com/

within the small display category. Devices with resolutions that fall below 10 000 pixels belong to a "low-end" category, whereas devices with over 10 000 pixels belong to a "high-end" category. The differences in resolution are illustrated in Chart 1.

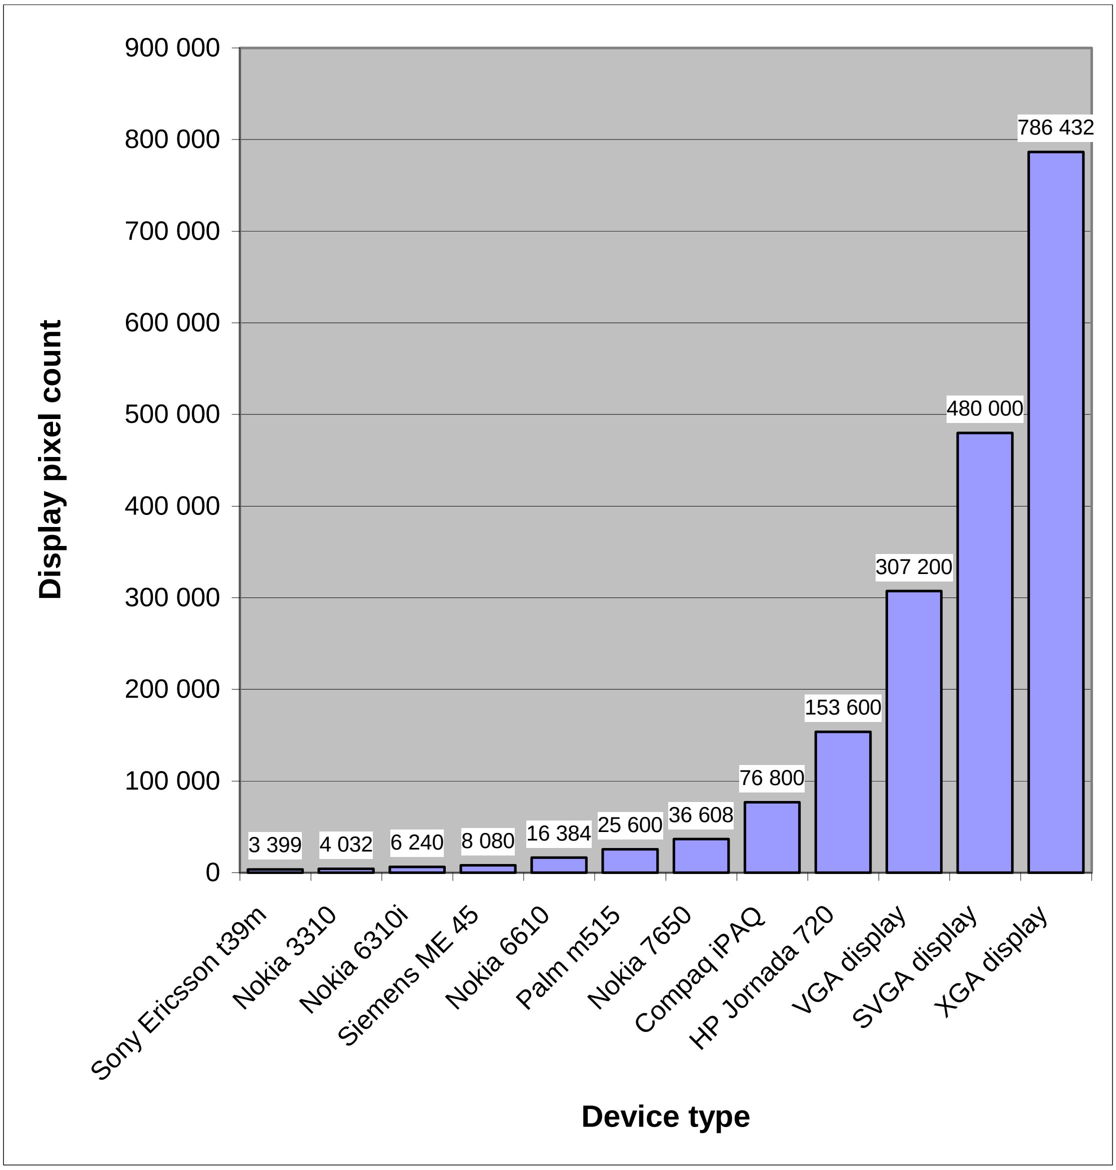

Chart 1. Display resolution (pixel count) by device type.

4.2 Devices

In order to assess the suitability of the visualization techniques presented in section 3.1, it is necessary to select a number of devices that can be used as evaluation platforms. The devices used in this thesis for the evaluation were selected based on three factors. First, the device must belong to the small display category as measured by the aforementioned display resolution categorization. Second, the device should support user-extensible applications, that is, it should be possible to develop visualization applications for the device by using publicly available development tools. Third, the device should include connectivity features that allow data to be brought into the device for visualization purposes.

Two devices, the Nokia 6310i GSM mobile phone and the Palm m515 PDA, were selected to represent typical small display devices in the own class. Both devices belong to the small display category and offer support for application development and connectivity. The 6310i and m515 feature different physical display measurements, display resolutions, input methods, and connectivity capabilities. Resolution-wise, Nokia 6310i belongs to the low-end display device category and Palm™ m515 to the high-end display device category of small displays. This makes it possible to assess the suitability of the visualization techniques as thoroughly in the context of small display device as possible.

4.2.1 Nokia 6310i

The Nokia 6310i[11] GSM mobile phone represents a typical professional mobile phone with a 96 × 65 pixel monochrome display and conventional nine key input system and additional application keys for menu browsing, selections,

---

[11] The Nokia 6310i at a Glance: *http://www.nokia.com/phones/6310i/*

and call management. The Nokia 6310 GSM mobile phone is featured in Figure 21. The connectivity features of the phone include, for example, GPRS (General Packet Radio Service[12]) and Bluetooth™ technologies[13].

Figure 21. Nokia 6310i.

The phone also features Java™ application development[14] support. This means that it is possible to implement custom visualizations for the phone by using the Mobile Information Device Profile (MIDP) application programming interface[15] and the Connected Limited Device Configuration (CLDC) implementation[16] of the Java™ programming language. The 6310i has only 178 kilobytes of memory

[12] Nokia GPRS: *http://www.nokia.com/gprs/*

[13] The Official Bluetooth™ Website: *http://www.bluetooth.com/*

[14] Java™ Programming Language: *http://java.sun.com/*

[15] Mobile Information Device Profile: *http://java.sun.com/products/midp/*

[16] Connected Limited Device Configuration: *http://java.sun.com/products/cldc/*

available for Java™ applications, which naturally constraints the complexity of the visualization applications that can be developed for the device.

4.2.2 Palm™ m515

The Palm™ m515[17] handheld device is a representative of a typical PDA with connectivity features. It has a 160 × 160 pixel display supporting more than 65,000 colours, and it is controlled with a stylus in addition to five application keys. The Palm™ m515 handheld device is featured in Figure 22.

Figure 22. Palm™ m515.

The Palm™ Operating System (Palm OS) provides its own application programming interface for developers to produce applications for Palm™ devices. A number of development environments are also available for creating Palm OS applications in a number of different languages such as C, C++, Visual BASIC, or Java. The m515 model includes 16 megabytes of memory and a Motorola Dragonball VZ 33 megahertz CPU (Central Processing Unit), which

[17] Palm™ m515 Handheld: *http://www.palm.com/products/palmm515/*

makes it possible to develop even quite computationally intensive visualizations for the device.

5. Evaluation framework

In this chapter a framework is presented for evaluating the suitability of the visualization techniques that have previously been proposed as solutions to the presentation problem. The framework is introduced through a number of constraint categories, which bind together the properties of the visualization techniques and the properties of small display devices.

5.1 Overview

There are several criteria that must be observed when evaluating the suitability of a visualization technique for use on a small display device. On one hand, there are the constraints imposed by the device used to display the visualization. On the other hand, there are the constraints that are imposed by the visualization technique itself in terms of interaction and visual aspects, such as the use of colour.

The constraints are divided into three categories: (1) visual constraints, (2) interaction constraints, and (3) implementation constraints. Visual constraints and interaction constraints are the most critical, given the two integral elements of the visualization process: presentation and human interaction. Visual constraints address issues such as the types of visual structures and graphical properties that can be used to construct the views. Interaction constraints include the different types of methods that can be used to interact with the visualizations at different points of the visualization process. Implementation constraints include issues such as processing power and storage space requirements.

Each category also includes the constraints imposed by the device in question, namely the resolution and type of the display, different interaction methods available in the device, and the hardware and software features that allow for the implementation of the visualization techniques.

Thus, the important questions that must be asked when considering the suitability of a given visualization technique on a specific device are thus: “Can this technique be adequately graphically presented on this device?”, “Can this technique be adequately controlled given the interaction controls available on this device?”, and “Is it possible to implement this technique on this device, given its technological limitations?”

5.2 Visual constraints

5.2.1 Graphical properties

The way in which graphical properties are used to encode information contributes to the effectiveness of a given visualization. These graphical properties have been originally classified by Bertin [1983, pp. 60-61] to include size, orientation, value (greyscale), colour, texture, and shape. Bertin calls these properties *retinal properties*, based on the fact that the human eye is sensitive to the properties regardless of location. Other graphical properties, such as crispness, transparency, and arrangement, have been proposed by MacEachren [1995, p. 275].

Greyscale, colour and texture are of particular interest when considering the use of visualization techniques on devices with limited display space and limited colour reproduction capabilities. On one hand, the type of the display (colour, greyscale or monochrome) determines which of the above graphical properties can be reproduced on the display; on the other hand visualization techniques make use of these selfsame graphical properties to encode information.

An overview of the relative effectiveness of the different graphical properties has been presented by Card *et al.* [1999, p. 30], based on the data presented by MacEachren [1995, p. 279]. Greyscale is most suitable for encoding ordinal data, followed by qualitative data. In contrast, colour and texture are best suited for encoding nominal data.

Graphical properties have been further separated by Card *et al.* into *spatial* and *object* properties based on which part of the brain they are processed in, and further according to the principal use: either the expression of the *extent* of a scale of values (*dissociative*) or the *differential* between objects (*associative*). Extent is expressed by size and greyscale, whereas differential is expressed by orientation, colour, texture, and shape.

According to Bertin [1981; p. 197, 231] dissociative variables have variable visibility while associative variables have constant visibility. In essence, this means that associative variables can be used to group objects that have the same dissociative variables, for example objects of same size. In contrast, objects of different size sharing the same value of some associative variable will be seen as separate entities.

A subset of the overview by Card *et al*. is presented in Table 3. It includes only those variables the use of which is dependent on the capabilities of the display elements found in typical handheld devices. In the table filled circle (●) indicates that the property is good for encoding that type of data, half-circle (◗) indicates that the property is marginally effective, and an open circle (○) indicates poor effectiveness.

Type of data

Principal use	**Object**	**Quantitative**	**Ordinal**	**Nominal**
Extent	Greyscale	◗	●	○
Differential	Colour	◗	◗	●
	Texture	◗	◗	●

Table 3. Relative effectiveness of different graphical (retinal) properties.

In this thesis it is assumed that the visualization techniques use such graphical properties for encoding purposes that are best suited for the data that is being visualized. Thus, the evaluation concentrates on determining whether the use of a given graphical property is essential for the given visualization technique, and whether the device in question provides support for the property. For example, greyscale can be used on a device that supports colour, but colour cannot be used on a device with a monochrome display.

The degree to which the graphical property is essential to the visualization technique can in some cases be evaluated based on the aforementioned types of data (quantitative, ordinal and nominal). For example, if the visualization technique uses greyscale to indicate the ordinal differential between objects, the visualization technique is only marginally suitable for monochrome displays, which can at best reproduce textures.

5.2.2 Resolution

Resolution, in addition to defining the size of a display, is used to determine if a given visualization technique can be used on a given device. For example, many visualizations use text or different kinds of shapes to distinguish objects, or distortion to emphasize the difference between the focus area and context area. This sets certain visualization technique specific minimum requirements for the size of the display area that is required to present the visualization in such a way that the objects are still legible and distinguishable.

However, it is often difficult to determine what the minimum or optimal resolutions are for a given visualization technique as user testing would be required to establish such boundaries empirically. Thus, by necessity, this criterion cannot be expressed in exact terms within the scope of this thesis and subjective approximations have to be made.

5.3 Interaction constraints

One of the key stages in the information visualization process is human interaction – the user’s ability to explore data interactively by rearranging it in different ways. Therefore effective visualizations should support interaction so that users are able to efficiently gain insight into the data they are exploring.

On the other hand, the devices that are used to present the visualizations offer many methods for interaction, such as physical or virtual keyboards, cursor keys, voice commands, touch-screen input, joysticks or keypads, to name a few features. These interaction methods may or may not be convergent with the methods required by the visualization techniques. This convergence, or the lack thereof, can be used as a measure to evaluate the suitability of a given visualization technique on a given device.

According to Card *et al.* [1999, p 231], interaction techniques can be divided into three categories: (1) interaction with data transformations, (2) interaction with visual mappings, and (3) interaction with view transformations. A categorization of a number of proven interaction techniques is presented in Table 4.

Modifies data transformation	**Modifies visual mappings**	**Modifies view transformation**
Dynamic queries	Dataflow	Direct selection
Direct walk	Pivot tables	Camera movement
Details-on-demand		Magic lens
Attribute walk		Overview + detail
Brushing		Zoom
Direct manipulation		

Table 4. Classification of interaction techniques.

To carry out the evaluation under this criterion, the visualization technique in question is first analyzed to determine which interaction techniques it utilizes. These interaction techniques are then compared to the interaction methods available in the device in question. The degree to which these interaction demands and available methods overlap is be used to determine the relative suitability of the technique for the device in terms of interaction constraints.

However, as was the case with resolution, it may be difficult to give exact estimations, since there is no simple way of mapping the interaction specific to the visualization technique to the specific interaction methods implemented in the device. Thus, the evaluation can at best say whether there is a feasible match between the visualization technique and the device.

5.3.1 Interactions with data transformations

Interactions with data transformations select which objects and variables are to be visualized.

Dynamic queries are a way of specifying desirable ranges for the variables in data tables, so that only the objects that fall within the specified ranges are displayed.

Direct walk is a method of navigating through the visualizations serially with the aim of searching for or modifying information. The use of a web browser is a typical example of direct walk interaction.

Details-on-demand can be used to ‘drill’ into a small subset of objects, which allows for a more efficient way of using the available space to show the attributes of selected objects. Figure 23 gives an example of the operation of details-on-demand.

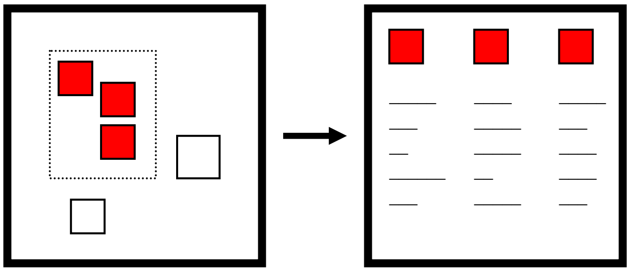

Figure 23. Illustration of details-on-demand.

Attribute walk is similar to direct walk, however in attribute walk the serial progression has the goal of finding other objects with similar attributes as the one that has been selected as a starting point.

Brushing is a method of using multiple visualizations for the different attributes of an object in such a manner. The objects selected in one visualization view are also selected or highlighted in the other visualization views. The operating principle of brushing is illustrated in Figure 24. The objects selected in the upper left-hand corner view are highlighted in the three other views.

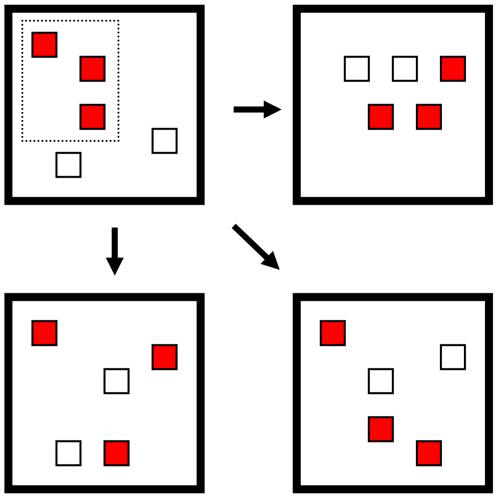

Figure 24. An illustration of brushing.

Direct manipulation in the context of information visualization (as opposed to direct manipulation in graphical user interfaces in general) allows for the modification of the parameters of data transformation itself, for example via a user interface component embedded into the visualization.

5.3.2 Interactions with visual mappings

Interaction can also be used to modify the visual mappings between the data tables and visual structures.

Dataflow techniques use a flow chart metaphor to map the data into visual form. One application area is visual programming, in which the structure of a software program can be visualized by using node-link diagrams, in which the nodes represent programming operations and links the execution (flow) of the program from one stage to the next.

Pivot tables is a technique found in many spreadsheet programs. It lets the user quickly manipulate the mapping of the data to the rows and columns on the spreadsheet. For example, in Microsoft Excel a pivot table report can be used to combine and compare large amounts of data. Rows and columns can be rotated to see summaries of the source data, which can be used to make comparisons, detect emerging patterns in and relationships in the data, or analyze trends.[18]

5.3.3 Interactions with view transformations

The final stage in which interaction can have an effect on the visualization is the modification of view transformations.

Direct selection involves the methods for selecting and highlighting objects, or groups of objects, in the visualization. This selection is used to identify the desired objects that are the targets of some other interaction.

Camera movement is used to move the position of the viewpoint, through which the visualization is viewed. Examples of camera movement include panning in two-dimensional visualizations or the change of viewing angle in three-dimensional visualizations.

Magic lenses, zoom and pan and *overview + detail*, which are introduced in section 2.4, can be considered as interaction techniques in addition to being independent visualization techniques in their own right.

5.4 Implementation constraints

As this thesis is concerned with visualizing information on actual devices, a number of implementation related issues must also be considered. On these devices the abstract ideas of visualization techniques are implemented as software programs that by nature require that the device has enough processing power, storage space, and memory to run the said program as intended. In terms of mobile devices connectivity as an extension to storage space must also

[18] PivotTable Reports 101 for Excel 2002: *http://office.microsoft.com/assistance/2002/ articles/xlconPT101.aspx*

be considered, given that it is an inherent aspect of mobility. Additional constraints are imposed by the display technology used. For example, some techniques make use of transparency, three-dimensional graphics or overlapping windows, which as of yet are not supported by most handheld devices. These constraints can be categorized as *hardware* constraints.

Additionally, the application programming interface (API) of the device's operating system governs how visualizations can be developed on the device. The API might, for example, limit access to certain parts of the operating system or the physical features of the device. The selection of user interface components (*widgets*) is more limited on mobile phone and PDA operating systems than on desktop operating systems. These constraints can be categorized as *software* constraints.

In this thesis certain assumptions are made when evaluating the implementation constraints related to specific visualization techniques and/or devices. This is due to the fact that it would be exceedingly difficult to accurately gauge the impact of these constraints on the suitability of a visualization technique without analysing the implementation in detail. This is not often possible, especially in the case of commercial visualization applications.

5.5 Summary

The constraint categories presented in the previous sections form a framework within which the suitability of visualization techniques for use on small display devices can be evaluated.

The type of the device in question establishes what resources, such as display type or interaction methods, the device can offer for the visualization techniques. The visual properties and interaction techniques that the visualization techniques make use of can be thus evaluated against the affordances of the device in question.

Table 5 illustrates this two-way relationship between the devices and visualization techniques by constraint category.

Visualization makes use of	Constraint	Device offers
Greyscale Colour Texture Resolution	*Visual*	Greyscale Colour Monochrome Resolution
Data transformations Visual mappings modifications View transformation modifications	*Interaction*	Input methods
Processing power Memory Storage space Display technology User interface components	*Implementation*	Processing power Available memory Available storage space Connectivity Display technology Programming interfaces

Table 5. The relationship between device capabilities and visualization requirements as defined by the evaluation constraints.

6. Evaluation of visualization techniques

The evaluation will concentrate on determining how well the types of visualization techniques presented in section 3.1 are suited for use on the type of devices presented in section 4.2. The evaluation will be carried out using the framework presented in chapter 5.

To do this, a three-step scale is used across the different constraint categories. If there is a direct match between the visualization constraints and the device constraints, the visualization technique is well-suited for use on the device in terms of that constraint. If there is a marginal match, meaning that requirement can be met at least partially, the visualization technique is marginally suited for use. If there is no match, meaning that the requirement cannot be met at all, the visualization technique cannot be used on the device. As discussed in the previous sections concerning the constraints, it depends on the visualization technique how exact an assessment of the match between the technique and the device can be made.

The result of the evaluation will be a presented in the format presented in Table 6. Each visualization technique is given a value that represents the match between the visualization requirement and the device capability in each category of evaluation criteria. Filled circle (●) indicates that the property is good match between the requirements, half-circle (◗) indicates a marginal match, and an open circle (○) indicates poor match. A dash (-) is used to denote a situation in which the constraint does not apply to the given visualization technique.

	Visual constraints		Interaction constraints			Implementation constraints	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Technique 1	●	◗	○	-			
...							

Table 6. A tabular presentation of the evaluation results by technique.

The actual evaluation consists of two parts: an evaluation of the visualization techniques on a low-end device (Nokia 6310i) and an evaluation on a high-end device (Palm™ m515). A recapitulation of the findings of both evaluations will be presented in summary.

6.1 Low-end display devices: Nokia 6310i

Generally, low-end display devices such as the Nokia 6310i mobile phone impose very strict limitations for visualizations in terms of display space, interaction capabilities, and implementation possibilities. Due to these reasons there has not been a lot of research done in the field of designing visualization techniques exclusively for mobile phones. Related research, for example by Belz *et al.*[1998], Encarnação *et al.* [1995] and Rauschenbach *et al.* [2000], has focused on optimizing the delivery and use of static graphics and multimedia on mobile handsets, with very few actual prototypes.

Given this lack of existing visualization applications for mobile phones, this evaluation considers the suitability of existing visualization techniques in principle only. That is, the evaluation is focused on the underlying concepts of the visualization techniques and the possibilities of implementing functional visualizations based on these concepts on the small display devices.

6.1.1 Focus + context techniques

Clearly, one of the largest problems when dealing with low-resolution displays is the ability to provide adequate context in addition to focus. Due the lack of overall display resolution it might prove difficult to show enough of the context to the user for the context to have any real use.

Leung and Apperley [1994] divide focus + context visualization techniques in their taxonomy as follows: Polyfocal Display, Bifocal Display, different fisheye views (discussed in 3.1.2) and the Perspective Wall. This distinction is used as a basis for the evaluation of focus + context techniques. Fisheye views are excluded, however, as they are discussed in conjunction with other mainly suppression-based techniques. The Hyberbolic Browser is discussed in the context of focus + context techniques as an example of a non-standard two-dimensional distortion technique.

The focus + context techniques presented in section 3.1.1 are not very sensitive to the visual constraints defined by graphical properties as a whole, as the main effective means of visualization is distortion, instead of, for instance, colour coding. In terms of visual constraints, the main limiting factor to the usefulness of focus + context techniques on low-end displays is *resolution*. Polyfocal Display techniques make use of multiple foci, which effectively limits the area available for a single focus area. Due to the low resolution and the limitations of the display hardware, the use of a single focus technique like Bifocal Display would be more advisable. The use of perspective distortion, and hence reduced legibility, employed in the Perspective Wall limits its suitability in terms of resolution. The Hyperbolic Browser, on the other hand, effectively requires more display space than is available on the low-end small displays.

Only the last interaction category, interactions with view transformations, is applicable to focus + context techniques as such. The interactions consist of moving the focus along the cardinal axes to bring items of interest into the view and selecting items for perusal. The Nokia 6310i offers adequate means for accomplishing both the view manipulation and the selection tasks.

The main implementation constraint that comes into play with focus + context techniques is processing power, mainly due to the use of distortion, which is typically based on complex algorithms. However, it is probable that current devices have enough processing power to perform reasonably well with focus + context techniques. The Polyfocal Display and the Hyperbolic Browser are the techniques that are the most processing power dependent, the reasons being the use of multiple foci and mathematical (Euclidean to hyperbolic space mapping) calculations.

A summary of the evaluation of focus + context techniques on a low-end display device is presented in Table 7.

	Visual constraints		**Interaction constraints**			**Implementation constraints**	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Polyfocal Display	●	○	-	-	●	◗	-
Bifocal Display	●	◗	-	-	●	●	-
Perspective Wall	●	○	-	-	●	◗	◗
Hyperbolic Browser	●	○	-	-	●	◗	-

Table 7. Evaluation of focus + context techniques on a low-end display device.

6.1.2 Suppression techniques

Suppression techniques are in this evaluation divided into (1) non-graphical fisheye views, (2) graphical fisheye views, (3) z-thru mapping, and (4) mixed techniques.

When considering graphical properties as constraints, it seems that z-thru mapping techniques are be the most dependent on the use of colour as a visual indicator of information due to the type of information (geographical, weather and map data) the techniques are typically used to present. The other techniques typically make no special use of colour, so in that regard they should be fairly suitable for use on a mobile phone display. Resolution, however, is a limiting factor in almost all cases, the exception being non-graphical fisheye views. Graphical fisheye views employ distortion, which impairs legibility on such a small display as the one used in the Nokia 6310i, whereas z-thru mapping and mixed techniques require more display space than is available in order to present information in enough detail.

None of the above techniques is concerned with data transformation as such. However, the modification of visual mappings via suppression controls is an integral part of fisheye view techniques, as well as z-thru mapping and mixed techniques. With fisheye views and mixed techniques there is a need to modify the parameters of the DOI functions, while with z-thru mapping the suppression is controlled by modifying the transparency of layers. It is evident that the keypad and the menu buttons found in 6310i offer only marginal means for controlling the suppression parameters. In terms of view transformations the situation is better. All of the techniques mentioned here make use direct selection as the means to shift the focus of the visualization, which can be accomplished with the interaction methods offered by the mobile phone in question.

Implementation constraints affect the fisheye view techniques to some extent. The required DOI calculations might demand considerable processing power, even more so with graphical views that also require distortion calculations. Mixed techniques use both DOI and distortion, which also increases the processing requirements. Z-thru mapping is probably the least processing intensive as its main visualization method is the use of transparency; however the display technology and the limitations of the Java™ programming environment utilized in the Nokia 6310i quite likely preclude the use of layers.

A summary of the evaluation of suppression techniques on a low-end display device is presented in Table 8.

	Visual constraints		**Interaction constraints**			**Implementation constraints**	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Non-graphical fisheye views	●	●	-	◗	●	●	-
Graphical fisheye views	●	○	-	◗	●	●	-
Z-thru mapping	◗	○	-	◗	●	○	○
Mixed techniques	●	○	-	○	●	●	-

Table 8. Evaluation of suppression techniques on a low-end display device.

6.1.3 Magic lenses

Magic lenses is a technique that is used to reveal additional information on the underlying visual structure. It consists of the underlying visual structure, a smaller view area that is moved over the visualization, and the parameters of the lens that affect how the additional information is generated.

In terms of graphical properties magic lenses is not bound by any particular property. Of course, the use of colour would bring added value to the use of the visualizations as it would make it easier to highlight the objects under the lens, or the additional information generated by the lens. Resolution, however, is again a limiting factor. The display on the Nokia 6310i is too small to effectively accommodate both the visual structure and the magic lens view in such a manner that the overall visualization would still be legible.

View transformation is the only interaction constraint that comes into play when using magic lenses. View transformation is used to both modify the parameters of the lens and to move the lens over the visual structure. The modification of parameters is typically handled by using controls embedded to the magic lens itself. However, in a mobile phone the changes would have to be carried through a series of menu selections or separate views, which severely reduces the ease of use.

Magic lenses, especially in the Fishkin and Stone [1995] implementation, make use of multiple overlapping lenses. It would probably be difficult to implement an effective solution on the 6310i mobile phone, given the limitations of the Java™ programming environment and the display hardware. A more straightforward approach would be to use a single layer only, which would be easier to implement, and which would also work better in terms of the resolution constraints.

A summary of the evaluation of magic lenses on a low-end display device is presented in Table 9.

	Visual constraints		Interaction constraints			Implementation constraints	
	Graph. Prop.	Resol.	Data Trans.	Visual Map.	View Trans.	Hard-ware	Software
Magic lenses	●	○	-	-	◗	◗ / ○	◗

Table 9. Evaluation of magic lenses on a low-end display device.

6.1.4 Riffling techniques

Riffling techniques trade space for time by presenting information in a sequential manner in the same limited space. This makes riffling a favourable candidate for use on mobile phones where display space is at a premium.

The impact of the graphical constraints depends on the type of information being presented. The monochrome display of the Nokia 6310i is unsuitable for presenting images as they are presented by such riffling techniques as the PolyNav™ previewer [Wittenburg *et al.*, 1998] or the RSVP Browser [De Bruijn *et al.*, 2001]. The RSVP Browser is an extension of the RSVP principle suggested by De Bruijn and Spence [2000]. It is intended for browsing World Wide Web pages on mobile devices and PDAs. The RSVP Browser requires, however, a colour display and more display space than Nokia 6310i is able to provide. An illustration of the RSVP Browser is presented in Figure 25.

As the functionality and features of the RSVP Browser are similar to that of the PolyNav™ system, no differentiation is made in their evaluation.

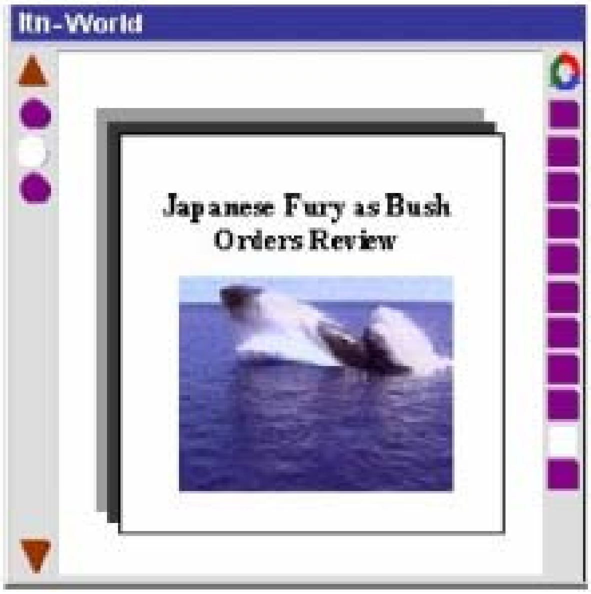

Figure 25. The RSVP Browser by De Bruijn *et al*. [2001].

Riffling techniques could be used, however, to present text or simple graphics (icons) that make use of textures. Provided that this kind of information is being presented, riffling techniques are highly suitable for low-end small displays. Resolution is not such a great concern with riffling, as it is with focus + context or suppression techniques, because no distortion or separate context and focus areas are used.

Interaction constraints are limited to view transformations, as the browsing itself, and the modification of the browsing parameters, is the main forms of interaction. These interaction tasks should be easily accomplished by using the numerical keypad and the control keys of the Nokia 6310i.

Implementation constraints are likely not to be a major concern, given that no extensive use of algorithms or mathematical calculations is required by the riffling techniques as such. For example, the processing required to turn HTML pages into a suitable image format is performed in the PolyNav™ system on the server side, from where the resulting image is transmitted to the previewer prototype. Storage space might, however, become a concern if the images are to be stored (cached) on the device. Image caching would, in fact, be a desirable solution due to the limited data transmission capabilities of the 6310i.

A summary of the evaluation of riffling techniques on a low-end display device is presented in Table 10.

	Visual constraints		**Interaction constraints**			**Implementation constraints**	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Riffling	●	●	-	-	●	◗	-

Table 10. Evaluation of riffling techniques on a low-end display device.

6.1.5 Zoom and pan techniques

Zoom and pan techniques are, along with riffling, the more promising approaches to solving the presentation problem on small displays. Instead of trading space for time, zoom and pan techniques provide more space through the use of zooming. The prominent technique to consider is semantic zooming, as its application areas are more varied than those of traditional geometric zooming interfaces.

Graphical properties are not a limiting factor when it comes to semantic zooming, as the technique is not bound to any given graphical property, be it colour, greyscale or texture. The lack of resolution, nevertheless, creates problems even with semantic zooming. With complex visual structures it brings about the need to pan the view across the visual structure, or alternatively zoom excessively deep into the visual structure. Panning across the visual structure occurs if the amount of zooming levels is restricted to a finite number, whereas zooming deep into the visual structure occurs if there is an infinite number of zooming levels available.

Interaction constraints do not come into play with zooming and panning techniques. The interaction tasks are generally limited to view transformations, such as zooming into and out of the visual structure and/or panning the view across the visual structure. These tasks can be accomplished quite adequately by using the numerical keypad of the Nokia 6310i.

In order to facilitate quick zooming into and out of the visual structure, it might be advisable to use only a finite number of discrete zooming levels and keep the details of the separate zooming levels in device memory rather than recalculate them each time a zooming action is triggered. This is mainly due to the fact that the calculations associated with zooming can be quite tasking, as is the case for example with the Pad++ system [Bederson and Hollan, 1994], and could thus unnecessarily slow down the interaction. This might be a concern in a device such as the Nokia 6310i, given the limited memory and processing power available to the application. One viable alternative would be to save the details for all the possible zooming levels and views as precalculated images on a server that is accessible by using the data transmission options of the 6310i.

A summary of the evaluation of zoom and pan techniques on a low-end display device is presented in Table 11.

	Visual constraints		**Interaction constraints**			**Implementation constraints**	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Semantic zooming	●	◗	-	-	●	◗	◗

Table 11. Evaluation of zoom and pan techniques on a low-end display device.

6.1.6 Summary

Based on the evaluation of the visualization techniques, it seems that the Bifocal Display, non-graphical fisheye views, riffling and semantic zooming are the most suitable techniques on low-end small display devices such as the Nokia 6310i.

Each of the four techniques has its strengths and weaknesses. Riffling and semantic zooming, while not expressly designed for use on small displays, make a trade-off between the available display space and other dimensions, such as time or the depth of zooming. When compared to these two techniques, the suitability of the Bifocal Display is reduced to a degree by the use of distortion, while non-graphical fisheye views is held back by the added burden of having to define the parameters of the suppression functions. On the other hand, the application domain of riffling is curtailed by the limited display area, whereas semantic zooming suffers from a number of implementation issues relating to the storage and processing of the views at different zooming levels.

The rest of the techniques are not suitable for low-end displays, mainly due to the limitations posed by the resolution of the Nokia 6310i display. Additionally, the suitability of z-thru mapping is also limited in terms of interaction and implementation constraints.

6.2 High-end display devices: Palm™ m515

The Palm™ m515 PDA is an example of a high-end small display device. It has several advantages over the low-end display devices, including a larger resolution display, a wider array of input methods, and greater processing and storage capabilities.

It can be expected, therefore, that a greater portion of the visualization techniques will be suitable for use on the Palm™ m515, merely due to the increased capabilities of the device. The exact situation will be assessed in the following sections.

6.2.1 Focus + context techniques

The added display resolution in the Palm™ m515 device over that of the Nokia 6310i mobile phone alleviates some of the problems that were associated with focus + context techniques in the context of low-end display devices. This is in part confirmed by the fact that applications of focus + context techniques for similar Palm™ devices already exist. Björk *et al.* have presented several applications of the focus + context principle, such as *WEST* [1999], a web browser for small terminals, and *PowerView* [2000], an application to provide access to common information on PDAs. Both approaches utilize the *Flip Zooming* visualization technique suggested by Holmquist [1997].

Flip Zooming is reminiscent of the Document Lens visualization technique, and is thus an extension of the Bifocal Display principle. The visualization area consists of several views, one of which is selected as the focus and presented in the middle in larger size. Each view corresponds to an object in a sequence, for example web pages or images in a folder. Flip Zooming uses linear scaling instead of continuous spatial distortion in the context area, which retains the legibility of the content of the individual views.

In the WEST browser Flip Zooming is applied to a hierarchical structure of information consisting of pages and decks, which are collections of individual pages. In addition, the application employs semantic zooming by allowing the user to change the content type of the views (thumbnail images – keywords – hyperlinks). The interface of the WEST browser is illustrated in Figure 26.

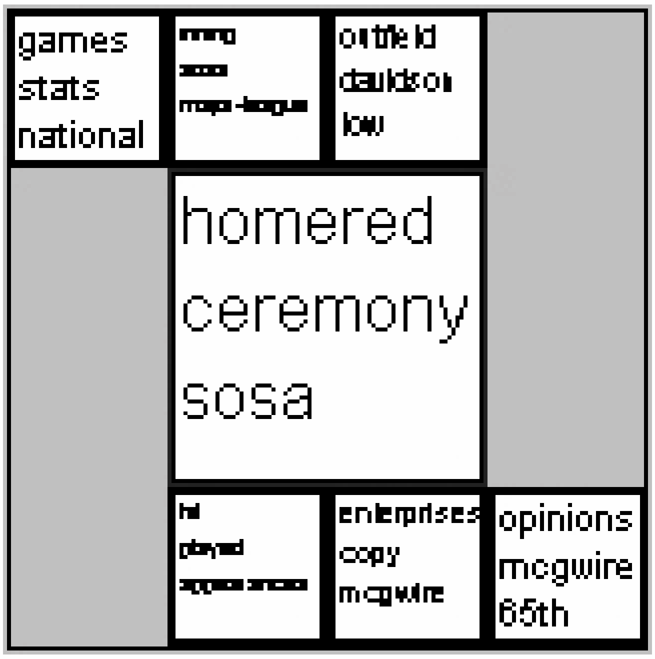

Figure 26. WEST interface in keyword mode with focus on the 4th card in the sequence [Björk *et al*., 1999].

As was previously mentioned, graphical properties do not come into play as such when considering the properties of focus + context techniques. Resolution does nevertheless pose some problems, especially with the Polyfocal Display (multiple foci) and the Hyperbolic Browser (hyperbolic distortion) techniques. Perspective Wall is also still held back by the lack of overall display space, mainly due to its use of perspective distortion. Bifocal Display continues to be the best technique in terms of resolution.

Interaction constraints are limited to those involving view transformations, including the movement of the viewpoint and the selection of objects, and, for example, zooming control in the case of Flip Zooming techniques. These actions can be quite well performed by using the hardware buttons and the stylus pointing device provided by the Palm™ m515 handheld.

The main implementation constraint observed in the case of mobile phones was processing power, mainly due to the fact that the majority of focus + context techniques presented here make use of distortion. The algorithms associated with computing the distortion effects are probably not as prohibitive when considering their use on handheld devices such as the Palm™ m515, thanks to the added processing power. In addition, the m515 has copious amounts of both memory and storage space when compared to the Nokia 6310i.

A summary of the evaluation of focus + context techniques on a high-end display device is presented in Table 12.

	Visual constraints		**Interaction constraints**			**Implementation constraints**	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hardware**	**Software**
Polyfocal Display	●	◗	-	-	●	●	-
Bifocal Display	●	●	-	-	●	●	-
Perspective Wall	●	◗	-	-	●	●	-
Hyperbolic Browser	●	◗	-	-	●	●	-

Table 12. Evaluation of focus + context techniques on a high-end display device.

6.2.2 Suppression techniques

An example of a suppression-based visualization application for high-end small display devices is the *Power Browser* by Buyukkokten *et al.* [2000]. While the Power Browser is not strictly speaking a visualization technique as such, it employs similar techniques for data suppression as are found in existing suppression-based visualization techniques, such as the logical (non-graphical) fisheye view by Furnas [1986]. In Power Browser the user is presented with a tree-like view of the structure of the web site that is currently being browsed.

This provides the user with both location information (focus) and information on the neighbouring navigation environment (context) at all times.

Figure 27. The Power Browser navigation interface.

Experiments showed significant gains in browsing speed and reduction in the required amount of pen movements when the Power Browser was compared to three traditional PDA browsers. The traditional browsers showed web content on the small display as similarly as they would on a desktop browser.

Extensions of z-thru mapping techniques based on the use of transparency have been presented by Kamba *et al.* [1996]. Their approach was to apply semi-transparency to widgets (the control objects in a user interface) in order to use the available display space more efficiently. In this regard their system resembles the Macroscope prototype by Lieberman *et al.* [1994].

Their implementation also featured variable delay by which objects in different layers of the screen responded to the users' selections. Experiments showed that their technique could help to increase the efficiency of display space usage on small display devices. Kamba *et al.* also suggest that their semi-transparency/delayed response model could be combined with techniques that are aimed at showing information content more efficiently, such as magic lenses or Table Lens. The experiment was run on a Macintosh personal computer with an active screen region of 320 × 240 pixels, but it can be assumed that the underlying principle of semi-transparency and delayed response would also work on a the Palm™ m515 display, albeit perhaps not as efficiently as on a larger desktop computer display. The interface of their prototype is presented in Figure 28.

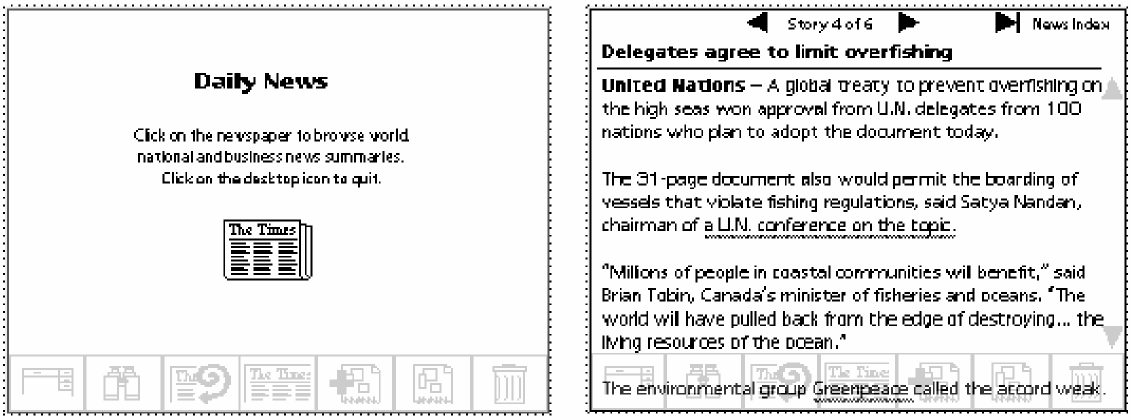

Figure 28. The interface of the prototype by Kamba *et al.*

In the context of mobile phone displays such as the one in Nokia 6310i it was observed that colour imposes constraints on the use of the z-thru mapping technique. In the case of the Palm™ m515 such limitations do not exist due to the high colour display used. Resolution is still a limiting factor but not as extreme as in the case of low-end displays. Non-graphical fisheye views can be presented quite successfully, as is evidenced by the Power Browser interface; however the rest of the techniques continue to suffer from the lack of available display space. Graphical fisheye views, mixed techniques and z-thru mapping techniques could nevertheless still be used on the m515 display resolution.

Interaction constraints as a whole do not come into play with the Palm™ handheld. The problems associated with modifying the visual mappings (fisheye views, mixed techniques and z-thru mapping), DOI functions (fisheye views) or layer transparency (z-thru mapping) on low-end display devices probably do not exist, given that Palm™ operating system and the m515 device offer mechanisms such as dialogs, different selection widgets and direct manipulation to control the different parameters. As with the 6310i, no problems should exist when interacting with view transformations.

The different fisheye view techniques and mixed techniques seem to be the least affected by implementation constraints, since the main constraint in their case is processing power, of which the Palm™ m515 should possess enough to deal with the necessary DOI and distortion calculations. However, with z-thru mapping the situation is different. Z-thru mapping techniques make use of transparency and layers, which might be quite difficult to implement adequately on the m515. They could, however, be simulated by extending the functionality of the user interface components available in the Palm™ operating system.

A summary of the evaluation of suppression techniques on a high-end display device is presented in Table 13.

	Visual constraints		Interaction constraints			Implementation constraints	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Non-graphical fisheye views	●	●	-	●	●	●	-
Graphical fisheye views	●	◗	-	●	●	●	-
Z-thru mapping	●	◗	-	●	●	◗	◗
Mixed techniques	●	◗	-	●	●	●	-

Table 13. Evaluation of suppression techniques on a high-end display device.

6.2.3 Magic lenses

Neither graphical properties nor resolution constitute any major obstacles for using techniques based on the magic lenses principle on the m515 device. Instead, the colour display of the m515 offers possibilities for using colour coding to distinguish the objects within the active area of the lens. The added display area enables, for example, the use of text labels or multiple overlapping lenses.

The view transformation interactions required for manipulating the parameters and position of the lens are also easier to handle on the PDA device. It is conceivable to picture an application in which the parameters are modified by using controls embedded on the magic lens itself by using the stylus.

Technical considerations, however, form an obstacle for the successful usage of magic lenses. While the programming environment on the Palm™ device offers more features than its counterpart on the 6310i, it is doubtful as to whether the device itself supports overlapping or transparent widgets adequately enough. The added resolution, in any event, makes it possible to consider the use of several lenses through an implementation that simulates lens transparency.

A summary of the evaluation of magic lenses on a high-end display device is presented in Table 14.

	Visual constraints		**Interaction constraints**			**Implementation constraints**	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Magic lenses	●	●	-	-	●	◗	◗

Table 14. Evaluation of magic lenses on a high-end display device.

6.2.4 Riffling techniques

The techniques introduced when discussing the use riffling techniques on mobile phone displays, the PolyNav™ previewer and the RSVP Browser, do not in this context suffer from the problems associated with graphical properties. It is even possible that the Java™ applet prototype of the PolyNav™ previewer could be used "as is" on an m515 equipped with an internet connection. The added resolution also removes the limitations regarding the type information that can be visualized. For example, it is reasonable to conjecture that the image-based web browsing example by De Bruijn *et al*. [2001] could be fitted on a display of the size of the Palm™ m515.

As noted previously, interaction constraints do not pose problems for riffling techniques. The situation is further alleviated by the introduction of direct manipulation via the stylus, as this allows for the straightforward selection of links and objects (RSVP Browser), or the modification of visualization parameters (PolyNav™ previewer).

The image caching and transmission issues discussed in conjunction with the Nokia 6310i are still present when examining the Palm™ device. However, given that the m515 features 16 megabytes of storage memory, which can be extended by using add-in memory cards, the issues with image caching are not as severe as they would be on a mobile phone such as the 6310i.

A summary of the evaluation of riffling techniques on a high-end display device is presented in Table 15.

	Visual constraints		Interaction constraints			Implementation constraints	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Riffling	●	●	-	-	●	●	-

Table 15. Evaluation of riffling techniques on a high-end display device.

6.2.5 Zoom and pan techniques

As mentioned earlier, zoom and pan techniques are one very appealing approach to solving the presentation problem. On mobile phones the lack of overall display space precludes the efficient use of zoom and pan techniques but the situation is different on the m515. As demonstrated by Björk *et al*. [1999, 2000], zooming can be very useful especially on high-end small displays.

Other zooming approaches, like the infinitely zoomable visualization interface of *Pad++*, by Bederson and Hollan [1994], are also worth considering. The basic principle of the Pad++ interface is to provide a smoothly zoomable interface into a very large set of data. It is a prime example of *semantic zooming*, in which objects change their representation depending on the zooming level. For example, in a file system visualization the files within a directory are shown as coloured dots within a square frame on the higher level, but as the user zooms into a given file, its contents are loaded into the frame where they be edited. Figure 29 provides an example of the operation of continuous zooming in Pad++ system. In the example the user is zooming into a graphical drawing from left to right and top to bottom.

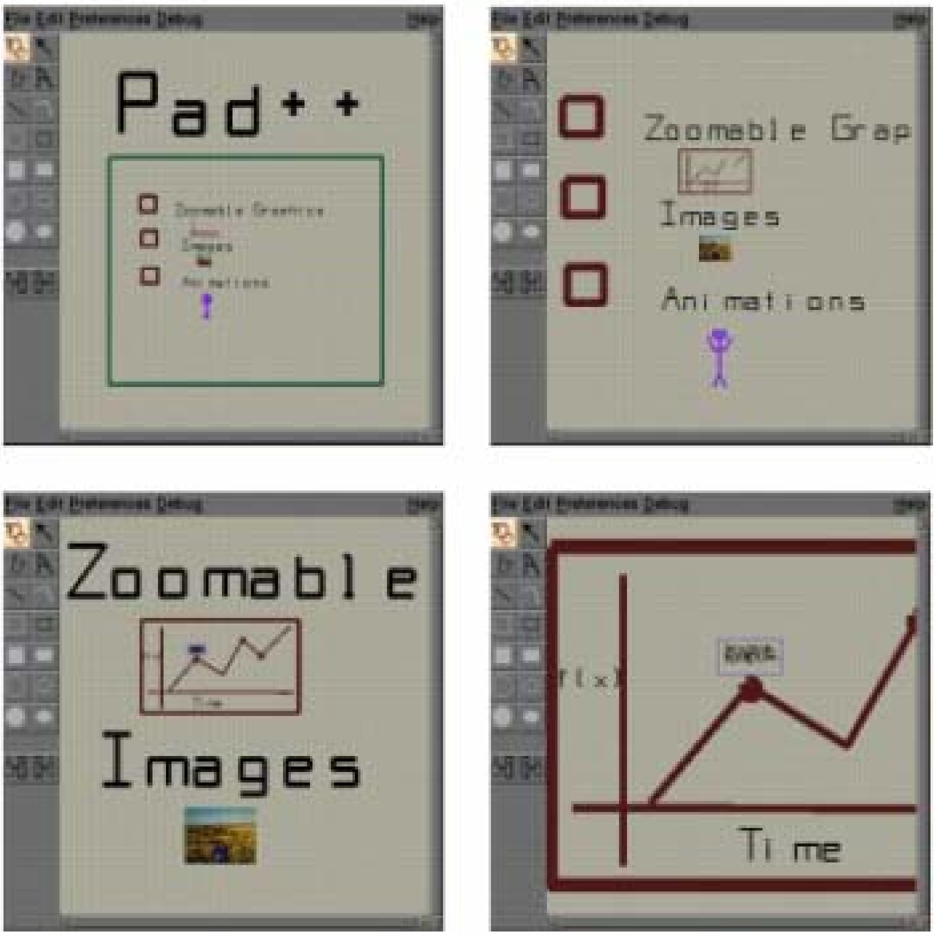

Figure 29. An example of the operation of the Pad++ interface.

The issues associated with visual constraints on low-end small displays, the need for excessive panning or zooming due to lack of display space, are remedied on the Palm™ m515 device. The reason is on one hand the additional resolution, and on the other hand the more sophisticated interaction methods available, such as the direct manipulation of different widgets with the stylus.

Also, the colour display increases the possible application domains of the zoom and pan techniques.

Interaction constraints are not a limiting factor with the m515, given the versatility of the stylus and the physical and virtual buttons available on the device. These can be used to good effect to carry out the zooming and panning tasks associated with zoom and pan techniques.

The implementation concerns that were raised when considering the use of zoom and pan techniques on the Nokia 6310i also apply to the Palm™ m515 device. While it might be possible to carry out the calculations required for rendering the views at different zooming levels, a more favourable approach would be to use a discrete number of zooming levels with precalculated views. The views could be stored, for example, on a dedicated server, as is done in the WEST browser implementation by Björk *et al.* [1999]. In any event, implementation constraints complicate the use of zoom and pan techniques.

A summary of the evaluation of zoom and pan techniques on a high-end display device is presented in Table 16.

	Visual constraints		**Interaction constraints**			**Implementation constraints**	
	Graph. Prop.	**Resol.**	**Data Trans.**	**Visual Map.**	**View Trans.**	**Hard-ware**	**Software**
Semantic zooming	●	●	-	-	●	◗	◗

Table 16. Evaluation of zoom and pan techniques on a high-end display device.

6.2.6 Summary

The evaluation of the visualization techniques on a high-end display device suggests that the majority of the techniques are suitable for use on the Palm™ m515 device.

Three groups stand out from the pool of proposed techniques. The non-graphical fisheye views, the Bifocal Display and riffling are clearly the most workable when one considers the application domain. These techniques do not necessarily require colour, make little to no use of distortion and are easy to interact with using the methods provided by the m515.

The second group includes graphical fisheye views, mixed techniques, semantic zooming, the Perspective Wall, the Hyperbolic Browser, and the Polyfocal Display. The problems with the focus + context techniques are mainly associated with the overall lack of display space. While the m515 features over four times more pixels than the 6310i, there still is not enough display area to make effective use of the above techniques. Semantic zooming, on the other hand, is still probably somewhat weighed down by the limited processing capacity of the m515.

The last group includes magic lenses and z-thru mapping. The main problems with these techniques have to do with implementation constraints. While it would be possible to make use of magic lenses and z-thru mapping techniques in terms of the visual and interaction constraints, it might be fairly difficult to come up with a satisfactory implementation. Z-thru mapping is based on the notion of transparent layers, while magic lenses employ overlapping (transparent) windows. It may be possible to overcome these problems by implementing specialized widgets that simulate transparency and layers. The advances in display and graphics acceleration technology will naturally in time negate any such problems.

6.3 Evaluation results

The possible solutions for overcoming the presentation problem were evaluated in the previous sections by using a framework that attempted to classify the different features of both visualization techniques and devices, and how well these features match one another in practice.

In the context of low-end display devices four distinctly suitable techniques could be distinguished, being the Bifocal Display, non-graphical fisheye views, riffling, and semantic zooming. What is shared by these techniques is the fact that their properties match the constraints of the device in question fairly well. They are viable solutions given the low display resolution, limited interaction capabilities and implementation possibilities of low-end display devices.

In the context of high-end display devices three groups of suitable techniques could be discerned, containing over two thirds of the proposed visualization techniques. It was not surprising to find out that almost all of the techniques became viable solutions on a high-end display device. It can thus be presumed that the situation will further improve as display technology improves.

As might have been expected, the techniques that were found suitable for low-end display devices were also found suitable for high-end display devices. This result can be attributed to the more advanced features of the high-end small display device, including a colour display with increased resolution, increased processing and storage capabilities, and the advanced interaction methods available.

It must be stressed, though, that these results are not definitely conclusive given the nature of the evaluation, and that prototyping and empirical testing would likely yield more tenable results. Through testing it would be possible to determine how well the different techniques perform in a real context of use. However, such an extensive evaluation process was not a feasible approach considering the nature of this thesis.

7. Conclusions

This thesis addressed issues relating to the presentation problem by introducing and evaluating several solutions suggested in the literature.

The solutions were bound on the practical plane by selecting two devices that acted as the context for the evaluation. The major factor in choosing the devices was their display resolution, which is the most practical way of distinguishing between different categories of devices. One device, the Nokia 6310i mobile phone, represented a low-end display device category and the other, the Palm™ m515 PDA, a high-end display device category. Other contributing factors to the selection of devices were the interaction techniques as well as the connectivity features available in the devices.

The evaluation was carried out by using a framework that was comprised of the features of the proposed visualization techniques and the target devices. The framework mapped the visual, interaction and implementation features employed by the visualization techniques to the practical constraints imposed by the devices.

Each visualization technique was ranked on a three step scale based on an assessment that measured how well the technique fulfils the demands of the target device. As a result, the visualization techniques could be ordered based on their rankings.

In terms of low-end display devices, the greatest role was played by visual constraints, namely colour and resolution. Interaction and implementation constraints were clearly a limiting factor in many cases, as well. In terms of high-end display devices, implementation constraints proved to have more effect than the other constraints. This was mainly due to the improved visual and interaction properties of the device in question.

The contributions of this thesis are threefold. First, it presents a crucial problem in the field of information visualization and possible solutions to the problem, based on relevant research results reported in the literature. Second, it introduces a framework for the evaluation of the proposed solutions. Finally, it reports the results that were gained by using the framework to evaluate the solutions. These results, and the proposed framework, can be used as guidelines when selecting the appropriate techniques for visualizing information on small display devices.

References

[Bederson and Hollan, 1994] Ben B. Bederson and J.D. Hollan, Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. In *Proceedings of UIST '94*, ACM, pp. 17-26.

[Bederson, 2000] Ben B. Bederson, Fisheye Menus, In *Proceedings of UIST 2000*, ACM, pp. 217-225.

[Belz *et al*, 1998] C. Belz, H. Jung, L. Santos, R. Strack, and P. Latva-Rasku, Handling of Dynamic 2-D/3-D Graphics in Narrow-Band Mobile Services. In *Virtual Worlds on the Internet*, IEEE Computer Society 1998, pp. 147-156.

[Bertin, 1981] Jacques Bertin, *Graphics and Graphic Information Processing*. Walter de Gruyter, 1981.

[Bertin, 1983] Jacques Bertin, *Semiology of Graphics: Diagrams, Networks, Maps* (W.J. Berg, Trans.). University of Wisconsin Press, 1983.

[Björk *et al.*, 1999] S. Björk, L.E. Holmquist, J. Redström, I. Bretan, R. Danielsson, J. Karlgren and K. Franzén, WEST - A Web Browser for Small Terminals. *In Proceedings of UIST'99*, ACM, pp. 187-196.

[Björk *et al.*, 2000] S. Björk, J. Redström, P. Ljungstrand and L.E. Holmquist, POWERVIEW: Using information links and information views to navigate and visualize information on small displays. *Handheld and Ubiquitous Computing*, Springer Verlag, 2000, pp. 46-62.

[Buyukkokten *et al.*, 2000] O. Buyukkokten, H. Garcia-Molina, A. Paepcke and T. Winograd, Power Browser: Efficient Web Browsing for PDAs. In *Proceedings of CHI 2000*, ACM, pp. 430-437.

[Card *et al.*, 1999] Stuart K. Card, Jock D. Mackinlay and Ben Shneiderman, *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, 1999.

[Colby and Scholl, 1991] G. Colby and L. Scholl, Transparency and blur as selective cues for complex visual information. *SPIE Vol. 1460, Image Handling and Reproduction Systems Integration*, 1991, pp. 114--125.

[De Bruijn and Spence, 2000] O. De Bruijn and R. Spence, Rapid Serial Visual Presentation: A space-time trade-off in information presentation. In *Proceedings of AVI 2000*, pp. 189-192.

[De Bruijn *et al*., 2001] O. De Bruijn, R. Spence and M.Y. Chong, RSVP Browser: Web Browsing on Small Screen Devices. In *Proceedings of Mobile HCI 2001: Third International Workshop on Human Computer Interaction with Mobile Devices*, M.D. Dunlop and S.A. Brewster (Eds.).

[Encarnação *et al.*, 1995] J.L. Encarnação, M. Frühauf, and T. Kirste, Mobile Visualization -- Challenges and Solution Concepts. In *Proceedings of CAPE'95*, Beijing, China, May 1995.

[Fishkin and Stone, 1995] K. Fishkin and M. Stone, Enhanced Dynamic Queries via Movable Filters. In *Proceedings of CHI '95*, ACM, pp. 415-420.

[Furnas, 1986] G.W. Furnas, Generalized Fisheye Views: Visualizing Complex Information Spaces. In *Proceedings of CHI '86*, ACM, pp. 16-23.

[Herot *et al.*, 1980] C.F. Herot, R. Carling, M. Friedell and D. Framlich, A Prototype Spatial Data Management System. *Computer Graphics*, 14(3), pp. 63-70.

[Holmquist, 1997] L.E. Holmquist, Focus + Context Visualization with Flip Zooming and the Zoom Browser. In *Proceedings of CHI '97 (Extended Abstracts)*, ACM, pp. 263-264.

[Hornbæk and Frøkjær, 2001] K. Hornbæk and E. Frøkjær, Reading of electronic documents: the usability of linear, fisheye, and overview+detail interfaces. In *Proceedings of CHI 2001*, ACM, pp. 293-300.

[Kadmon and Shlomi, 1978] N. Kadmon and E. Shlomi, A polyfocal projection for statistical surfaces. *The Cartographic Journal*, 15, 1, pp. 36-41.

[Kamba *et al.*, 1996] T. Kamba, S.A. Elson, T. Harpold, T. Stamper and P Sukariya, Using Small Screen space more efficiently. In *Proceedings of CHI '96*, ACM, pp. 383-390.

[Lamping and Rao, 1996] R. Rao and J. Lamping, The Hyperbolic Browser: A Focus+Context Technique Based on Hyperbolic Geometry for Visualizing Large Hierarchies. *Journal of Visual Languages and Computing*, 7, 1 (1996), pp. 33-55.

[Leung and Apperley, 1994] Y. Leung and M. Apperley, A review and taxonomy of distortion-oriented presentation techniques. In *Proceedings of CHI '94*, ACM, pp. 126-160.

[Lieberman, 1994] H. Lieberman, Powers of ten thousand: Navigating in large information spaces. In *Proceedings of UIST '94*, ACM, pp. 15-16.

[Lieberman, 1997] H. Lieberman, A Multiscale, Multilayer, Translucent Virtual Space. In *Proceedings of Information Visualisation '97*, IEEE, pp. 124-131.

[MacEachren, 1995] A.M. MacEachren, *How Maps Work*. The Guilford Press, 1995.

[Mackinlay *et al.*, 1991] J. Mackinlay, G.G. Robertson and S.K. Card, Perspective Wall: detail and context smoothly integrated. In *Proceedings of CHI '91*, ACM, pp. 173-179.

[Matlin, 2001] M. Matlin, *Cognition*, John Wiley & Sons, 2001.

[Rao and Card, 1994] R. Rao and S.K. Card, The table lens: Merging graphical and symbolic representations in an interactive focus+context visualization for tabular information. In *Proceedings of CHI '94*, ACM, pp. 318-322.

[Rauschenbach *et al.*, 2000] U. Rauschenbach, T. Weinkauf and H. Schumann, Interactive Focus and Context Display of Large Raster Images. In *Proceedings of WSCG'2000*.

[Robertson and Mackinlay, 1993] G.G. Robertson and J.D. Mackinlay, The document lens. In *Proceedings of UIST '93*, pp. 101-108.

[Robertson *et al.*, 1991] G.G. Robertson, J.D. Mackinlay and S.K. Card, Cone Trees: Animated 3D Visualizations of Hierarchical Information. In *Proceedings CHI '95*, ACM, pp. 189-194.

[Sarkar and Brown, 1994] M. Sarkar, and M.H. Brown, Graphical fisheye views. *Communications of the ACM*, 37(12), pp. 73-84.

[Siirtola, 2000] H. Siirtola, Scalable Visualization Techniques. Manuscript, March 13, 2000. Available upon request.

[Silvers, 1995] R. Silvers, Livemap: A System for Viewing Multiple Transparent and Time-Varying Planes in Three Dimensional Space. In *Conference Companion of CHI '95*, ACM.

[Spence and Apperley, 1977] R. Spence and M.D. Apperley, The interactive man-computer dialogue in computer-aided electrical circuit design. *Transactions of Circuits and Systems*, CAS-24, 2 (1997), pp. 49-61.

[Spence, 2001] Robert Spence, *Information Visualization*. ACM Press Books, Addison-Wesley, 2001.

[Ware and Lewis, 1995] C. Ware and M. Lewis, The DragMag Image Magnifier. In *Video Program and Companion Proceedings CHI '95*, ACM, pp. 407-408.

[Ware, 2000] Colin Ware, *Information Visualization: Using Perception to Think*. Morgan Kaufmann Publishers, 2000.

[Wittenburg *et al.*, 1998] K. Wittenburg, W. Ali-Ahmad, D. LaLiberte, and T. Lanning, Rapid-Fire Image Previews for Information Navigation. In *Proceedings of AVI'98*, pp. 76-82.