Achieving Higher Magnification in Context

Sheelagh Carpendale[1], John Light[2], Eric Pattison[1]

[1]Department of Computer Science
University of Calgary
Calgary, Alberta, Canada
{sheelagh,ericp}@cpsc.ucalgary.ca

[2]Intel Research
Hillsboro, Oregon
jjlight@ibeam.intel.com

ABSTRACT
The difficulty of accessing information details while preserving context has generated many different focus-in-context techniques. A common limitation of focus-in-context techniques is their ability to work well at high magnification. We present a set of improvements that will make high magnification in context more feasible. We demonstrate new distortion functions that effectively integrate high magnification within its context. Finally, we show how lenses can be used on top of other lenses, effectively multiplying their magnification power in the same manner that a magnifying glass applied on top of another causes multiplicative magnification. The combined effect is to change feasible detail-in-context magnification factors from less than 8 to more than 40.

Categories and Subject Descriptors: I.3.6 [**Methodology and Techniques**]: Interaction Techniques; H.5.2 [**User Interfaces**]: Graphical User Interfaces (GUI)

Additional Keywords and Phrases: magnification, distortion viewing, focus-in-context

INTRODUCTION

We store and manipulate vast amounts of information in a computer and have only comparatively small screens for which it can be viewed. This issue, which affects all aspects of computing, is intensified on the increasingly common smaller screens, and is still an issue on very large screens because their resolution is usually lower, thereby reducing their effective size. Screen space being in short supply makes it difficult to simultaneously present information details and its context. Claims have been made that focus-in-context presentations may better support visual memory and help address problems with search, navigation, and perhaps even reduce cognitive load [6, 17, 23, 24]. This in turn has induced considerable discussion as whether this is important [24] or useful [7, 8, 9, 10, 16, 25]. Since it appears that these types of presentations have their uses [9, 16, 25], this is an appropriate time to consider fine tuning our understanding of the geometry on which they are based.

UIST '04, October 24–27, 2004, Santa Fe, New Mexico, USA.

One consistent limitation in focus-in-context techniques is that they do not effectively provide access to very high magnification in context. This has been discussed as a fundamental limitation of the focus-in-context paradigm [14]. Providing more space for a selected focal region of necessity takes this space from somewhere else. Therefore maintaining context places limits on the amount of magnification. With techniques that maintain context, common magnification factors discussed and illustrated in papers are in the range of 2 to 4 times. One paper has shown an image of 8 times magnification [5] but even this can not be considered high magnification. For serious exploration of dense information spaces these low magnification factors are limiting.

We do not dispute the point that to provide more space for one region in a given presentation, it has to be taken from another region. To address this, we use a visibility based approach through which we can assess and adjust the functions used to make the most of available space. We present several methods that together finesse the manner in which this is done, creating focus-in-context presentations with magnification factors greater than 40.

RELATED WORK

Though there has been little discussion of how to provide high magnification while maintaining context, the seminal papers in this area [23, 6] contain interesting and pertinent discussions about the nature of the context that is to be maintained. Spence and Apperley [23] introduced the idea of maintaining a full but compressed context while viewing selected details. In their terms a compressed *full context* should preserve at least some visual symbolic vestige of all the information. They discuss how searching in physical space is supported by spatial memory, kinetic memory, and visual and verbal clues. Through physical searches all of these clues are continually reinforced by a reasonable degree of constancy. The idea is to support the use of spatial memory and visual scan.

Based on studies that revealed how people use their detailed knowledge set in enough domain knowledge to provide context, Furnas [6] suggested that detail-in-context may be a useful and intuitive way to present information. His studies suggest that, when enough is known about the domain to ensure preservation of crucial aspects, a filtered context may be sufficient. Furnas' context used a degree of interest (DOI) function. A DOI is based upon the distance from the current focus and an 'a priori importance' (API) that is domain specific and known for each item.

While the methods vary, the central purpose, which is to preserve the essential features of the context in less space, remains constant. Different methods of preserving context are: using compression that tries to maintain a full context and preserve spatial constancy [21, 25]; partially filtering the context [6]; combining filtering and distorting [1, 6, 20, 21]; and using distortion only [3, 4, 11, 12, 13, 15, 19].

Sufficient Context Methods

Sufficient context, as Furnas [6] defined it, meant context filtered according to one's degree of interest. What exactly does sufficient context mean for image data? Arguably, separate windows might be considered sufficient context for some image data or some tasks, but with separate windows the user must contend with the cognitive load of image reintegration. Insets may provide some context but, especially with higher magnifications, local context, that region immediately surrounding the inset, is occluded. However, both separate windows and insets do not place limitations on the degree of magnification.

DragMag [26] was the first example of a sufficient context technique designed for image data. Here an inset is moved slightly to one side, thereby revealing local context at the expense of some other, hopefully less relevant, context. In addition, visual cues in the form of connecting lines are included to help the user with the image reintegration. Folding [2, 4] is another sufficient context method very related to DragMag. Here the inset can be repositioned by the user, maintaining the connection between the inset and the context. However, while this provides greater support for cognitive reintegration, it also occludes much more of the context. While both DragMag and Folding provide partial context, the question of whether this is sufficient remains an open issue. Neither of these techniques place limits on the degree of magnification.

Full Context Methods

The idea behind full context is to provide the desired magnification without eliminating any of the context. When the value of full context, such as support for spatial memory and visual search, was first proposed [23], it was achieved by symbolic replacement. More recently this has been achieved through some type of distortion, creating the space for the magnified region by compressing some of the context. Since extremely compressed context becomes unreadable and previously a magnified region required more space according to the degree of magnification, these approaches have had limits for degree of magnification. There was a trade off between space for magnification and the amount of compression that is acceptable. To date published focus-in-context methods commonly present magnifications in the order of 2 or 4 times [3, 4, 11, 12, 13, 15, 19], one achieved slightly higher magnification of 8 [5]. It is this problem that we explore. What are the limits of focus-in-context magnification when one wants to maintain full context? What techniques can extend the limits?

WORKING TOWARDS HIGHER MAGNIFICATION

There are three frameworks for generalizing focus-in-context distortion methods. Leung and Apperley [14] categorized existing methods and unified them by showing how to derive 2D-to-2D transformation algorithms for 3D based techniques. Magnification Fields [12] describes an approximate integration based approach that, given a pattern to strive for, can create any magnification pattern. Elastic Presentation Framework, EPF, [4] demonstrates how to achieve previous 2D to 2D based approaches using a 3D intermediate step and indicated new methods provided by this approach.

We chose to work with the EPF [4] for two reasons. One is that this problem of achieving higher magnification can be framed in terms of visibility and this framework lends itself to a visibility approach. The second reason is that two sufficient context methods, DragMag and Folding, are readily available in this framework and are not in either of the other frameworks. This means that one can plan on extending the degree of magnification as much as possible while maintaining full context and if necessary resort to sufficient context.

The Basics

The following is the basic algorithm for 3D based approaches. A plane or surface that holds the 2D representation is manipulated in 3D then viewed through perspective projection. The transformation function results from the combination of surface manipulation and its perspective projection. This combination simplifies the mathematics of the relationship between magnification and transformation to the geometry of similar triangles. In a perspective framework the 2D surface is placed on an x, y plane that is parallel to the viewplane at a distance along the z axis from the viewpoint which defines unit magnification. Single point perspective projection in this orientation preserves angles, proximity, and parallelism on all planes x, y and has visual realism from the perspective foreshortening in z. The scale or magnification factor of planes parallel to the viewplane is a function of the distance from the viewpoint (for complete explanation see [3, 4, 5]).

The surface manipulation is achieved in the following manner. The focal region of a lens is defined positionally and parametrically so that it provides the desired magnification. Visual integration from the focal region into the context is provided by a drop-off function. Points on the surface are then translated depending on the drop-off function evaluated at the distance from the point in question to the centre of the focus. To ensure full visibility and uniform magnification response the foci are viewer-aligned and the translation vectors are z-normalized (see [4]). The extent of the spread of the distortion into the context can be controlled by adjusting the domain and range of the drop-off function. The manipulated surface is then viewed through perspective projection.

A Visibility Based Approach

Since this framework uses perspective projection and a normal viewing frustum, this problem can be discussed in terms of what is visible from the viewpoint. Figure 1 shows that for a given focus it is relatively simple to find the areas in which a drop-off function would be visible and the areas in which it would not. With this information, choosing a drop-off function and a lens diameter to obtain a visible drop-off function is relatively simple. However, if one increases the magnification, something must be done to maintain visibility (Figure 2). There are three ways we have approached this problem; working with the characteristics of the drop-off functions, working with the diameter of the lens at its base, and working with the size of the focal region.

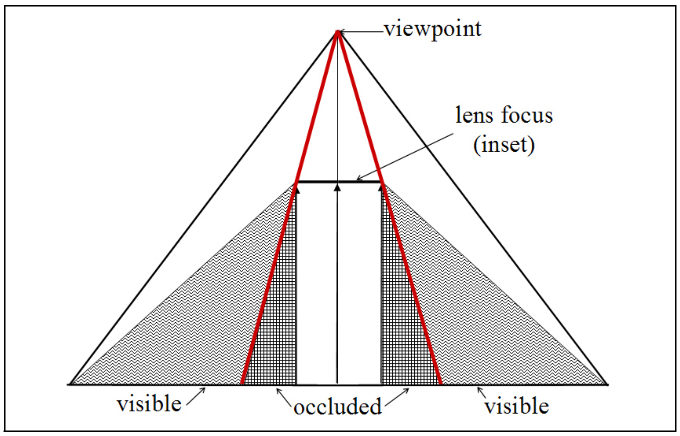

Figure 1: Visible and occluded regions for a given lens focus

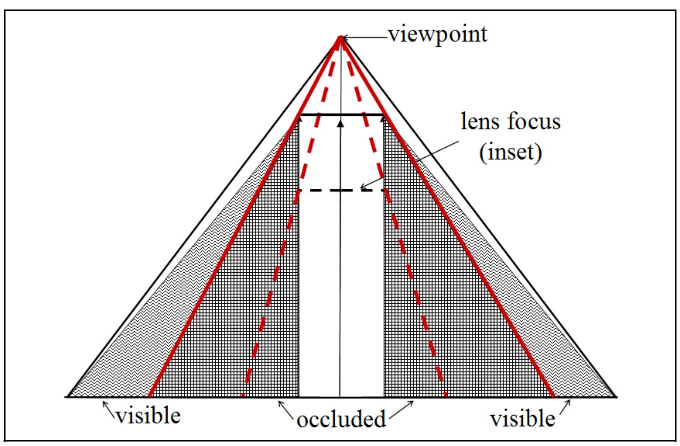

Figure 2: Increasing the magnification creates much large occluded regions

Drop-off Functions

All distortion techniques cause some pattern of expansion and compression, much of which is dependent on the choice of drop-off function. Let us start by looking at a few drop-off functions from a viewpoint visibility perspective. Figure 3 from left to right shows three drop-off functions placed in the *(0,1), (1,0)* square; a hemispherical drop-off, Gaussian drop-off, and a linear drop-off. Considering these drop-offs from a visibility perspective reveals much about the distortion patterns they will create. While the hemisphere will have a smooth connection with the focal area, it will almost immediately cause occlusion at its outer edge, making it a poor candidate for higher magnification. The Gaussian will have a smooth connection with both the focal area and the context, offering good focal and context integration properties and its area of maximum compression will be in the middle of its bell curve. The linear function provides only *C1* continuity at its connections to the focal area and context but does a better job at spreading out the compression across the distorted region. It is also possible that under some situations the clear demarcation of the start and stop of the distorted regions can be beneficial.

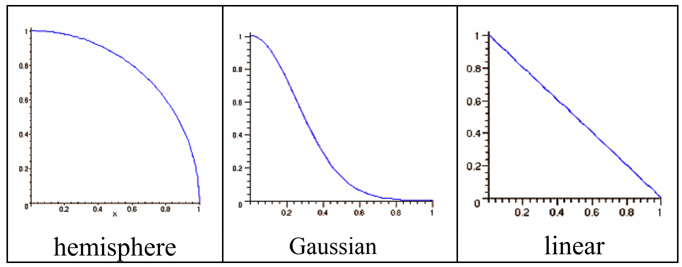

Figure 3: Three drop-off functions

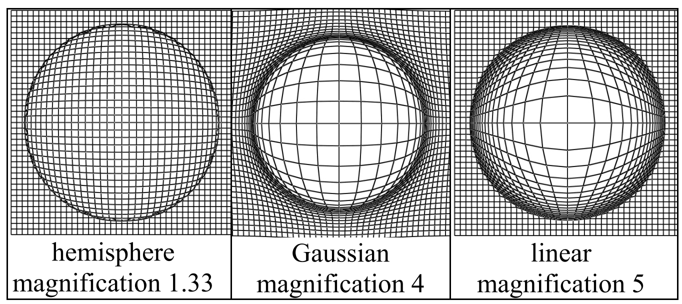

Figure 4: Distortion patterns for 3 drop-off functions

One can see that the choice of drop-off function makes a considerable difference. The hemisphere drop-off achieves a magnification of 1.33 before full context is no longer preserved. Gaussian achieves approximately 4 times magnification and the linear achieves magnification 5 (Figure 4). Both the Gaussian and the linear drop-off functions show promise.

A closer examination of the Gaussian reveals that the central region of the curve is the one that becomes most rapidly compressed as the focal magnification is increased (Figures 5, 6 and 7). When, as in Figure 5, the vectors, *viewpoint to A* and *viewpoint to B*, become coincident, the entire region between *A* and *B* will be projected as a single line (Figure 6a). This maximum compression occurs when the surface normal is orthogonal to the viewvector and causes a ring of maximum compression around the lens (Figure 6b). Continuing to increase the focal magnification will cause the angular order of the viewvectors to *A* and *B* to be reversed and the region that contains point *B* will no longer be visible from the viewpoint.

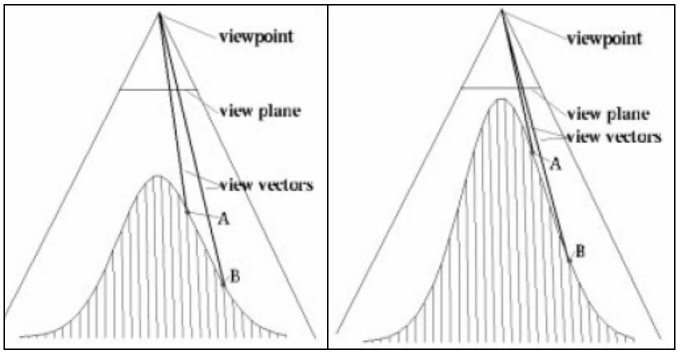

Figure 5: As focal magnification increases the region between *A* and *B* is increasingly compressed

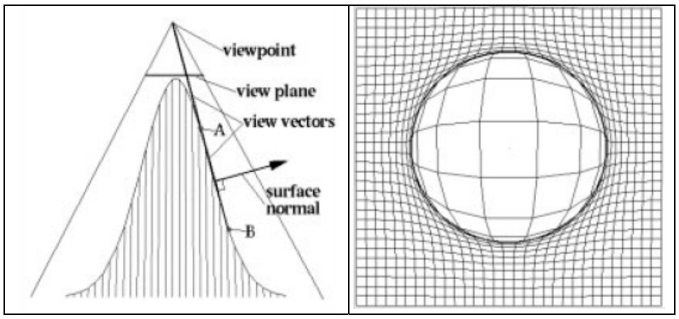

Figure 6: Maximum compression occurs when surface normal is orthogonal to viewvector

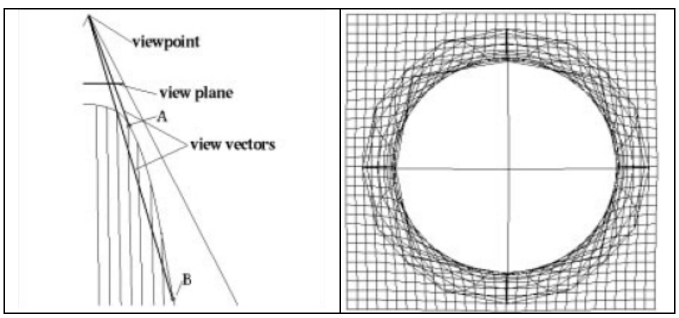

Figure 7: A viewvector passing through surface causes reversal and occlusion

AUXILIARY FUNCTIONS

A possibility noted in [3] is to introduce an auxiliary function, which modifies an existing drop-off function and was tried with the Gaussian drop-off only. For ease of implementation, all drop-off functions are considered from a standardized form decreasing in some manner from *(0,1)* to *(1,0)* (see Figure 3). As a lens radius can vary, we work with a relative distance, *reld* where $reld = dis_{pt} / lr$ and dis_{pt} is the distance from the lens centre to the point in question and *lr* is the lens radius. Similarly one obtains a relative height, *relh*, between 0 and 1, which must be multiplied by the height calculated to provide the selected magnification *hmag*. *hmag* can be calculated from the desired magnification by $hmag = dis_{bp} - (dis_{bp} / mag)$, where dis_{bp} is the distance from the viewpoint to the baseplane. The baseplane is set at the distance to provide a magnification of 1 (see 4).

Auxiliary Function for Gaussian Drop-off

The Gaussian auxiliary function was chosen to be the first two quadrants of the sine function, with varying domain and range to adapt it to a specific purpose. The auxiliary function for the Gaussian drop-off takes the form of:

$$aux_G = 1 - k * sin(reld * \pi/2)^{0.8}$$

This has the effect of reducing the upper "shoulder" of the Gaussian curve, thus, reducing the steep slope in the middle of the curve (Figure 8). Also, it preserves a strength of the Gaussian drop-off function which is that it is highly asymptotic to the *XY* plane.

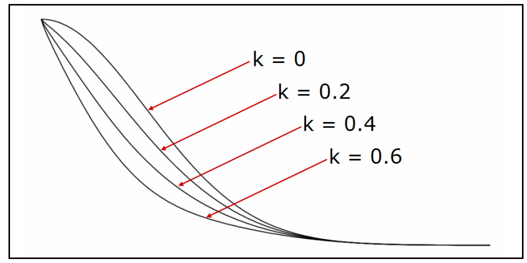

Figure 8: Effects of Gaussian auxiliary function

Auxiliary Function for Linear Drop-off

The auxiliary function to the linear drop-off function takes the form:

$$aux_L = 1 - k * sin(reld * \pi / 1.5),$$

where $0 <= k <= 1$. The *relh* returned by the linear drop-off function is multiplied by aux_L. This has the effect of making the linear drop-off more of a concave curve whose depth varies with *k* (Figure 9).

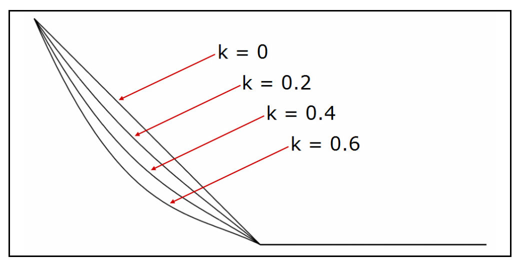

Figure 9: Effects of linear auxiliary function

Note in Figure 9 that the connection to the context is less than ideal and does not improve with the application of this auxiliary function. Both of these auxiliary functions have an adjustable parameter *k* through which the adjustment of the curves created by the drop-off function is possible. The auxiliary function has no effect when $k = 0$. The auxiliary functions do successfully increase the amount of useful full context magnification, but they only work up to a point. The value of k has to be carefully chosen. Figure 10 shows that use of an auxiliary function can give magnification of 16 and 20 with Gaussian and linear lenses respectively. Note Gaussian still has good asymptotic qualities but the linear drop-off has abrupt visual transitions between the lens and its context. That, along with the arbitrary nature of these auxiliary functions, encouraged us to look further.

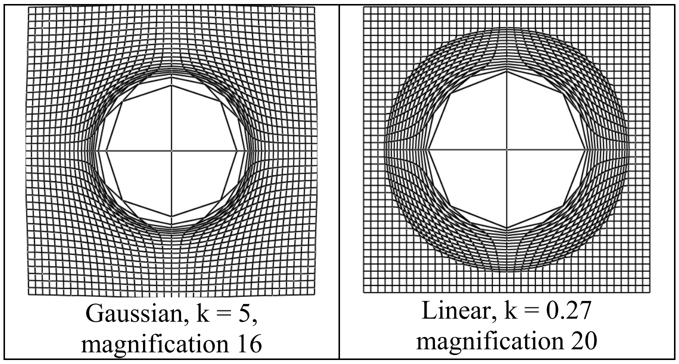

Figure 10: With use of auxiliary functions magnifications of 16 (left) and 20 (right) are possible

CONCAVE FUNCTIONS

Hyperbolic Section Drop-off

Noting the success of the concave function, which results from the application of an auxiliary function, we explored using a hyperbolic section: $xy = k$, as shown in Figure 11.

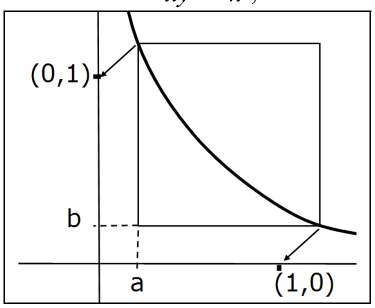

Figure 11: Translating the hyperbolic section

Translation in x is $a = 1/2\left(1 - \sqrt{1 + 4k}\right)$ and translation in y is $b = 1/2\left(1 - \sqrt{1 + 4k}\right)$. The hyperbolic drop-off function is:

$$relh_H = [k/(reld - a)] + b$$

The ability to adjust the curve is incorporated in this function. $k = 1$ gives approximately the curve in Figure 11, $k < 1$ creates more concave curves and as k approaches ∞ this drop-off approaches linear slope (Figure 12). Note in Figure 12 that the connection to the context is not smooth.

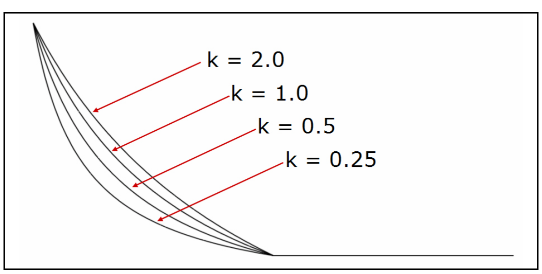

Figure 12: Various hyperbolic drop-off functions

Parabolic Section Drop-off

Another possibility is a parabolic section or Power drop-off function. Translating the section *(-1,1), (0,0)* to the regular *(0,1), (1,0)* (see Figure 13) provides a Power drop-off function:

$$relh_P = (1 - reld)^k$$

For this drop-off, the flexibility comes from k. When $k = 1$ this is the linear drop-off. When $k = 2$ this is a parabola (Figure 14). Setting k to approximately *2.7* provides a visually appealing curve that is capable of displaying much higher magnification in full context (Figure 14).

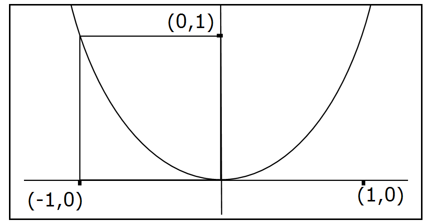

Figure 13: The power drop-off uses the (-1,1), (0,0) square translated to the normal location

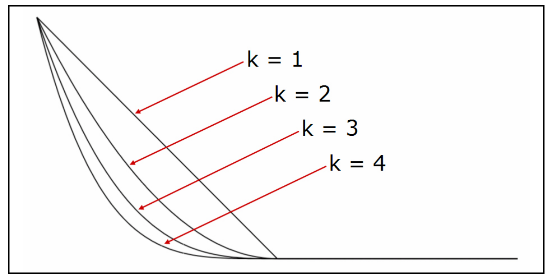

Figure 14: Various Power drop-off function

Figure 15a shows a hyperbolic lens (mag = 20, point focus, k = 0.13) and Figure 15b shows a power lens (mag = 20, point focus, k = 0.7). Both of these drop-off functions provide very high magnification in context but this is for point foci. Figures 15c and 15d show the same two lenses, hyperbolic section and power, with a tiny focal radius. In order to have focal radii of any size, we need to address the lens width factors.

HANDLING WIDTH FACTORS

The images and the explanations in the previous sections all pertained to lenses with a point focus and a reasonable, but fixed lens radius. Even with an effective drop-off function, problems occur due to lens radius and focal radius. This effect is not pronounced at lower magnification, but as magnification increases it becomes a severe problem.

Working with Lens Radius

The lens radius was defined as the sum of the focal radius and the drop-off radius at the base of the lens (Figure 16). Changing the magnification simply raised the focal region and did not affect the lens radius. As a result occlusion occurred fairly rapidly. Figure 17 illustrates this. Here the lens parameters are; drop-off, power function with n = 2.7, focal radius = 20, lens radius = 220. A magnification of 4 (Figure 17a) provides a reasonable lens; a magnification of 6 is at the breaking point when occlusion starts to occur (Figure 17b) and a magnification of 8 has considerable occlusion (Figure 17c).

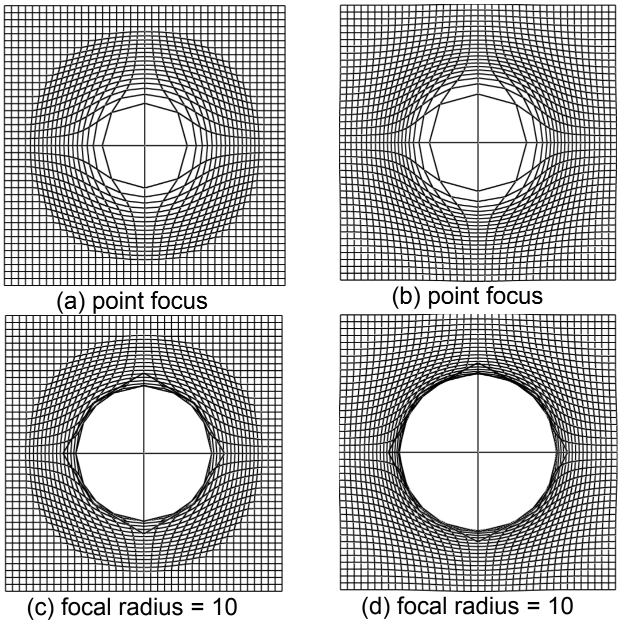

Figure 15: Hyperbolic (left) and power (right) lenses with magnification of 20

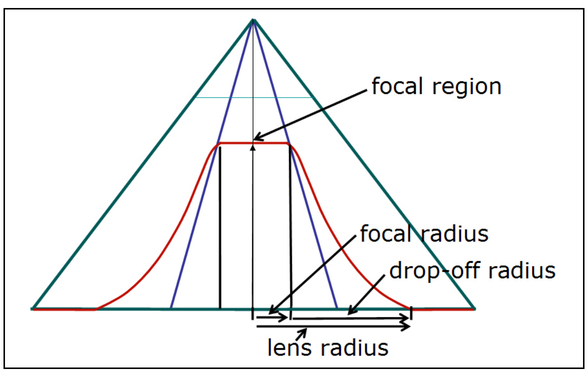

Figure 16: Lens radius

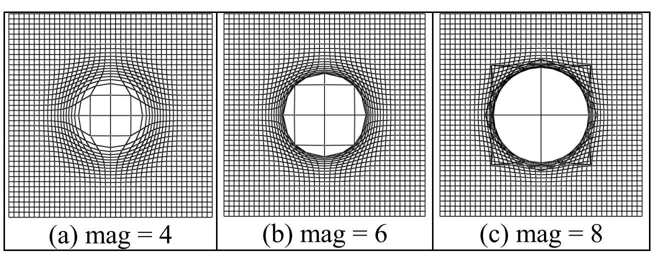

Figure 17: Increasing magnification with a small focal region

The drop-off radius is the distance, at the base of the lens, from the edge of the central focal region to the apparent edge of the distorted region surrounding it. The word apparent is used because with some drop-off functions such as a Gaussian, the function actually continues beyond the point where it has visual effect. As the magnification increases, the height of the focal region above the baseplane increases. In 3D perspective space this means that the focal region appears larger and more of the *XY* plane is occluded. The projected size of the focal region is increasing rapidly. This is a factor that is known and can be used.

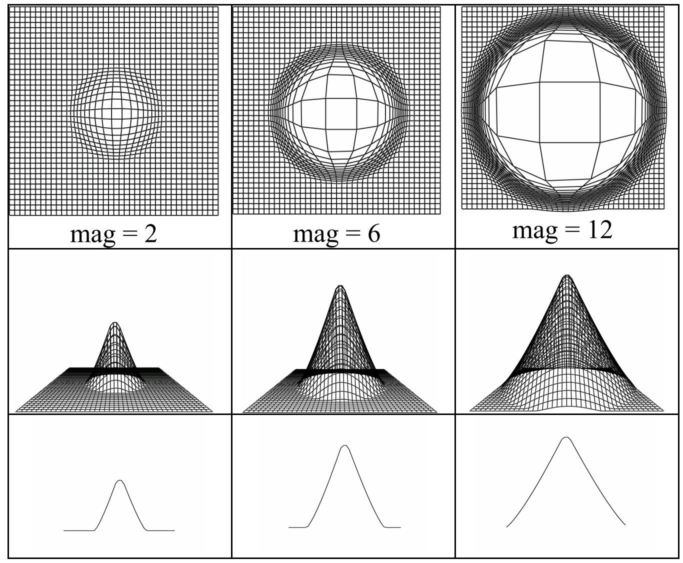

Figure 18: Power drop-off, increasing magnification with expanding lens radius

If the lens radius is defined as the back projected focal radius plus the drop-off radius, then it will expand more appropriately as focal magnification increases. The two definitions for lens radius are the same at magnification of 1 and diverge as focal magnification increases.

$$lensrad = dropoffrad + focalrad * (mag - 1)$$

As Figure 18 illustrates, this offers a considerable improvement but is not the complete answer. Now the break over point, at which occlusion starts, is magnification of 10 instead of 6. The larger the focal region the less effective this is. A similar problem arises from directly expanding the focal region. As is shown in Figure 15, if the focal radius is expanded even with the magnification constant, occlusion will soon occur. Using an expandable definition for the lens radius can improve this (Figure 19). In Figures 18, 19 and 20 the middle and bottom row show the mesh and profile side views of the lenses in the top row.

Working with Focal Radius

In presentations with non-uniform magnification, size is a relative term. For this discussion we will use the terms information-size when referring to the original non-magnified size in information space. For information-size to remain constant, the same information should be present, as magnification increases. However, the apparent size on the screen will increase. In some circumstances this is the desired behavior. For example, if the focus region exactly frames an object of interest, this behavior will keep the same framing as the object is magnified.

Screen-size refers to the amount of space occupied on the display. As seen in the previous discussion the focal region typically expands in terms of screen-size as it is magnified. However, at times it may be preferable to keep the focal screen-size constant while its magnification increases. For example, if you find a small object of interest at low magnification, you might center it in the lens and increase the magnification.

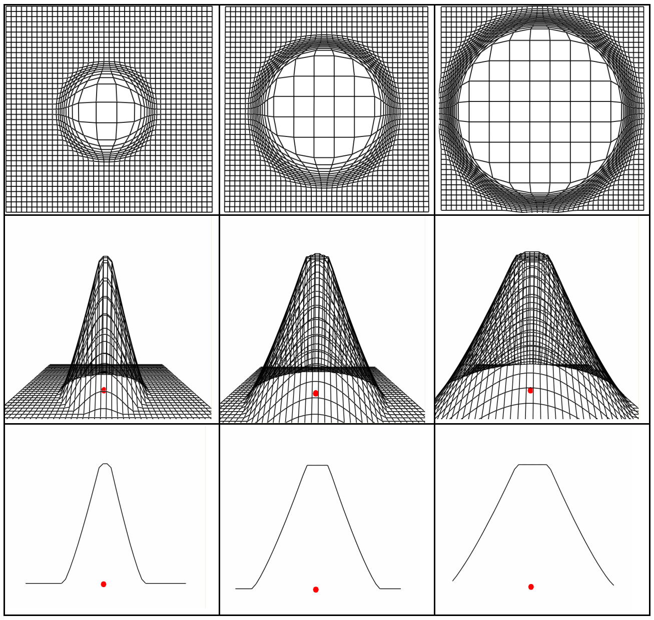

Figure 19: The expanding lens radius also helps when the focal radius is increased

If the lens focus region stays the same screen-size, the object will appear to grow in size until it fills the focus region, while the context remains consistent. To achieve constant focal screen-size, the focal radius can be clamped. To implement this clamping the actual dimensions are scaled to maintain the apparent dimensions in the 3D projection.

ssf = isf [(mag-start)/(mag-current)]

where *ssf* is screen-size of the focus, *ifs* is the information-size of the focus, *mag-start* is the magnification at the time of clamping, and *mag-current* is the magnification just requested. Figure 20 shows in the top row a lens with increasing magnification (*4, 8, and 30x*) and constant focal screen size. The context connection remains effective.

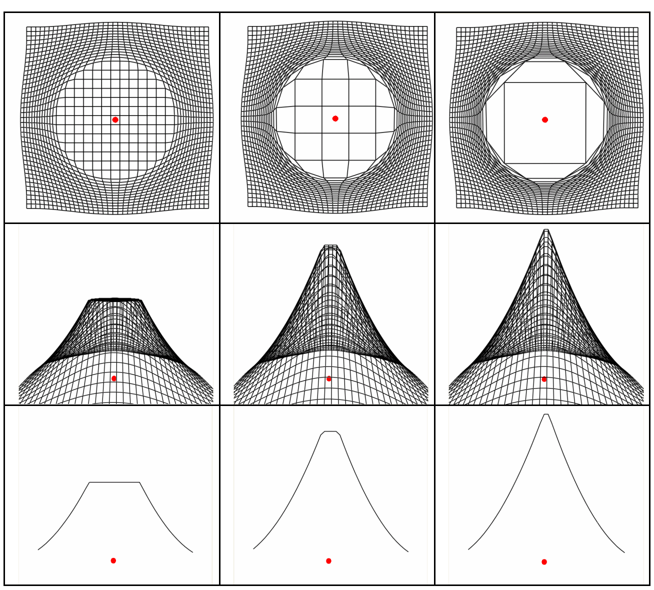

Figure 20: Increasing magnification with constant focal screen size

When the clamp is removed, the lens will behave as it did before it was clamped, increasing in apparent size as magnification increases. Maintaining focal screen-size removes the need for clamping the lens radius to the expanding focal screen-size. The result of maintaining focal screen-size is that lenses appear very stable, as if they are physical constructs in the display space.

HIGH MAGNIFICATION LENSES

Using an effective drop-off function such as the Power drop-off function and setting a constant focal screen size makes quite high focus-in-context magnifications possible. Figure 21 shows a magnification factor of 50. Figure 22 shows an aerial view of Calgary (Calgary image credit [18]), first with no lens, then with lenses of 4, 16, and 27 times magnification.

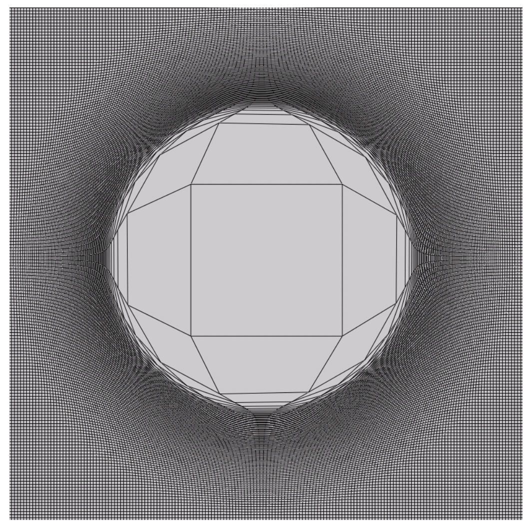

Figure 21: Focus-in-context with magnification 50

Note how in Figure 22d once sufficient magnification is reached instead of more detail becoming available all that is displayed are pixels of increasing size. This can be addressed by bringing in higher resolution information. Figure 23 shows two series. The series on the left shows increasing magnification just resulting in larger pixels, the series on the right shows the same set of lenses in the same location but this time displayed with additional resolution.

LENSES ON TOP OF LENSES

Placing a new lens on top of an existing lens effectively magnifies the results of the lens underneath it. This parallels the effect of magnifying glasses in the real world. Magnifying glass A provides magnification A and magnifying glass B provides magnification B. Applying magnifying glass B over magnifying glass A results in a magnification of BxA. This multiplicative lens effect could be achieved with Keahey and Robertson's [11] compound lens transformation with sequential composition, though they introduce sequential composition as a method for handling spaces in between two lenses where there is partial overlapping making no mention of multiplicative magnification.

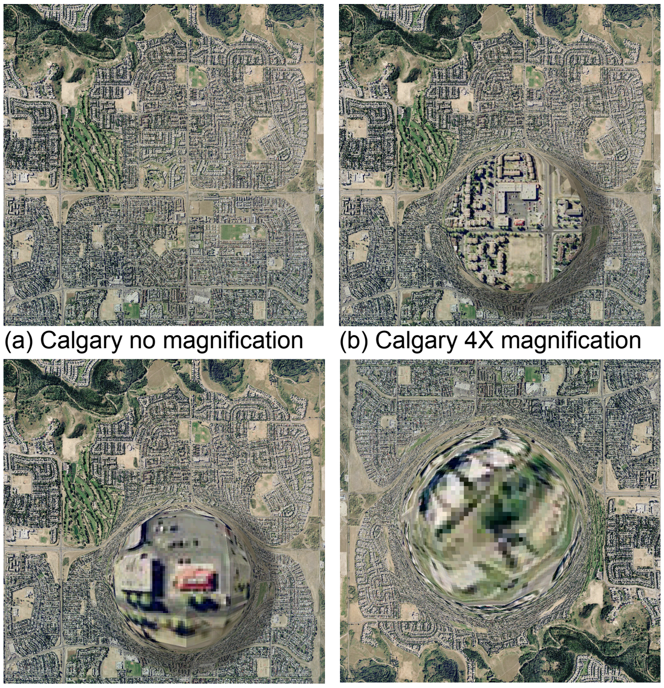

(a) Calgary no magnification (b) Calgary 4X magnification

(c)Calgary, 16X magnification (d)Calgary, 27X magnification

Figure 22: Lenses of increasing magnification (Calgary image credit [18])

A simple method for achieving this is to apply the new lens to the 2D result of the existing lenses (see [11]). However, this solution eliminates the possibility of folding [2, 4] which can be used to clarify the connection of focus to context with high magnification lenses. To enable folding a 3D method is needed. This can be done through use of the relationship between height and magnification in perspective viewing. The magnification *mag* of any particular point in 3D can be obtained through:

$$mag = dis_{bp} / (dis_{bp} - hpt)$$

where dis_{bp} is the distance from the viewpoint to a baseplane and *hpt* is the height above the baseplane of the point of interest. Then the cumulative magnification *Cmag,* the product of the magnification of existing lenses and the new lens, can be used to project the point towards the viewpoint to the new height *newhpt* obtained by:

$$newhpt = dis_{bp} - (dis_{bp} / Cmag)$$

The multiplicative lens is applied after all other lenses have been applied. The result is a lens whose structure appears constant as it moves over the existing lenses.

Figure 24a shows a smaller *4x* lens approaching a larger *2x* lens. In Figure 24b the smaller lens is 'climbing' the side of the larger lens. Instead of blending with the larger lens, this multiplicative lens magnifies portions of the existing lens. For example, Figure 24b shows how the smaller lens can be used to examine the compressed region around the larger lens, locally mitigating the loss of detail caused by compression. Figure 24c shows the smaller lens's magnification multiplying the magnification of the lens underneath creating combined magnification of 8.

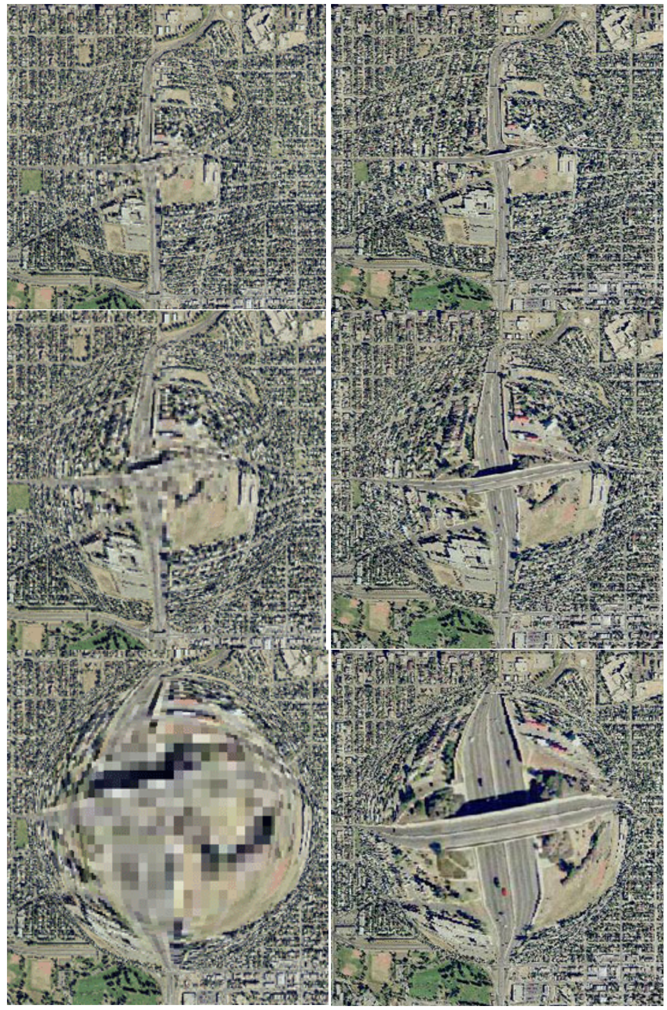

Figure 23: Left top to bottom series, magnification produces fat pixels; right, top to bottom series, introducing more image resolution into the lenses

Figure 25 shows a practical use of multiple lenses. Initially two lenses have been applied to the Calgary image to emphasize two features: a golf course on the left and an east west thoroughfare below it. The smaller multiplicative lens is being used to examine that feature in more detail. Instead of blending with the larger lens, this multiplicative lens magnifies portions of the already emphasized feature, in particular the golf course parking lot and clubhouse buildings at a total 8x magnification.

A multiplicative lens allows for the examination of the drop-off region of another lens. It allows the detailed examination of a part of the focal region of another lens. Such lenses can be stacked because their effects are cumulative. If a given lens is used for relatively stable emphasis, an interactive multiplicative lens can be used for detailed examination. The use of a smaller mobile lens set to a higher magnification allows temporary access to features at very high magnifications without requiring all lenses to be adjusted to higher magnification levels with the attendant high compression. This contributes to using high magnification levels in practical applications.

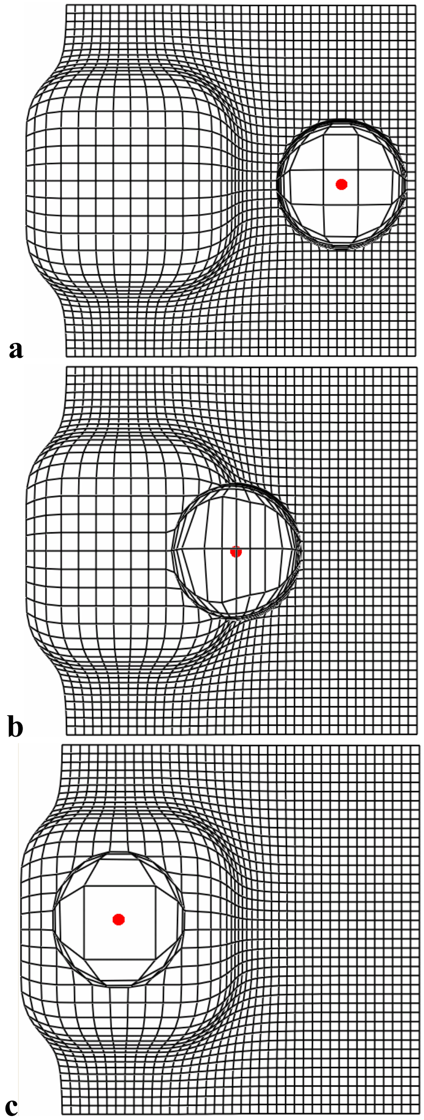

Figure 24: A multiplicative lens being moved across and over an existing lens

DISCUSSION

The creation of high magnification lenses causes no higher burden on processing power than lenses with less magnification. The EPF algorithm complexity is essentially linear with the number of control points that are affected by the lens, invariant with changes in magnification and only somewhat dependent on the type of drop-off function being used. Providing higher resolution data to show in the higher magnification lenses can be a substantial computational burden, though this burden is mitigated by the relatively small area of higher magnification that is typically created.

All of the images (grid and image) in this paper allow the lenses to be moved and reconfigured interactively. While these lenses are algorithmically interactive, precise navigation at extremely high magnifications remains an open problem. Input of greater than pixel level accuracy, such as twips, may be necessary to resolve this.

Figure 25: The small multiplicative lens shows the golf course clubhouse and parking lot (for details about non-circular lens shapes see [4])

Having the capability to achieve higher magnification, it remains to be seen exactly when, where and how much magnification will be useful. The next steps will be to see how these results can be used in practice. This will necessitate addressing navigation issues mentioned above, and issues of information access and readability.

CONCLUSION

While maintaining full context does place limits on the magnification that can be achieved in a focus-in-context presentation, the previously assumed limits of approximately *4x* magnification were much too low. We have shown that full context can be preserved under some circumstances with magnifications as high as *50x*. High magnification is facilitated by judicious choice of drop-off function, and we have shown two new functions that work well for this purpose: hyperbolic and parabolic sections. New methods that provide high levels of detail in the expanded areas created by high magnification will help to overcome the resolution limitations.

It is always the case that there is a trade-off between the factors of focal size, amount of undistorted context and the degree of magnification. Through the use of lenses with expandable lens radii or set focal screen size the interactive interference of these trade-offs can be minimized. Automatic adjustment of drop-off radius and focus screen-size enables the creation of lenses that are able to handle the higher magnification levels. Roaming multiplicative lenses allow the drawbacks of high magnification to be localized in time and screen space.

ACKNOWLEDGMENTS

We would like to thank Intel Inc., the Natural Sciences and Engineering Research Council (NSERC) for support and members of the Interactions Lab, University of Calgary, and anonymous reviewers for insightful comments.

REFERENCES

1. Bartram, L., Ho, A., Dill, J., and Henigman, F. The continuous zoom: A constrained fisheye technique for viewing and navigating large information spaces. In *Proceedings of UIST'95*, ACM, NY, 1995, pp. 207–215.
2. Carpendale, M. S. T., Cowperthwaite, D. J., Fracchia, F. D., and Shermer, T. Graph Folding: Extending detail and context viewing into a tool for subgraph comparisons. In *Symposium on Graph Drawing*, GD'95, Springer-Verlag, 1995, pp. 127-139.
3. Carpendale, M. S. T., Cowperthwaite, D. J., and Fracchia, F. D. 3-Dimensional pliable surfaces: For effective presentation of visual information. In *Proceedings of UIST'95*, ACM, NY, 1995, pp. 217–226.
4. Carpendale, M. S. T. and Montagnese, C. A Framework for Unifying Presentation Space. In *Proceedings of ACM UIST'01*, ACM, NY, 2001, pp. 61-70.
5. Carpendale, M. S. T., Tigges, M., Cowperthwaite, D. J., and Fracchia, F. D. Bringing the Advantages of 3D Distortion Viewing into Focus. In *Proceedings of INFO-VIS'98: Hot Topics,* 1998, pp. 17-20.
6. Furnas, G. Generalized fisheye views. In *Proceedings of ACM CHI'86*, ACM, NY, 1986, pp. 16–23.
7. Gutwin, C. and Fedak, C. Interacting with Big Interfaces on Small Screens: a Comparison of Fisheye, Zoom, and Panning Techniques. In *Proceedings of Graphics Interface*, 2004, pp. 145-152.
8. Gutwin, C. and Fedak, C. A Comparison of Fisheye Lenses for Interactive Layout Tasks. In *Proceedings of Graphics Interface*, 2004, pp. 213-220.
9. Gutwin, C., and Skopik, A., Fisheye Views are Good for Large Steering Tasks. In *Proceedings of ACM CHI'03*, ACM, NY, 2003, pp. 115-123.
10. Hornbaek, K., and Frokjaer, E., Reading of Electronic Documents: The Usability of Linear, Fisheye, and Overview+Detail Interfaces. In *Proceedings of ACM CHI'01*, ACM, NY, 2001, pp. 293-300.
11. Keahey, T. A. and Robertson, E. L. Techniques for non-linear magnification transformations. In *Proceedings of the IEEE Symposium on Information Visualization*, IEEE Computer Society Press, 1996, pp. 38-45.
12. Keahey, T. A. and Robertson, E. L. Nonlinear magnification fields. In *Proceedings of the IEEE Symposium on Information Visualization*, IEEE Computer Society Press, 1997, pp. 51-58.
13. Lamping, J., Rao, R., and Pirolli, P. A focus and context technique based on hyperbolic geometry for visualizing large hierarchies. In *Proceedings of ACM CHI'95*, ACM, NY, 1995, pp. 401–408.
14. Leung, Y. K. and Apperley, M. D. A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer-Human Interaction,* 1(2), 1994, pp. 126–160.
15. Mackinlay, J. D., Robertson, G. G., and Card, S. K. The perspective wall: Detail and context smoothly integrated. In *Proceedings of the ACM CHI'91*, ACM, NY, 1991, pp. 173–180.
16. McGuffin, M. and Balakrishnan, R. Acquisition of Expanding Targets. In *Proceedings of ACM CHI'02,* ACM, NY, 2002, pp. 57-64.
17. Misue, K., Eades, P., Lai, W., and Sugiyama, K. Layout adjustment and the mental map. *Journal of Visual Languages and Computing*, 6(2), 1995, pp. 183–210.
18. Orthoshop (image: Corporate author) Title: City of Calgary 2001 [electronic resource]: digital colour orthophoto, 0.5 m. pixel, on 1/4 township basis Publication info: Calgary : Orthoshop, 2002
19. Robertson, G. and Mackinlay, J. D. The document lens. In *Proceedings of ACM UIST'93*, 1993, pp. 101–108.
20. Sarkar, M., Brown, M. H. Graphical fisheye views. *Communications of the ACM*, 37(12), 1994, pp. 73–84.
21. Sarkar, M., Snibbe, S., Tversky, O. J., and Reiss, S. P. Stretching the rubber sheet: A metaphor for viewing large layouts on small screens. In *Proceedings of ACM UIST'93*, ACM, NY, 1993, pp. 81–91.
22. Schaffer, D., Zuo, Z., Greenberg, S., Bartram, L., Dill, J., Dubs, S., and Roseman, M. Navigating hierarchically clustered networks through fisheye and full-zoom methods. *ACM Transactions on Computer-Human Interaction*, 3(2), 1996, pp. 162–188.
23. Spence, R., and Apperley, M. D. Data Base navigation: An office environment for the professional, *Behaviour and Information Technology*, 1(1), 1986, pp. 43-54.
24. Storey, M-A.D., Wong, K., and Müller, H. A. How Do Program Understanding Tools Affect How Programmers Understand Programs? In *Science of Computer Programming,* 36(2-3), 2000, pp. 183-207.
25. Kuederle, O., Inkpen, K.M., Atkins, S., and Carpendale, M.S.T. Interacting with Image Sequences: Detail-in-Context and Thumbnails. In *Proceedings of Graphics Interface*. 2001, pp. 111-118.
26. Ware, C. Lewis, M. The DragMag image magnifier. In *Proceedings of ACM CHI'95; Conference companion*, 1995, pp. 407 – 408.