FRUSTUM : A Novel Distortion Oriented Display for Demanding Applications

Paul Anderson, Ray Smith and Zhongwei Zhang

Gippsland School of Computing and Information Technology,
Monash University,
Switchback Road, Churchill Victoria 3842
Australia.

ABSTRACT

Distortion oriented displays (DOD) are an interface approach for supporting navigation through large visual datasets (maps) without losing context. The traditional approaches of windowing and zooming lead the user to lose context within the overall map. DOD present the user with a movable virtual magnifying glass within which a detailed view of the point of focus is presented. Surrounding this the rest of the map is presented in a visually compressed view to ensure context is retained.

An important feature of a DOD is that the user should be able to move the point of focus around the screen and experience no discernible delay in the redisplay of the map. Therefore the computation overhead is very important when considering the implementation of a DOD. This paper describes FRUSTUM, a novel form of DOD with low apparent distortion and minimal computational overhead.

Experience with the FRUSTUM display has indicated that considerably higher magnification factors are possible than the generally accepted maximum for previously described DOD.

Keywords

Distortion-oriented presentation techniques, Information visualisation, Geographical Information Systems.

1. INTRODUCTION

Considerable research has been carried out in the last 10 to 15 years on the navigation problems involved with applications such as Geographical Information Systems (GIS) which present the user with a large visual dataset (map). All GIS present the user with an undistorted but restricted view of the map through a window and allow them to zoom in to obtain a more detailed (and more restricted) view. Unfortunately there is some evidence [1] to suggest that users of this type of interface spend significant amounts of time zooming out in order to re-establish their context to the rest of the map.

An early solution to this problem in the form of a bifocal display was suggested by Apperley and Spence in 1982 [3]. Their solution was to present the user with a movable virtual magnifying glass, within which a detailed view of the point of focus is presented. Surrounding this detail *focus* region the rest of the map is presented in a visually compressed view with superfluous detail filtered out. The image of this *context* view has to be distorted in such a way as to ensure feature matching across the boundary with the region of focus. In this way the user is able to see the detail they require presented within a view that is able to convey its context.

Since the early work on bifocal displays a number of different schemes have been proposed which can collectively be described as distortion oriented displays (DOD). In 1993 Leung and Apperley [2] described a unified theory of these distortion-oriented presentation techniques. Recent experimental work by Leung, Spence and Apperley [5] has confirmed the potential of DOD as presentation and navigational aids for large topological networks.

An important property of DOD is that the user should be able to move the point of focus around the screen and experience no discernible delay in the redisplay. In order to achieve this fluid interaction the display needs to be redrawn at least as fast as the human visual processor cycle time of 100 msec [4] and preferably faster. Unfortunately the amount of computation required for the redisplay of a DOD is considerable. Apperley and Spence resorted to a hardware based solution to this problem. More recent examples have been software based but have been restricted to fairly small datasets in order to achieve the required dynamic response. This is unfortunate since it is applications such as GIS's which are characterised by very large datasets that are most in need of this type of navigation aid.

 0-8194-2030-1/96/$6.00

This paper presents a novel DOD called the FRUSTUM display which has been developed in an attempt to overcome some of these GIS navigation problems.

2. FRUSTUM DISPLAY

The FRUSTUM display is a DOD which involves mapping the dataset onto a chopped off cone with the focal point located on its longitudinal axis. This results in a circular focal region with no distortion but magnified to show the required detail. The context region is rotated onto the surface of the cone as shown in Figure 1. It should be noted that the angle of any point from the focal point is maintained with this transformation.

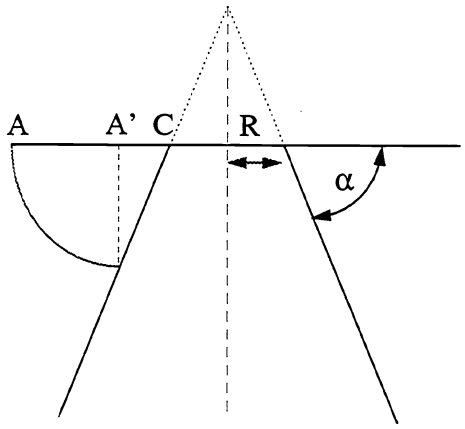

Figure 1. Section showing FRUSTUM mapping onto the surface of a truncated cone.

In Figure 1 the point A in the undistorted view is transformed to point A' by rotating point A about the centre of rotation point C until it strikes the surface of the cone and projecting back to the horizontal plane. The transformation is dependent on both the radius R of the truncated top of the cone and the angle α.

All points that lie within the truncated top (ie, within distance R of the longitudinal axis) map onto the same point therefore ensuring that the focal region presents an undistorted view.

3. DOD MAGNIFICATION FUNCTION

DOD can be categorised in terms of their magnification functions. Firstly those with non-continuous magnification functions such as the bifocal [3] and perspective wall [2] displays. The second group are those with continuous magnification functions which includes the graphical fisheye view [6] and the polyfocal display [2].

The non-continuous functions are generally easier to compute and usually have little or no distortion about the focal point but can have considerable distortion at the boundary with the context region. Whereas the continuous functions produce less apparent distortion and a smoother transition from the focal point but usually have a higher computational overhead.

The FRUSTUM display has a non-continuous magnification function as shown in Figure 2.

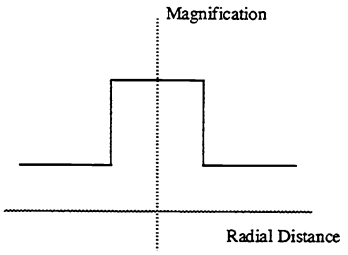

Figure 2. FRUSTUM Magnification function

4. IMPLEMENTATION

The FRUSTUM display has been implemented in WARP (Widely Adaptive and Responsive disPlay) [8] a distortion oriented GIS browser which is based on the Self Tuning ARchitecture (STAR) [7].

4.1 The STAR Architecture

The STAR architecture can functionally be divided into three distinct sections, responsible for data acquisition (uploading), data transformation and display, and performance monitoring and tuning (Figure 3). Changing from one form of DOD to another only affects the data transformation and display module.

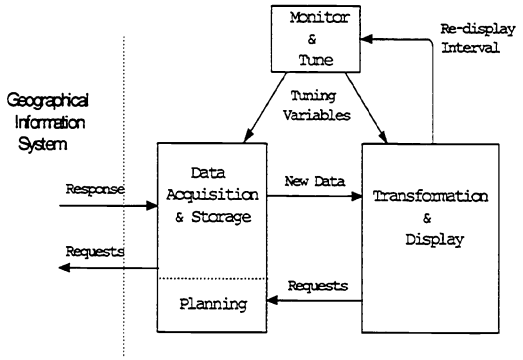

Figure 3: The STAR architecture.

The backbone of the STAR architecture is a combination of static and dynamic filtering regimes which adjust the quantity and/or quality of data displayed based on the measured system performance. The time taken to service requests is monitored and compared with an ideal response time and the quality of the data displayed is modulated in line with the difference. This dynamic adjustment is referred to as *tuning*.

As the context region of a DOD is only there to show context it follows that superfluous detail not only could, but should be removed. Therefore tuning of a DOD interface concentrates on degradation of the context region in preference to the focal region.

The filtering regimes used depend on the type of DOD used, not the STAR architecture. To achieve this STAR only dictates that a single tuning variable (T_V) is used and any particular DOD implementation may map T_V onto any number of tuning variables (eg, one for each filtering regime used).

STAR takes advantage of the normal burst mode of user interaction. Users do not continually move the focus region around but have short intense periods of interaction interleaved with periods of complete inactivity. The rate of moving the focal

region is monitored and if stationary T_V is set to a maximum to achieve a display in its most complete form. However when movement is occurring the value of T_V is modulated to produce the best quality display while still maintaining specified response times. It should be noted that not only is the users activity bursty in nature but also the data can be highly localised. Therefore location of the focal region can also be an input to the tuning process.

4.2 WARP Implementation

WARP is a multiple transformation distortion oriented GIS browser written in C++ and based on the STAR architecture. Currently WARP forms the front end for the INFOCAD GIS. It is a shared memory multi-threaded application with an X window user interface and a shared memory connection to the GIS. WARP has been implemented on a HP9000 715/75 with a graphics accelerator.

A number of contrasting displays have been implemented including :

- Graphical Fisheye (Cartesian) display [6]
- Bifocal display [3]
- Experimental displays (including the FRUSTUM)

4.3 The FRUSTUM Implementation

To allow the FRUSTUM display to be fully investigated the implementation includes the ability to modify the 3 crucial factors governing the appearance of the FRUSTUM display :

- The magnification (distortion) factor
- The radius of the top of the truncated cone
- The angle of the cone surface to the horizontal plane (angle α in Figure 1)

Figure 4 shows the basic interface of WARP with an undistorted view of a map showing Australia and its constituent states and territories. The dataset contains considerable detail for a mythical city near the Queensland/New South Wales border coastline. Details included are road, subdivision, town planning, water and gas service information. The dataset contains in the order of 25000 vectors and is used throughout the rest of this paper to demonstrate the attributes of the FRUSTUM display.

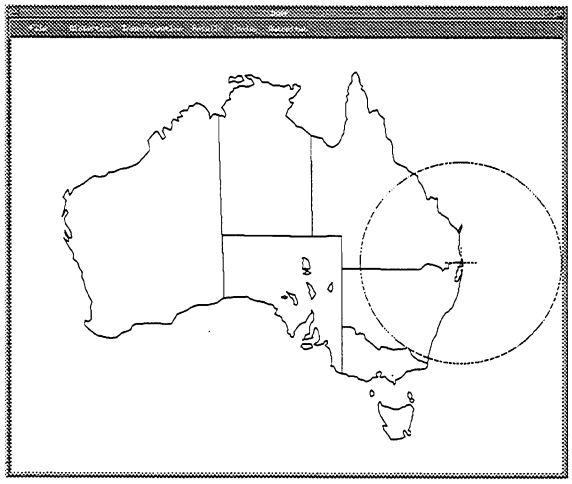

Figure 4. WARP interface with undistorted Australia dataset.

The distortion (magnification) factor is a measure of the size of the focal region for a display compared to its size in the normal undistorted view. WARP allows the user to choose values in the range of 1 (undistorted) to 1000 (very highly distorted).

The normal windowing and zooming used in GIS is shown in Figure 5 which is focused on the mythical city with a distortion factor of 100. The angle α has been set to zero which means no distortion for the context area. The top radius has been specified to be 25% of the window height.

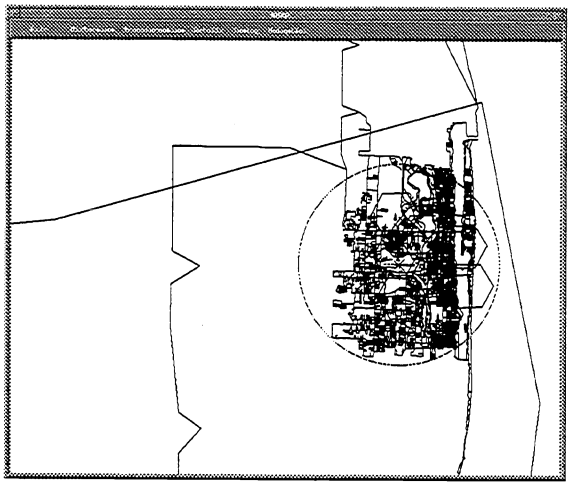

Figure 5. Distortion factor 100, Radius 25% and $\alpha = 0$

Figure 6 demonstrates the dramatic difference that occurs with the FRUSTUM display when the cone surface angle α is increased so as to constrain the whole dataset to appear within the window. The focal point, distortion factor and radius remain the same as for Figure 5.

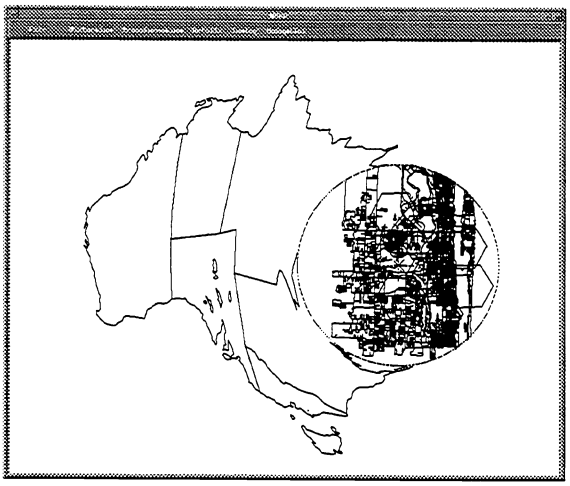

Figure 6. Distortion factor 100, Radius 25% and $\alpha = 89.5$ degrees.

Finally Figure 7 shows the effect of reducing the top radius to 10% of the window height while maintaining all other settings from Figure 6.

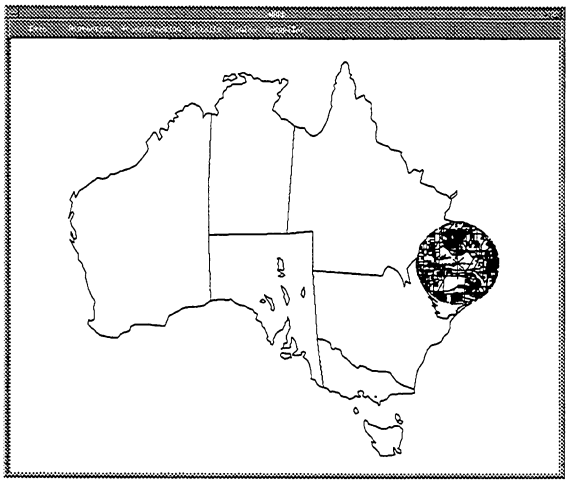

Figure 7. Distortion factor 100, Radius 10% and $\alpha = 89.5$ degrees.

5. PRACTICAL IMPLICATIONS

5.1 Limits of magnification

Experience with the FRUSTUM display has indicated that much higher magnification factors are possible than the generally accepted maximum for other DOD. Other DOD can become unusable above magnification factors in the order of 10 times whereas FRUSTUM produces good displays at magnification of 100 times or greater. The displays are by no means perfect and the search for better transformations will continue but the FRUSTUM is a significant improvement.

5.2 Dynamic response

The large datasets encountered in GIS (typically 25000 vectors or more) make the task of meeting the temporal constraints of a DOD extremely difficult to achieve. The approach used in the STAR architecture to address this problem is to degrade the display quality to reduce display overhead. It would appear obvious that most of the degradation should occur in the context region not the focal region in which the user is vitally interested. When the user stops moving the focal point the display can then be presented at full quality.

Minimising the overhead involved in the recalculation and redisplay of the display is essential. The overhead involves both the more obvious complexity in generating the transformed position and size of each object to be displayed which is DOD dependent and also the overhead in getting the object data from the datastore which is DOD independent.

The FRUSTUM display can be functionally divided into the circular focal region and the surrounding context region. The focal region is always shown undistorted and can therefore be recalculated relatively quickly compared to the context region. Therefore more detail can be included in the focal area in the same unit of time compared to the context area which is as desired. As the focal region is undistorted it is also possible to incorporate techniques using pixmaps stored at the current distortion factor which can be used directly in the correct position.

Vectors within the context region can be transformed into curved lines when mapped onto the cone surface. One approach to achieve this mapping is by fitting a polynomial to the vector data. The greater the number of terms calculated the better the mapping. However, the calculation of more terms incurs a computational overhead. Reducing the number of terms used can produce substantial performance improvements and is incorporated into the tuning process.

Only vector data has been considered in the present implementation but preliminary investigations into the distortion of bitmap data have started.

The current FRUSTUM implementation on the HP9000 715/75 is acheiving the required dynamic response time of 100 msec with a most satisfactory and usable display.

6. DISCUSSION AND FUTURE WORK

The FRUSTUM display is an improvement over other DOD previously described for navigation within the dataset of GIS applications. However it still displays the characteristic distortion at the boundary of the focal and context regions of all transforms with a non-continuous magnification function. A number of modifications to the basic FRUSTUM display are possible to try to overcome this restriction and are currently under investigation.

Further optimisations of the context region are also possible with the most beneficial being the reuse of previously calculated display data. Small movements in the point of focus lead to very small changes in some areas of the context region. Therefore if these areas can be quickly identified it is not necessary to re-transform the data in these areas and the existing display can be reused with no loss of quality. We are investigating efficient ways of identifying these areas for reuse.

7. REFERENCES

1. Tolsby, H. : "Navigating in a Process Landscape". HCI Proceedings of 3rd International Conference EWCHI '93 Moscow Russia. Selected Papers, August 1993 pp 141-151

2. Leung, Y. and Apperley, M. : "A Unified Theory of Distortion-Oriented Presentation Techniques". Massey University School of Mathematical and Information Sciences Report Series A No 93/1.

3. Apperley, M.; Tzavaras, I. and Spence, R. : "A Bifocal Display Technique for Data Presentation". Eurographics '82 pp 27-43., DS Greenaway, EA Warman (eds) North-Holland Publishing Company

4. Card, S.; Moran, T. and Newell, A. : *The Psychology of Human-Computer Interaction.*, Erlbaum Ass., Hillsdale, New Jersey, 1983

5. Leung, Y.; Spence, R. and Apperley, M. : "Applying Bifocal Displays to Topological Maps". International Journal of Human-Computer Interaction 7(1) 1995 pp 79-98.

6. Sarkar, M. and Brown, M. : "Graphical Fisheye Views". Technical Report CS-93-40, Dept. of Computer Science, Brown University September 1993.

7. Anderson, P.; Smith, R.; and Zhang, Z. : "STAR : A general architecture for the support of Distortion Oriented Displays". Proceedings of the KDD'95 conference, Montreal Canada August 1995.

8. Anderson, P. and Smith R. : "WARP : A Distortion Oriented Display Implementation based on the STAR Architecture". Proceedings of the OZCHI'95 Conference, Wollongong Australia. November 1995 pp 278-282