

The Effects of Low Frame Rate on a Measure for User Performance in Virtual Environments

Martin Reddy*

Technical Report ECS-CSG-36-97
Department of Computer Science
University of Edinburgh

1 December 1997

Abstract: This report investigates the effect of low frame rate on human performance in a desktop virtual environment. This is done using an objective measure of users' ability to perform a simple heading task at various frame rates (2.3–14.2 Hz). Two principal experiments are presented. The first experiment shows that for a drop in frame rate from 11.5 Hz to 2.3 Hz, users' accuracy and time to complete the task degraded significantly. The second experiment reveals a continuous, asymptotic relationship between frame rate and performance for the chosen task. At low frame rates (up to 10-15 Hz) there is a sharp improvement in performance as frame rate increases. After around 15 Hz this increase is substantially less rapid. The results provide evidence for reinforcing that a minimum frame rate of around 15 Hz is necessary for virtual environments, but also that further increases in frame rate will continue to yield greater performance levels, albeit at a much reduced rate.

Keywords: frame rate, transport delay, user performance, simulated pursuit fixation, virtual environments.

1 Introduction

The latency that can be experienced within a virtual reality (VR) system can be categorised into two principal components. These are: frame rate and transport delay (Bryson, 1993). Frame rate (or update rate) is the rate at which the visual display is refreshed with a new image, whereas transport delay (or lag) is the period of time between a user's input and the effect of that input being represented on the display device (Bryson and Fisher, 1990),

*The author has since changed affiliation. Any enquiries should be sent to: Martin Reddy, SRI International, Artificial Intelligence Center, 333 Ravenswood Avenue, Menlo Park, CA 94025-3493.

e.g. the delay between a user panning their head and the display being refreshed with the new perspective of the virtual environment (VE).

A degradation of either of these two quantities can exhibit a detrimental effect on the ability of the user to interact with the system in an efficient manner. This can be explained, at least in part, by the temporal aliasing and perceptual distortions caused by high latency systems which can deteriorate a user's depth and egomotion perception (Piantanida *et al.*, 1993). As a result, it is generally agreed that frame rate is more important than display fidelity in VEs (e.g. Smets and Overbeeke, 1995; Swartz *et al.*, 1992).

For example, taking the case of transport delay, lags of between 30–120 ms have been shown to degrade user performance, depending upon the application (Held and Durlach, 1993). Singhal and Cheriton (1995) state that humans can detect network lags of around 100 ms in a distributed VR system, and will only tolerate maximum inconsistencies in the order of 200 ms. Gregory (1990) also reports that a lag of around 500 ms can seriously degrade hand-eye coordination tasks such as drawing and writing.

Frame rate degradation can also affect user performance in a VE (Lampton *et al.*, 1994). (In fact, Bryson, 1993, reports that the effect of both frame rate and transport delay on the user's ability to perform a particular tracking task was quantitatively similar). For example, Tharp *et al.* (1992) found that performance degraded below 10 Hz for their immersive tracking experiments, and that performance levels reached a plateau between 10–20 Hz. Also, Watson *et al.* (1997) report that moving from a frame rate of 20 Hz down to 10 Hz significantly degraded user performance, both in terms of accuracy and response time.

2 Goal of Current Study

In this report we are concerned solely with the effect of frame rate on the ability of users to perform a prescribed task in a VE: transport delay will not be investigated (i.e. it will be assumed to be constant). As a result, a desktop system was favoured—rather than an immersive one—in order to minimise the influence of transport delay. We present two principal experiments which encapsulate the goal of this work. These attempt to achieve the following objectives:

Experiment One : given a specific task, confirm the basic premise that a user's performance will degrade when they must work in a low frame rate environment (e.g. < 10 Hz).

Experiment Two : once this premise has been validated, investigate the actual relationship between the frame rate of the simulation and the performance of users with respect to the particular task chosen.

The remainder of this report will be organised as follows. First we will describe the task that will be used to objectively assess a user’s performance. We will then proceed to describe the above two experiments in turn; presenting the stimuli and procedures employed, the results obtained, and a discussion of these results in each case. Finally we will present our conclusions and assess the findings of this study.

3 The Performance Task

The measure that was elected to assess user performance was a heading task based upon the simulated pursuit fixation work of Cutting *et al.* (1992) and others from the field of visual perception (e.g. Warren and Hannon, 1990; Cutting, 1986; Rieger and Toet, 1985). This task involves the subject being passively transported through an environment of objects. The display is updated so that the observer is always looking towards a certain fixation point which is deviant from their heading direction. To illustrate this task in terms of a real world scenario, imagine that you are on the back of a jeep manning a TV camera which is free to pan left and right. The jeep is being driven through an environment cluttered with various objects, e.g. trees. You look through the camera’s viewfinder and track one specific tree as the jeep moves through the environment, keeping the tree always in the centre of the viewfinder. This concept is illustrated in Figure 1.




Figure 1: *Overview of the heading task in which the subject is passively navigated through an environment with their direction of fixation oriented differently from their direction of movement. Angle δ represents the gaze/movement angle at one point on the navigated course.*

The task of the subject is to deduce whether they are being navigated to the left or to the right of the fixation point (based upon the radial motion cues which they acquire from the surrounding objects). For example, in Figure 1 the heading direction is to the left of the target. Intuitively, this task will become more difficult as the differential between

the fixation and the heading vectors (referred to as the gaze/movement angle) decreases. Typically, the gaze/movement angle will increase as one progresses along the navigated path. The maximum value of this angle for any trial is the independent variable of interest, i.e. the final gaze/movement angle.

From a number of subjects' responses to different navigation scenarios, Cutting *et al.* produced psychometric curves which plot the final gaze/movement angle against the percentage of correct responses. The point at which this curve drops below a certain threshold (e.g. 95% correct) can be used to compare the subject's task performance under different situations.

The frame rate of the simulation was varied for different trials in order to investigate the variance of users' performance with this quantity. N.B. the simulated forward velocity remained constant for all frame rates, i.e. the maximum virtual distance travelled was always the same, but at higher frame rates the progression appeared more smooth. (It should be noted that all values for frame rate, Hz, presented in this report are averaged over one entire trial and are assumed to be constant throughout that trial. In reality, the potential variance of frame rate during a single trial was of the order ± 0.1 Hz.)

The format of this experiment has obvious parallels with various generic VR applications, such as driving and flight simulators, where users must navigate a course through a VE. It is also worth noting that this type of heading experiment, as well as having been used by numerous vision scientists in the past, has also been used by researchers in the field of VR to assess the performance of subjects within a VE (e.g. Wann *et al.*, 1995). The experiment is of course completely passive: the user does not manipulate the navigated course in any way. The resultant data therefore represents findings for the basic perceptual affordances that the system provides to enable the task of a user, rather than assessing the ability of a user to interact with the system using any one specific interaction metaphor.

4 Experiment One

4.1 Method

Purpose. To investigate the performance of users (accuracy and time) to complete the above heading task for two test frame rates. The chosen frame rates represent an excessively low value (2.3 Hz) and a moderately acceptable one (11.5 Hz).

Stimuli. The test VE was generated using the IRIS Performer graphics library on a Silicon Graphics, Inc. (SGI) Onyx RealityEngine² workstation with one 200 MHz R4400 micro-processor. The content and dimensions of the environment were modeled in order to replicate closely Cutting *et al.*'s experimental setup. Specifically, each environment

contained 150 objects, randomly positioned at ground level (across the x - z plane), and randomly rotated (about the y axis). The fixation object was coloured light purple and a crosshair was positioned over it to guide the user's fixation.

Subjects were navigated through the environment for up to 5 seconds, but could submit their decision at any point. The fixated object was initially positioned at a distance of 50 m from the viewpoint. Each object was 4.32 m high and the simulated forward velocity was 4.36 m/sec. The object which was used to populate the environment was the 'temple' model supplied with the SGI machine. An overview and example screenshot of the test environment are provided in Figure 2.




Figure 2: *An overview of the simulated pursuit fixation task. (a) presents an example screen shot of the test environment, and (b) provides a plan overview of the environment layout. The fan shape in (b) illustrates the various paths through the environment which were traversed. The fixation object can be noted at the centre of the 50m line.*

Images were displayed at a resolution of 1280×1024 pixels. The subject viewed the screen such that it occupied 43.6×33.4 degrees of their field of view (and these FOV values were used by the graphics renderer for all perspective calculations). Steps were taken to ensure that no other users could remotely log into the workstation during the course of an experiment. This was done in order to limit the effect of any background processes which might interfere with the update rate of the experiment.

Sixty-four trials were randomly presented to each subject. These were composed of 8 final gaze/movement angles (0.25, 0.5, 1, 2, 4, 8, 16, and 32 degrees) \times 2 movement directions (left and right of the fixated object) \times 2 frame rates (2.3 and 11.5 Hz) \times 2 attempts.

Procedure. Twenty subjects participated in the study, drawn from a breadth of backgrounds including undergraduate students, postgraduate students, staff, and graduates of the University of Edinburgh. All subjects had normal or corrected-to-normal vision and were naïve to the experimental hypothesis. They were encouraged to keep their gaze fixated on the crosshair at the centre of the screen, but no attempt was made to monitor eye movements. A chin rest was used to restrict subjects' head movement and to maintain the viewing distance (which was set to 50 cm).

Subjects pressed either the left or right mouse button to indicate whether they felt that they had been transported to the left or to the right of the fixation point, guessing when unsure (i.e. the experiment was performed as a 2 alternative forced choice). Reaction times were recorded for each trial. No feedback was given to the subject on their success rate during the experiment to reduce the chances of them learning some nonsense visual task unrelated to wayfinding. Subjects were given a number of practice trials beforehand until they were satisfied that they understood the task. The experiment lasted about 20 minutes and participants were paid £5.00.

4.2 Results and Discussion

Figure 3 presents the results of the first simulated pursuit fixation experiment, averaged over the 20 subjects. These are compared for the two test frame rates.

The psychometric curve in Figure 3(a) shows the average proficiency of subjects to correctly ascertain their heading. We can see from this graph that at large gaze/heading angles, the user is very proficient at correctly resolving their heading direction; but this ability drops to chance level ($\sim 50\%$) for smaller angles. It is immediately evident that subjects' efficiency was maintained higher for longer for the higher frame rate case. Using Cutting *et al.*'s 95% threshold as a measure of performance, we find that subjects could discriminate angles which were 2.8 times smaller (9.5 deg / 3.4 deg) for the five-fold difference in frame rate from 2.3 Hz to 11.5 Hz.

Figure 3(b) shows the average response time of subjects during the task, i.e. how long it took them to resolve their heading direction. Again it is obvious from a cursory inspection that subjects had a distinctly faster response during the higher frame rate trials. From these data we find that, on average, users performed 1.67 times faster at 11.5 Hz than at 2.3 Hz.

To summarise, the findings of this task performance experiment show conclusively that a low frame rate can detrimentally affect user performance in a VE. Specifically, for the five-fold drop in frame from 11.5 Hz down to 2.3 Hz, we found that users' accuracy dropped significantly, and that the time it took them to perform each task decreased markedly.




Figure 3: Averaged results from the first simulated pursuit fixation experiment. (a) shows the overall proficiency of subjects to perform the task. (b) presents the overall response times of subjects. In both graphs, the solid line represents a frame rate of 11.5 Hz; whilst the broken line represents a frame rate of 2.3 Hz. The horizontal bars annotating the curves in (a) represent the line of 95% correct.

5 Experiment Two

5.1 Method

Purpose. To augment the results of the previous experiment with two further test frame rates. With these data we can attempt to draw conclusions on the actual relationship between frame rate and user performance (accuracy and time) for the simulated pursuit fixation task.

Stimuli. Exactly the same experimental setup and stimuli were used for this experiment as were used in the previous case, with the exception being that the two test frame rates were chosen to be 6.7 and 14.2 Hz. This provides us with a value intermediate to the previous values, and another marginally higher value. Each subject therefore performed 64 trials, consisting of the same 8 final gaze/movement angles as before \times 2 movement directions \times 2 frame rates (6.7 and 14.2 Hz) \times 2 attempts.

Procedure. In this instance, five subjects performed the heading experiment. Each subject had normal or corrected-to-normal vision and had not participated in the previous experi-

ment. As before, users pressed one of two mouse buttons to state their perceived movement direction. Reaction times were recorded and no feedback was given. Subjects were allowed a number of trial runs to acquaint themselves with the experimental technique.

5.2 Results and Discussion

The averaged results from this second experiment are plotted in the graphs of Figure 4. The results from the previous study are overlaid here to allow a comparative inspection. From these data we can see that the results from the current experiment follow the same general trends revealed by the previous study, i.e. users' accuracy and time to complete the task degraded for the lower frame rate case. More specifically, at a 95% correct threshold, users could perform the task for final gaze/heading angles that were 2.1 times smaller (6.0 deg / 2.8 deg) for the $\times 2.1$ improvement in frame rate from 6.7 Hz to 14.2 Hz. Also, on average, users could respond 1.37 times faster for the higher frame rate trials.




Figure 4: Averaged results from the second simulated pursuit fixation experiment, displayed along with the previous experiment's results. (a) shows the overall proficiency of subjects to perform the task for the four frame rates: 2.3, 6.7, 11.5, and 14.2 Hz. (b) presents the overall response times of subjects for these same frame rates.

We can also note from Figure 4 that the magnitude of user performance in each case appears to be proportional in some way to the frame rate of the simulation. Taking the data for all four frame rates, we are in a position to describe the actual relationship between user performance and frame rate for our performance measure. This is achieved in Figure

5.

Figure 5(a) plots the frame rate of the simulation against the final gaze/heading at which the 95% correct threshold was reached for the task. This therefore conveys the threshold accuracy of users to perform the simulated pursuit fixation task in relation to frame rate. A curve fit process was applied in order to find a simple, general relationship. This is represented by the broken curve in Figure 5(a).




Figure 5: *The relationship between the frame rate of the system and (a) the users' 95% threshold accuracy to perform the task, and (b) the average time it took users to perform the task. The broken curve in (a) defines a good representative curve fit to the data, $17.18x^{-0.67}$ ($r = 0.978$). The broken curve in (b) represents a best curve fit given by the equation, $-4.005 \log_{10}(x) + 8.232$ ($r = 1.000$).*

We can observe from this graph that user accuracy improves precipitously at first as the frame rate is improved, but that this increase soon begins to level off, presumably to some asymptotic threshold where further frame rate improvements do not affect performance. By extrapolating and generalising these results we can state coarsely that user performance accuracy drops significantly below the 10–15 Hz range: whilst above roughly 15 Hz relatively slight improvements in performance are to be gained for higher frame rates. This lends credence to the general maxim that a frame rate of 15 Hz is a minimum requirement for VEs (e.g. Falby *et al.*, 1993; Barfield and Hendrix, 1995).

N.B. whilst a frame rate of around 15 Hz may produce an acceptable level of user performance for many applications, this should not necessarily be taken as a statement that 15 Hz is all you need. It is evident from the curve in Figure 5(a) that further increases in frame rate will marginally improve a user's ability to perform the task. So for an accuracy critical application, such as a flight or surgical simulator, the improvement in performance

to be gained by going from 15 Hz to 30 Hz is likely to be significant. That is, the 15 Hz neighbourhood does not guarantee optimum performance, merely a generally acceptable degree of performance.

In the case of user response time, we can observe from Figure 5(b) that a similar relationship appears to exist here as for user accuracy. That is, we find that the speed with which users can perform the task degrades logarithmically for lower frame rates. We also find that the general statement that we made earlier, that performance degrades most precipitously below roughly 10–15 Hz, appears to hold true for response time as well as accuracy.

6 Conclusions

The product of this report has been the proposal of the simulated pursuit fixation experiment as an objective measure for user performance in a VE, and the usage of this metric to assess the effect of low frame rate on the ability of users to interact with a VR system.

We confirmed the fact that forcing a user to work in a low frame rate VE (e.g. < 10 Hz) will substantially degrade their ability to perform tasks, both in terms of accuracy and speed. We also found that above roughly 15 Hz, the rate of improvement in user performance is much less dramatic. We therefore use these results to support the aphorism that a frame rate of around 15 Hz should be taken as a minimum requirement for VR applications; but we also make the addendum that higher frame rates will continue to improve performance and should be strived for in performance critical applications.

The results from this study compare favourably with those found by other researchers using different approaches (e.g. Tharp *et al.*, 1992; Watson *et al.*, 1997). It is also worth noting that Barfield and Hendrix (1995) report that frame rates of less than 15 Hz can significantly diminish a user's sense of presence within a VE. It would therefore appear that their level of presence results correlate well with our task performance results.

References

- Barfield, W. and Hendrix, C. (1995). The Effect of Update Rate on the Sense of Presence within Virtual Environments, *Virtual Reality: Research, Development, Applications*, **1**(1): 3–16.
- Bryson, S. (1993). Effects of Lag and Frame Rate on Various Tracking Tasks, *Proceedings of the SPIE - The International Society for Optical Engineering*, Vol. 1915 of *Stereoscopic Displays and Applications IV*, Bellingham, WA, pp. 155–166.

- Bryson, S. and Fisher, S. S. (1990). Defining, Modeling, and Measuring System Lag in Virtual Environments, *Proceedings of the SPIE - The International Society for Optical Engineering*, Vol. 1256, Bellingham, WA, pp. 98–109.
- Cutting, J. E. (1986). *Perception with an Eye for Motion*, MIT Press, Cambridge, MA.
- Cutting, J. E., Springer, K., Braren, P. A. and Johnson, S. H. (1992). Wayfinding on Foot From Information in Retinal, Not Optical, Flow, *Journal of Experimental Psychology: General*, **121**(1): 41–72.
- Falby, J. S., Zyda, M. J., Pratt, D. R. and Mackey, R. L. (1993). NPSNET: Hierarchical Data Structures for Real-Time Three-Dimensional Visual Simulation, *Computer and Graphics*, **17**(1): 65–69. Pergamon Press, UK.
- Gregory, R. L. (1990). *Eye and Brain: the Psychology of Seeing*, fourth edn, Weidenfeld and Nicolson, London.
- Held, R. and Durlach, N. (1993). Telepresence, Time Delay and Adaption, in S. R. Ellis, M. K. Kaiser and A. J. Grunwald (eds), *Pictorial Communication in Virtual and Real Environments*, second edn, Taylor and Francis, chapter 14, pp. 232–246. ISBN 0-74840-0082-6.
- Lampton, D. R., Knerr, B. W., Goldberg, S. L., Bliss, J. P., Moshell, J. M. and Blau, B. S. (1994). The Virtual Environment Performance Assessment Battery (VEPAB): Development and Evaluation, *Presence: Teleoperators and Virtual Environments*, **3**(2): 145–157.
- Piantanida, T., Boman, D. K. and Gille, J. (1993). Human Perceptual Issues and Virtual Reality, *Virtual Reality Systems, Application and Research*, **1**(1): 43–52.
- Rieger, J. H. and Toet, L. (1985). Human Visual Navigation in the Presence of 3-D Motions, *Biological Cybernetics*, **52**: 354–360.
- Singhal, S. K. and Cheriton, D. R. (1995). Exploiting Position History for Efficient Remote Rendering in Networked Virtual Reality, *Presence: Teleoperators and Virtual Environments*, **4**(2): 169–193.
- Smets, G. J. and Overbeeke, K. J. (1995). Trade-Off Between Resolution and Interactivity in Spatial Task Performance, *IEEE Computer Graphics and Applications*, **15**(5): 46–51.
- Swartz, M., Wallace, D. and Tkacz, S. (1992). The Influence of Frame Rate and Resolution Reduction on Human Performance, *Proceedings of the Human Factors Society 36th Annual Meeting*, Santa Monica, CA, pp. 1440–1444.

- Tharp, G. K., Liu, A. M., French, L., Lai, S., Deung-po-gu, Y. and Stark, L. W. (1992). Timing Considerations of Helmet-Mounted Display Performance, *Proceedings of the SPIE - Human Vision, Visual Processing, and Digital Display III*, Vol. 1666, pp. 570–576.
- Wann, J. P., Rushton, S. K. and Lee, D. N. (1995). Can You Control Where You Are Heading When You Are Looking At Where You Want To Go?, in B. Bardy, R. Bootsma and Y. Guiard (eds), *Proceedings of the 8th International Conference on Event Perception and Action*, pp. 171–174.
- Warren, W. H. and Hannon, D. J. (1990). Eye Movements and Optical Flow, *Journal of the Optical Society of America*, **A**(7): 160–169.
- Watson, B., Spaulding, V., Walker, N. and Ribarsky, W. (1997). Evaluation of the Effects of Frame Time Variation on VR Task Performance, *Technical Report GIT-GVU-96-17*, Graphics, Visualization and Usability Centre, Georgia Institute of Technology, Atlanta, GA.