User Performance With Gaze Contingent Multiresolutional Displays

Lester C. Loschky
Department of Psychology &
Beckman Institute
University of Illinois at Urbana-Champaign
405 N Mathews Ave.
Urbana, IL 61801
1 217 333 0970
loschky@uiuc.edu

George W. McConkie
Department of Educational Psychology &
Beckman Institute
University of Illinois at Urbana-Champaign
405 N Mathews Ave.
Urbana, IL 61801
1 217 333 7634
gmcconk@uiuc.edu

ABSTRACT

One way to economize on bandwidth in single-user head-mounted displays is to put high-resolution information only where the user is currently looking. This paper summarizes results from a series of 6 studies investigating spatial, resolutional, and temporal parameters affecting perception and performance in such eye-contingent multi-resolutional displays. Based on the results of these studies, suggestions are made for the design of eye-contingent multi-resolutional displays.

CR Categories and Subject Descriptors:

H.5.2 [Information Interfaces and Presentation]: User Interfaces – User-centered Design, Windowing Systems, Screen Design, Evaluation/Methodology; I.4.2 [Image Processing and Computer Vision]: Compression (Coding); General Terms: Human Factors, Experimentation.

Additional Keywords:

multiresolutional displays, dual-resolution displays, eyetracking, eye movements, visual perception, peripheral vision, wavelets, bandwidth

1. INTRODUCTION

Since the introduction of computer technology, electronic displays have become almost ubiquitous. Initially these displays were small with coarse resolution, but both the overall size and spatial resolution have increased dramatically. However, there is always a desire for larger and higher resolution displays, capable of presenting larger and more visually complex CAD models, representing larger terrains and more personnel units for crisis management and military operations, providing more detailed medical images, producing larger and finer grained art work, or giving more realistic immersive experiences. Such large, high-resolution displays require greater computing power and bandwidth to support them, often exceeding realistic limits of what is likely or possible to have available.

Eye Tracking Research & Applications Symposium 2000
Palm Beach Gardens, FL, USA

In this regard, it is important to observe that most of the resources that are used to produce a large, high-resolution display are actually wasted, especially with single-user displays. This is based on perhaps the most widely studied aspect of human vision: that is, the fact that the human visual system can only resolve detailed information within a very small area at the center of vision, with resolution dropping off rapidly as one moves into the visual periphery (e.g., [2, 12, 15, 21]). The biological basis of this spatial variation in human visual resolution has become increasingly well known over the past 20 years, and appears to be due to the spatial density of cone receptors in the fovea and of retinal ganglion cells in the periphery (for a concise review, see [17]; for detail see [4, 7, 8]). Importantly, such differences in resolution across the visual field are well described in terms of a loss of contrast sensitivity to higher spatial frequencies as one moves away from the center of vision [4, 15, 21]. Knowing this, one can therefore save image bandwidth by producing images in which higher spatial frequency information is only encoded at the center of vision, leaving only lower frequency information in the periphery. To do this, however, requires that one know at any given moment where in the image the viewer is looking so that high frequency information from that location in the image can be made available to the viewer. This means that the image must be dynamically updated based on real-time gaze direction information. We refer to this as a Gaze Contingent Multiresolutional Display (GCMRD), or "moving window". The simplest type of multiresolutional display, of course, is to have only two regions with different resolutions; more complex displays could reduce the resolution in steps at more visually eccentric locations, or smoothly degrade the image from the center of vision outward, attempting to match the spatial resolution function of the retina.

GCMRDs can reduce bandwidth requirements in any single-user resource-limited application (for a review, see [14]). These include indirect vision devices, teleoperation of remotely piloted vehicles, teleconferencing and telemedicine. The earliest and best known application of GCMRDs is in flight simulators. Studies have evaluated the characteristics of GCMRDs needed for effective flight simulations [9], such as the field of view requirements for different types of flight tasks to be trained [6, 18]. However, due to limits on the time available in flight simulators for research purposes, most such studies have been unable to carry out careful psychophysical investigations of the perceptual effects of various GCMRD parameters. Thus, more indepth research on GCMRD parameters has generally been done within the context of the psychophysical laboratory.

To our knowledge, only a small number of studies have looked at the perceptual effects of GCMRDs. Two studies have found a functional tradeoff between size of the high-resolution window at the center of vision and the level of peripheral degradation [16, 23]. Specifically, the more degraded the periphery is, the larger the window of high resolution must be to achieve equivalent performance, whether measured in terms of recognition memory [16] or visual search times [23]. However, Shioiri and Ikeda [16] suggest that this relationship does not simply reflect the fall-off of acuity with foveal eccentricity. These authors compared the "useful resolution" function derived from their moving window condition with an "available resolution" function based on Landolt ring acuity at various foveal eccentricities. They found that "useful resolution was substantially less than available resolution" [16, p. 358]; that is, measures of peripheral acuity appear to over-predict viewers' ability to use information from the periphery.

Two other studies investigated the effects of varying the type of information available outside the high-resolution window [20, 22]. Given that people's ability to resolve high frequency information is known to decrease with foveal eccentricity, one would expect that low frequency information in the periphery would be of most value. However, contrary to these expectations, Wampers et al. [22] found that the most disruptive form of peripheral degradation is that in which there is only low-frequency information—a periphery with only high-frequency information was less disruptive to both search times and eye movement measures. The authors have argued that low-frequency information may be sufficient to identify a scene, but that high-frequency information may be more important for identifying individual objects [20].

The above studies also suggest that degrading the information outside the high-resolution window can alter picture-viewing patterns measured in terms of eye movements. Having a moving window can reduce saccade lengths [22], with smaller windows leading to shorter saccades [16], similar to results found with moving window studies using text displays [11]. This reduction in average saccade length suggests that the process of selecting potential targets for eye movements is being affected. Thus, degrading those parts of a scene for which the visual system needs high resolution may disrupt the process of locating new saccade targets and integrating information across eye movements. It has also been found that changing the image in the periphery during eye fixations, even when this involves replacing a degraded image with one that is at full high resolution disrupts processing and results in longer eye fixations [20].

The above-cited studies leave open a number of questions. Two key issues yet unaddresssed are the timing of GCMRD updates and detectability of the GCMRD. For example, how soon after the end of an eye movement does the window need to be updated in order to avoid disrupting processing? Also, is there a difference between the window sizes and peripheral degradation levels that are visually detectable and those that produce behavioral effects? Other unresolved issues are tied to methodological details of the above studies. Two of these studies used very simple stimuli [16, 23], thus raising questions about their generalizability to situations with realistic images. Furthermore, one of the studies used head rather than eye tracking, thus making its window size effects difficult to compare with eye tracking results [23]. Finally, the other two studies, while they used more realistic images and involved eye tracking, did not independently manipulate window size and peripheral degradation level [20, 22], thus making it impossible to determine the interaction of these factors.

The current paper summarizes the results of a series of six studies that examined some of the psychological issues involved in using GCMRDs with realistic photographic scenes. The studies both provide information about characteristics of the human visual system, and information needed by engineers who wish to produce human-friendly displays of this type. These studies were conducted with a specially designed system that includes a high-accuracy eyetracker, and a large image memory that allows the use of precomputed images for fast updating of the Gaze Contingent Display in order to investigate how the dynamic control of the image affects the visual system.

2. METHOD

2.1 Participants

In all the studies described below, participants were college students or adult members of the community with normal uncorrected vision. The number of participants in each study is given in parentheses (e.g., n = 20).

2.2 Stimuli

All experiments used a set of 15 complex, monochromatic, photographic scenes as stimuli (see Figure 1). Images were transformed such that a circular, high-resolution region (the "window") was surrounded by a degraded peripheral region through the use of the discrete wavelet transform. These transformations were accomplished by, first, carrying out a wavelet decomposition of each image into 4 bands of increasingly higher spatial frequencies, each an octave apart, using the 9/7 symmetric biorthogonal wavelet basis function described by Antonini, et al. [1, p. 208, filter tap values given on p. 209]. This resulted in a total of 13 subbands, each having an associated set of coefficients. We then reconstructed each image while dropping out varying numbers of higher frequency subbands. Using all 13 subbands essentially recreated the original image; including only some of them (10, 7, 4, or 1) resulted in a degraded image (Figure 2)(see [13] for a detailed description of the image filtering). An important benefit of using the discrete wavelet transform to create multiresolutional images is that it allows one to selectively include all of the spatial frequency subbands in the region where high resolution is desired, and only a subset outside that region.

Using the wavelet transform is an efficient method of saving image bandwidth; discarding sets of wavelet coefficients results in lossy image compression. The simplest way of expressing the degree of compression is in terms of measures based on the proportion of coefficients remaining, *p*. Thus, 1-*p* is the proportion of coefficients saved, and

$$BRS = B - ((p*B)/B)$$

is a measure of bit rate savings, where B is the bit rate [13. p. 33]. Niu used this measure to arrive at rough estimates of the coefficient savings for a 512 x 512 image, finding savings ranging from nearly 50% to 94% for the various combinations of degradation level (1, 4, 7, and 10 subbands included) and window size (2°-5°) used in the studies reported below (see Table 1). Note, however, that increasing the size of the image *outside of* the gaze contingent high-resolution window further increases savings. In fact, only one of the experiments below used the 512x512 image (subtending $14^{\circ 2}$ visual angle) assumed by Niu. All other experiments used 768x512 images (subtending 18° x 12° visual angle), for which savings would be greater since the degraded proportion of the image would be higher.

Table 1: Rough Estimates of Proportion of Coefficients Saved for a 512x512 image (from [13, p. 34])

Radius of High-Res Window	Degradation Level			
	10	7	4	1
2°	0.706	0.882	0.927	0.938
3°	0.651	0.813	0.854	0.864
4°	0.574	0.717	0.753	0.762
5°	0.475	0.593	0.613	0.631

Figure 1. Example Stimulus Picture with Full High-resolution Used in Control Condition. Picture Shows a Guatemalan Street Market Scene.

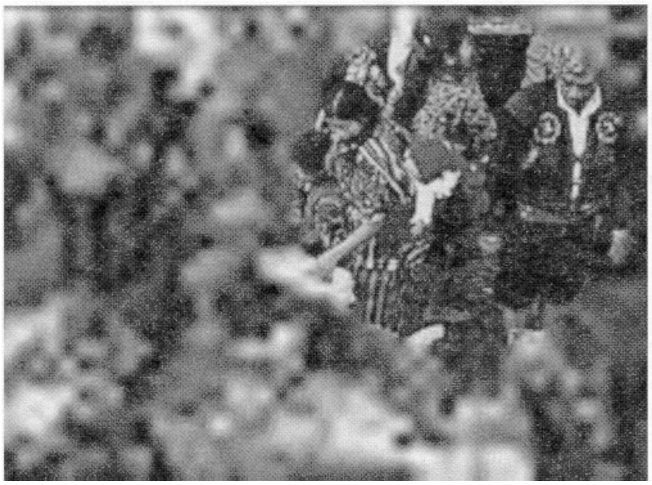

Figure 2. Same Picture as Shown in Figure 1, but with Largest High-resolution Window (4.1° radius) and Highest Peripheral Degradation (Level 1).

All studies included normal, high-resolution versions of the images as a control condition in order to determine baseline behavior and/or false alarm rates. Sizes of the high-resolution windows are reported as the radius of the circle in degrees visual angle. In all studies, different subjects were presented with the same pictures under different conditions, so the results are counterbalanced across pictures.

2.3 Apparatus

Experiments involving eye tracking used a Dual Purkinje Image generation V eye tracker, sampling eye position 1,000 times/sec, and with a spatial resolution of 1/16° visual angle, using a bitebar to stabilize the participant's head. The experiments were controlled by a 486 microcomputer that was interfaced with the eyetracker. An online algorithm used the time and position values produced by the eyetracker to determine whether the eye was in a saccade or a fixation on any given 1 ms sample. This involved checking if the difference between the current x and y positions and those 4 ms earlier were above a threshold of 3.6 minutes of arc. Three difference measures are used: change in x, change in y, and change in x and y. If any threshold was exceeded, the algorithm then checked the next 4 samples to determine if they were also above threshold and moving in the same direction. If so, it determined that a saccade had begun with the first above threshold sample. The algorithm for determining the end of a saccade differed in that all three difference measures had to be below a threshold of approximately 1.5 minutes of arc/4 ms. Fixations were defined as the intervals between consecutive saccades. The location of a fixation was determined by taking the average location during the fixation period.

Experiments using the Gaze Contingent Display (or "moving window") utilized a ViewGraphics 2-Gbyte semiconductor image memory in which was stored 330 precomputed versions of each picture, each having a high resolution area at a different location, corresponding to a 15x22 array of screen locations, each less than 1° apart. A 60 Hz refresh rate was used. The change from one image to another could begin at any point during the refresh cycle of an image; therefore, it was not necessary to wait until the end of a refresh cycle before a new version was displayed. Thus, a new image was completely scanned onto the monitor within a real 17 ms of the time an image change was requested.

3. EXPERIMENTS

In a gaze contingent multi-resolutional display ("moving window"), the image on the screen can be changed either continuously (updated on each frame based on current eye position) or only at the end of every saccadic eye movement, keeping the image stable during eye fixations even when the actual position of the eyes drift somewhat. In this research we opted for the second approach, updating the image only at the end of each saccade, in order to reduce the likelihood of actual perception of the image changes taking place. Image changes made during a saccadic eye movement go unnoticed because the visual system has a temporarily higher perceptual threshold, or "saccadic suppression" (e.g., [5, 10]). However, if the change takes place very long after the eyes have stopped moving, it will likely be noticed. As a precursor to the moving window studies, we performed a study to determine the earliest time following an eye movement that a change to the display, of the type planned, could be detected.

In this experiment (n= 12), a high-resolution version of a scene was initially displayed on the screen. However, at the beginning of selected saccades, this image was replace by a highly degraded version (created from 4 sets of coefficients) of the same picture. We then changed the degraded version back to the high-resolution version after a variable number of milliseconds (ms). Participants were asked to press a button if they detected a flicker or blurred image. The results showed that for an image change to go undetected, it must be started within 5 ms after the end of the eye movement. Detection likelihood rose quickly beyond that point. In this study the monitor was refreshed at a rate of 144 frames per sec, thus requiring 7 ms to change an image. This is an important result for our moving window studies, because it gives us a deadline for updating the image.

Our next question concerned the perceptibility of peripheral degradation in a multi-resolutional display. Specifically, how large a window is needed to avoid detection of different levels of peripheral degradation? Two studies were conducted to investigate the detection of peripheral image degradation with a stationary window briefly flashed at the center of vision for 150 ms (enough time to perceive the image but not to make an eye movement). Since the participant's only task was to push a button if he/she detected peripheral degradation, we considered this to be an ideal condition for maximizing the likelihood of detecting the window. In the first such study (n = 10), the high-resolution window sizes ranged from 2°-5° radius, and there were 4 levels of peripheral degradation (using 1, 4, 7, or 10 sets of wavelet coefficients). Figure 3 shows that the least peripheral degradation (level 10) went undetected even at the smallest window size (2°), whereas the opposite was true with the highest level of degradation (level 1)—it was quite detectable at even the largest window size (5°). At intermediate levels of degradation, window size made a difference, such that the larger windows led to less detection. Another study of the same type (n = 20) used the same levels of peripheral degradation, with a slightly different range of window sizes (1.6-4.1°) and added two window edge sharpness conditions (sharp edge vs. smoothed edge). We replicated the results of the previous study but found that edge type had no effect on detectability. Similar to the previous results, under ideal conditions for detecting peripheral degradation, only minimal degradation (level 10) was completely undetectable; however, moderate peripheral degradation (level 7) was seldom detected at the largest window size (4.1°).

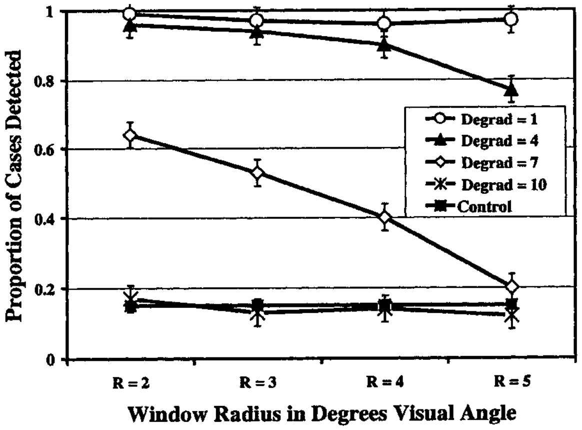

Figure 3. Detection of Peripheral Degradation in Briefly Flashed Images (Stimulus Duration = 150 ms) as a Function of Window Size and Degradation Level. Error Bars Represent Standard Error. Control Condition was a Full High-Resolution Image.

In our fourth study we asked how performance and eye movement behavior in natural tasks would be affected by use of a gaze contingent, multi-resolutional display. In this experiment (n = 15), within 5 ms of the end of each eye movement the image on the screen was updated, selecting from among 330 versions of the image to center the high-resolution window on the viewer's fovea.

Participants viewed a picture for varying times up to 20 sec. The window sizes and degradation levels were a subset of those used in the detection studies described above (windows = 1.6°, 2.9°, 4.1° radius; degradation levels = 1, 4, 7). Participants' tasks were to search for an object in the scene, or to remember the scene.

As shown in Figure 4, search times were significantly slower in the smallest window condition (1.6°) than in the largest window condition (4.1°), which did not differ from the control baseline. In contrast to the detection results, however, search times were less affected by the level of peripheral degradation. Time spent in each eye fixation (fixation durations) produced a similar pattern of results. On the other hand, Figure 5 indicates that both window size and peripheral degradation level significantly affected the mean distance traveled by the eyes in saccadic movements (saccade length). Specifically, saccade lengths were shorter in the smallest window condition than in the largest window condition, which did not differ from the control baseline. The smaller the window and the greater the peripheral degradation, the greater the tendency for the eyes to be sent to locations within the high-resolution window. Overall, a window size of roughly 4° radius with moderate peripheral degradation (level 7) resulted in eye movements and search performance that were quite comparable to the situation in which there was no window at all.

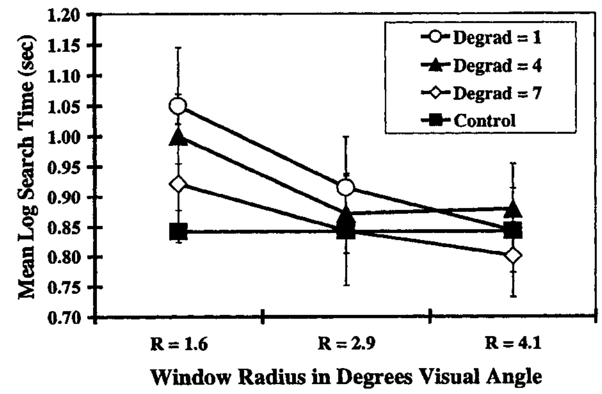

Figure 4. Mean Log Search Time (Seconds) in a Gaze Contingent Multi-resolutional Display as a Function of Window Size and Degradation Level. Error Bars Represent Standard Error. Control Condition was a Full High-Resolution Image.

The fifth study investigated whether detection of peripheral degradation is reduced when the viewer is engaged in realistic, ongoing tasks such as visual search, as compared to just trying to detect the degradation. In the study (n = 30), participants were engaged in the same tasks as in the moving window experiment (i.e., search and memory), but the window of high resolution with peripheral degradation appeared only occasionally on single fixations, with a full high-resolution image present at all other times.

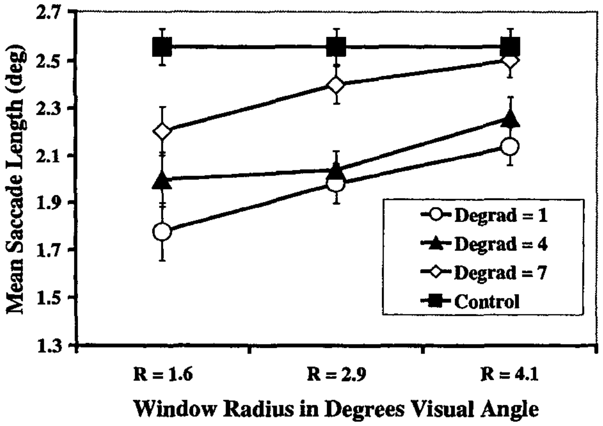

Figure 5. Mean Saccade Length in a Gaze Contingent Multi-resolutional Display as a Function of Window Size and Degradation Level. Error Bars Represent Standard Error. Control Condition was a Full High-Resolution Image.

Participants were given a secondary task, which was to press a button any time they noticed peripheral degradation. Window sizes and degradation levels were the same as in the moving window study. The results were surprising in that they were virtually identical to those of the previous detection studies using tachistoscopic presentation. Note that in those studies, participants' only task was to detect degradation, whereas in this study their primary task was presumably much more demanding—search for an object or conduct a detailed inspection of the image in preparation for a memory test. As shown in Figure 6, having such additional tasks did not lessen the detectability of peripheral degradation. Thus we conclude that a moving window need not affect eye movements or performance in realistic tasks to be detectable, nor does engaging in such tasks narrow the attentional span for peripheral degradation. As with the prior detection study, a roughly 4° radius window with moderate degradation (level 7) was detected less than half of the time.

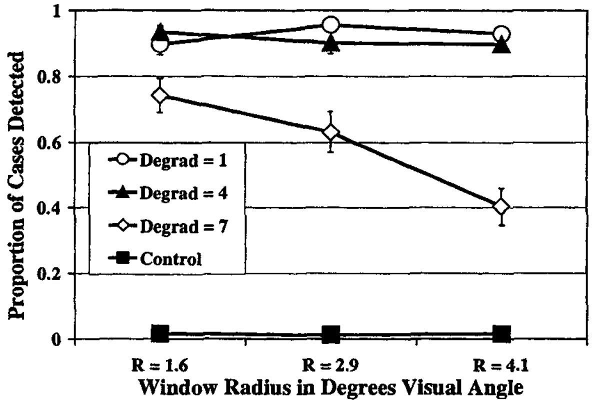

Figure 6. Detection of Peripheral Degradation (Presented on Occasional Fixations in a Gaze Contingent Multi-resolutional Display) as a Function of Window Size and Degradation Level. Error Bars Represent Standard Error. Control Condition was a Full High-Resolution Image.

Our sixth and final study investigated how quickly the gaze contingent multi-resolutional display must be updated at the end of each eye movement in order to maintain natural task performance and eye movements.

In previous experiments we had used the deadline of 5 ms following the end of a saccade for updating the image. However, we were only able to do so because of two technologies unlikely to be available in a head-mounted, gaze contingent, multi-resolutional display: (1) an eyetracker with high temporal resolution (sampling every ms), and (2) precomputed images stored in a very large image memory (see Apparatus section). Assuming that a head-mounted system will be slower than the system we have used, we need to know how slow it can be without adversely affecting performance and eye movements. The study (n = 30) used the same tasks and measures as the moving window study, but with two degradation levels (1 and 7), two window sizes (1.6° and 4.1°) and three image update deadlines (5, 15, and 45 ms following the end of a saccade). The search time results showed no effect of updating deadline, though there were strong effects for both window size and peripheral degradation level. As in the moving window experiment, a window size of 1.6° produced longer search times than a window size of 4.1° radius, which was very similar to the control condition. However, contrary to our previous results, degradation level 1 also produced longer search times than level 7, which was also similar to the control. We attribute this difference in results to the greater power of our current study because we included twice as many participants as the previous study.

Fixation durations provided a more sensitive measure of visual processing than overall search times. The results indicated that a 45 ms delay in updating a gaze contingent multi-resolutional display may be too long, especially with the window sizes used in our study. As shown in Figure 7, viewers' average fixation times were reliably longer for 45 ms delays than for either 5 or 15 ms delays, which did not differ from each other. The results for window size and degradation level were consistent with the earlier moving window study—window size produced an effect, with the 4.1° window yielding fixation durations nearly equivalent to the control condition, but level of peripheral degradation did not make a reliable difference. Finally, it is worth noting the apparent trade-off in fixation duration between update delay and high-resolution window size. We see in Figure 7 that the mean fixation duration for the smallest window with the shortest delay (R = 1.6°, Delay = 5 ms) is approximately equal to that of the largest window with the longest delay (R = 4.1°, Delay = 45 ms), though both differ from that of the control condition. Thus, the longer the update delay is, the less benefit one gets from having a larger sized window of high resolution, i.e., it becomes perceptually similar to being in a small window condition. This is most likely due to an increasing probability of perceiving peripheral degradation after the eyes have moved, but before the window is updated to a new position, as the update delay increases.

Overall, the results of this study indicate that while a delay of as much as 45 ms in updating a multi-resolutional display may not appreciably affect gross measures of performance on visual tasks (e.g., total time taken to find an object in a photograph), it nevertheless does affect visual processing (e.g., fixation durations). However, delays on the order of 15 ms seem to affect neither type of measure. We do not know from this study at what point in the range of 15-45 ms these effects on processing begin. However, we would recommend that updating delays be kept to less than 45 ms if possible in gaze contingent multi-resolutional displays. Furthermore, the apparent trade-off between update delay and window size suggests that if one must use a relatively long update delay in a GCMRD, one may need a relatively larger window size to make up for it.

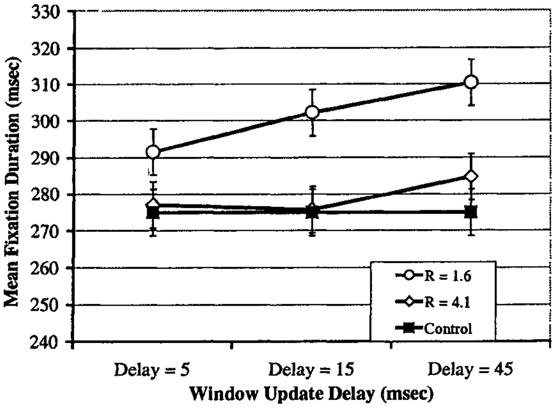

Figure 7. Mean Fixation Durations as a Function of Window Update Delay and Window Size. Error Bars Represent 95% Confidence Intervals Based on Within-subject Error. Control Condition was a Full High-Resolution Image.

4. CONCLUSIONS

There are two possible goals of a Gaze Contingent Display system: to produce a display that is indistinguishable from a full-resolution display, or to produce a display that, even though distinguishable from a full-resolution image, none-the-less does not deteriorate visual task performance. We found that the first goal is quite difficult to achieve: the display must be updated within 5 ms of the time an eye fixation begins, and the peripheral degradation (that lies outside the high-resolution area) must be minimal. However, while greater delays (e.g., 15 ms) and greater degradation (including just the first 7 levels of coefficients from the wavelet decomposition of the image in the periphery) produce detectable visual artifacts, in fact they appear to have minimal impact on the performance of visual tasks when there is a 4.1° radius, high-resolution area centered at the point to which the gaze is directed. This illustrates a more general point, which is that the effects of a particular window size and degradation level will vary depending on the tasks and measures one uses. This should be borne in mind whenever evaluating GCMRDs.

An important general principle that is illustrated by the findings in these experiments is that various factors in a GCMRD can be traded off against one another. For example, consistent with the results of previous psychophysical studies of GCMRDs [16, 24] we found that as the level of peripheral degradation increased, the high resolution window size also had to increase in order to maintain the level of performance exhibited in the control condition. This reflects the general finding that the sensitivity of the visual system varies as a function of the interaction between spatial frequency and retinal eccentricity. We also found a trade-off between update delay and window size, such that longer delays require larger window sizes in order to maintain control level performance. This seems to reflect both the temporal dynamics of visual sensitivity following a saccade (i.e., saccadic suppression)[5, 10], together with the more general relationship between visual sensitivity and retinal eccentricity. More generally, then, we find that the size of the window of high resolution in a dual resolution GCMRD can be traded off against various other factors that have an impact on performance.

Whether the edge between the high and lower resolution areas of the image is sharp or softened made no difference in the detectability of the visual artifacts, so is a variable that apparently need not be of concern. This result is in conflict, however, with those of previous GCMRD flight simulator studies [3, 19] that indicated that the blending region was important to the perceptual quality of the display. Thus, it is also possible that the lack of difference found in our study was due to both the 'sharp' and 'soft' edges being perceptually salient.

A surprising result is that detection of peripheral degradation is just as great when the participant is searching freely for objects in a picture, as when the image is presented briefly to participants in a peripheral degradation detection study. We had expected that under dynamic viewing conditions, participants might frequently constrict their region of visual attention and often miss the presence of peripheral degradation. That does not appear to be the case.

The fact that near-normal performance on tasks can be achieved with a high-resolution area of only 4.1° radius, and with substantial loss of detail in the periphery, indicates that the Gaze Contingent Display approach has potential for saving considerable computation and bandwidth in single-user displays, while still providing the user with apparently high-resolution images. Nevertheless, the current studies were conducted under conditions in which the viewer had a limited range of movement and in which the images in the multi-resolutional display were prestored. Thus, future studies are needed to extend these findings through the use of eyetracking systems that allow greater freedom of movement (though these tend to have with lower spatial and temporal resolution) and on-line image computation algorithms. Our prediction at this point is that the lower spatial and temporal resolution of such systems will require a larger window size to achieve natural performance and eye movements, with low detectability.

5. ACKNOWLEDGMENTS

Research funding for this project was provided in part by the US Army Federated Laboratory under Cooperative Agreement DAAL01-96-2-0003. Some of the equipment used was provided under NSF Infrastructure Grant No. CDA 96-24396. The gaze contingent multiresolutional image display programs for this study were written by Gary Wolverton. The program to produce dual resolution images was written by Ed Niu.

6. REFERENCES

[1] Antonini, M., Barlaud, M., Mathieu, P. and Daubechies, I. Image coding using wavelet transform. *IEEE Transactions on Image Processing, 1.* 205-220.

[2] Aubert, H. and Forster Beitrage zur Kenntnisse der indirecten Sehens. *Graefes Archiv fur Ophthalmologie, 3.* 1-37.

[3] Baldwin, D., Area of interest: Instantaneous field of view vision model. in *Image Generation/Display Conference II*, (, 1981), 481-496.

[4] Banks, M.S., Sekuler, A.B. and Anderson, S.J. Peripheral spatial vision: Limits imposed by optics, photoreceptors, and receptor pooling. *Journal of the Optical Society of America, 8* (11). 1775-1787.

[5] Burr, D.C., Morrone, M.C. and Ross, J. Selective suppression of the magnocellular visual pathway during saccadic eye movements. *Nature, 371* (6497). Oct 1994, 511-513.

[6] Collyer, S.C., Ricard, G.L., Anderson, M., Westra, D.P. and Perry, R.A. Field of view requirements for carrier landing

training (AFHRL-TR-80-10), Air Force Human Resources Laboratory, Williams Air Force Base, AZ, 1980.

[7] Curcio, C.A. and Allen, K.A. Topography of ganglion cells in human retina. *Journal of Comparitive Neurology* (300). 5-25.

[8] Curcio, C.A., Sloan, K.R., Packer, O., Hendrickson, A.E. and al., e. Distribution of cones in human and monkey retina: Individual variability and radial asymmetry. *Science, 236* (4801). 579-582.

[9] Lund, R.M., Technology: An operational requirements case study - The Harrier GR MK5/7 Mission Simulator Visual Display System. in *Flight Simulation Technology, capabilities and benefits*, (London, United Kingdom, 1995), Royal Aeronautical Society, 20.1-20.8.

[10] Matin, E. Saccadic suppression: A review and an analysis. *Psychological Bulletin, 81* (12). Dec 1974, 899-917.

[11] McConkie, G.W. and Rayner, K. The span of the effective stimulus during a fixation in reading. *Perception & Psychophysics, 17* (6). 578-586.

[12] Munn, N.L. and Geil, G. A note on peripheral form discrimination. *Journal of General Psychology, 5*. 78-88.

[13] Niu, E.L.-C. Gaze-based video compression using wavelets *Electrical and Computer Engineering*, University of Illinois at Urbana-Champaign, Urbana, IL, 1995, 76.

[14] Reingold, E.M., Stampe, D.M., Loschky, L.C. and McConkie, G.W. Variable-resolution gaze-contingent display applications: An integrative review. *Manuscript submitted for publication, University of Toronto*.

[15] Robson, J.G. and Graham, N. Probability summation and regional variation in contrast sensitivity across the visual field. *Vision Research, 21* (3). 1981, 409-418.

[16] Shioiri, S. and Ikeda, M. Useful resolution for picture perception as a function of eccentricity. *Perception, 18*. 347-361.

[17] Thibos, L.N. Acuity Perimetry and the Sampling Theory of Visual Resolution. *Optometry & Vision Science, 75* (6). 399-406.

[18] Thomas, M. and Geltmacher, H. Combat simulator display development. *Information Display, 9*. 23-26.

[19] Turner, J.A., Evaluation of an eye-slaved area-of-interest display for tactical combat simulation. in *The 6th Interservice/Industry Training Equipment Conference and Exhibition*, (, 1984), 75-86.

[20] van Diepen, P.M.J. and Wampers, M. Scene exploration with Fourier-filtered peripheral information. *Perception, 27* (10). 1998, 1141-1151.

[21] Virsu, V., Naesaenen, R. and Osmoviita, K. Cortical magnification and peripheral vision. *Journal of the Optical Society of America, 4* (8). Aug 1987, 1568-1578.

[22] Wampers, M., Diepen, P.M.J.V. and d'Ydewalle, G. The use of coarse and fine peripheral information during scene perception, Laboratory of Experimental Psychology, University of Leuven, Belgium, 1998.

[23] Watson, B., Walker, N. and Hodges, L. Effectiveness of peripheral level of detail degradation when used with head-mounted displays, Georgia Institute of Technology, Graphics, Visualization & Usability Center., 1995.

[24] Watson, B.A., Walker, N., Hodges, L.F. and Worden, A. Managing level of detail through peripheral degradation: Effects on search performance with a head-mounted display. *ACM Transactions on Computer-Human Interaction, 4* (4). 323-346.