Specification and Evaluation of Level of Detail Selection Criteria[*]

Martin Reddy[†]

Department of Computer Science
University of Edinburgh
Edinburgh, UK

Abstract: Level of detail (LOD) is a technique where geometric objects are represented at a number of resolutions, allowing the workload of the system to be modulated on-line. There are numerous schemes for implementing LOD, using selection criteria based upon an object's distance, size, velocity, or eccentricity. However, little is known about how to specify optimally when a particular LOD should be selected so that the user is not aware of any visual change, or to what extent any particular LOD scheme can improve an application's performance. In response, this paper produces a generic, orthogonal model for LOD based upon data from the field of human visual perception. The effect of this model on the system is evaluated to discover the contribution that each component makes towards any performance improvement. The results suggest that both velocity and eccentricity LOD should be implemented together (if at all) because their individual contribution is likely to be negligible. Also, it is apparent that size (or distance) optimisations offer the greatest benefit, contributing around 95% of any performance increment.

Keywords: computer graphics, level of detail, performance optimisation, visual acuity, visual perception.

---

[*] Published as: Reddy, M. (1998). "Specification and Evaluation of Level of Detail Selection Criteria". Virtual Reality: Research, Development and Application, 3(2): 132-143. © Springer–Verlag, London Ltd.

[†] Enquiries or requests for reprints should be sent to: Martin Reddy, Artificial Intelligence Center, SRI International, 333 Ravenswood Avenue, Menlo Park, CA 94025-3493. E-mail: `reddy@ai.sri.com`.

1 Introduction

The technique of level of detail (LOD) provides virtual environment (VE) designers with a powerful tool to modulate the workload of a graphics system in real-time, and hence to improve the performance of their applications. Essentially, this is done by providing a number of different representations for certain objects, each varying in complexity (e.g. polygon count, degree of texture-mapping, lighting model, etc.). The system can then select the most appropriate representation to use at any point using some selection criterion, or criteria. Many selection criteria have been developed over the past 20 years since the initial work of Clark (1976). We can encapsulate all of these using four principal selection criteria:

Distance : an object's level of detail is based upon its distance from the viewpoint (Carey and Bell, 1997; Chrislip and Ehlert Jr., 1995; Vince, 1993; Kemeny, 1993).

Size : an object's level of detail is based upon a measure of its pixel size, or area, on the display device (Roehl, 1995; Wernecke, 1993).

Eccentricity : an object's level of detail is based upon the degree to which it exists in the periphery of either the display device or the user's field of view (Watson *et al.*, 1995; Reddy, 1995; Ohshima *et al.*, 1996; Funkhouser and Séquin, 1993; Hitchner and McGreevy, 1993).

Velocity : an object's level of detail is based upon its velocity relative to the user, i.e. its velocity across the display device or the user's retina (Ohshima *et al.*, 1996; Amselem, 1995; Funkhouser and Séquin, 1993; Hitchner and McGreevy, 1993).

In addition to these, various other mechanisms have been used to constrain or augment the LOD modulation in some way. For example, fixed frame rate schedulers are used to sustain a desired frame rate (Rohlf and Helman, 1994; Wloka, 1993; Airey *et al.*, 1990), object priority schemes prevent important objects being degraded (Holloway, 1991), hysteresis techniques reduce the effect of objects scintillating at an LOD threshold (Astheimer and Pöche, 1994), and transparency fading regions can be used to smooth the visual transition between LODs (Rohlf and Helman, 1994; Vince, 1993).

Although there is substantial perceptual evidence to support the use of each of these selection criteria, it is often unclear how best to implement any specific scheme, and even how multiple schemes may be combined, thus causing distracting visual artifacts—or popping effects—when an object switches resolution. This is particularly the case for velocity and eccentricity LOD. For example, Watson *et al.* (1995) state that they had no way to decided the extent to which detail could be degraded in the periphery of their display. Also, Ohshima *et al.* (1996) were unsure whether they should take the product of their velocity and eccentricity scaling factors, or whether a minimum function should be used.

In addition to this problem, few quantitative data exist to inform the VE designer about the fundamental utility of any particular selection scheme. For example, Funkhouser and

Séquin (1993)'s seminal paper describes a system that incorporates most of the above selection criteria, but no information is provided about the degree to which any of the components contribute towards the overall performance increment. Hitchner and McGreevy (1993) developed a similarly general LOD model, but again we have no results illuminating the individual merit of each selection criterion.

This paper attempts to provide the VE community with solutions to both of these problems. It aims to provide answers to questions such as, 'How can I implement any specific set of LOD selection criteria in a principled and orthogonal manner?', and, 'How much faster can my application run if I implement, say, eccentricity LOD?'. This information is of obvious importance and relevance to developers and designers of VEs.

2 Aims and Organisation

This paper is organised as follows. First we present a generic system for implementing LOD that is based upon principled models of visual perception. We then embark upon a theoretical evaluation to assess the implications that follow directly from this model. This leads on to an empirical evaluation where we present data from a prototype implementation of this LOD model. Note that the empirical evaluation is concerned with system performance only: the effect of the model on users is dealt with elsewhere (Reddy, 1997a). Finally, we discuss the findings of this work and present our conclusions.

Before proceeding, we should clarify a few points. Many LOD schemes have been developed in the past, so why develop another one? The reason is that practically all attempts to produce general LOD models have been *ad hoc* in nature and normally incorporate various arbitrary variables which have to be instantiated through experimentation. Here we will present a principled model for LOD based upon data from the field of human visual perception. This is done to ensure that the LOD system is founded on the threshold characteristics of our visual system; and therefore that an object's LOD can be selected with the user being able to perceive little or no visual change. That is, we wish to investigate the optimal degree of detail that a computer graphics system need display due to limitations of the user's visual system.

Secondly, it will be apparent that size LOD and distance LOD are essentially the same: as an object moves further away, it subtends a smaller visual angle. In this paper we favour the use of size LOD because it does not require choosing an arbitrary point in the object for the calculation, and also because it is robust to changes in object scale and display resolution. However, distance LOD does offer the advantages of being very efficient and simple to implement, so which approach the VE designer should adopt will depend upon the specific requirements and constraints of an application. For example, distance LOD might be more applicable for an application on a low-end machine, with many multi-resolution models, where the eradication of LOD popping effects is not a critical goal. The reader should note that all results that we present for size LOD are automatically applicable to distance LOD.

Finally, it is worth illustrating that we should not need to incorporate techniques such as hysteresis, fading regions, or priority schemes into a true perceptually-based system. These techniques have been developed to counter the visual anomalies which occur when LOD is used in a non-principled and arbitrary manner. These methods are therefore redundant if we can produce a system which will modulate detail without the user being able to perceive any change.

3 Describing the Model

3.1 Perceptual Background

The human ability to perceive detail is determined by the relative size and contrast of a stimulus (Campbell and Robson, 1968). The size of a stimulus is normally defined in terms of spatial frequency, in units of contrast cycles per degree of visual field (c/deg). This is a measure of the rate of change of intensity over a region, where a high spatial frequency implies a stimulus of high detail.

Through empirical study, it has been confirmed that our ability to resolve spatial frequency varies with respect to velocity across the retina (Kelly, 1979), the degree of displacement into the peripheral field (Rovamo and Virsu, 1979), orientation (Campbell *et al.*, 1966), and also the level of background illumination (Kelly, 1975). The phase of a spatial frequency has no effect on its detectability (Lamming, 1991). It is therefore clear that we have substantial quantitative data available to describe the ability of the human visual system to resolve spatial detail.

However, in order to optimise detail in a VE based upon any model of visual perception we require a computer system with the ability to describe the perceptual content of a scene in terms of this model, and also to quantify the degree of detail that an observer can perceive in terms of this model. With these two facilities, the computer system can judge which detail a user can and cannot see in a computer-generated scene. We therefore require the following functions:

1. A machine-computable mechanism to describe concisely any LOD model in a VE in terms of its component spatial frequencies (c/deg). Only the unique spatial frequencies between two successive LODs are of interest, i.e. those representing a visual change between two models.
2. An efficient mathematical system to predict the highest visible spatial frequency (c/deg) that a standard observer can resolve under various visual conditions (e.g. variable velocity and eccentricity). This is referred to as visual acuity.

3.2 A General Model for Perceptual LOD

One of the pivotal design decisions that we must address is that of where to calculate the spatial frequencies in an object. We have two choices: either we base our analysis on the three-dimensional (3D) geometric definition of an object (its polygons and vertices), or the two-dimensional (2D) rendered image of that object. It is our determination that we must advocate a system which extracts detail from the rendered image of an object because this is the actual information which is presented to the user's visual system. Looking at the geometrical description of an object does not give a reliable indication of what the user eventually sees because the geometry can be displayed differently depending upon a number of factors such as: the particular shading model being used, the effect of any light sources, the use of texture maps or environment (reflection) maps, the simulation time of day, the use of fog or haze effects, transparency, etc.

As a result of this decision, we must compute the spatial frequency content of each object off-line because we are required to know the spatial frequency content of a model before we display it: once a model has been rendered, we have already expended the computational resources that we were seeking to preserve. This process will involve sampling the object from a number of different viewpoints in order to capture its 3D profile (see Figure 1), and then storing these data along with the model's geometric description. A similar approach was adopted by Maciel and Shirley (1995) for their LOD system.

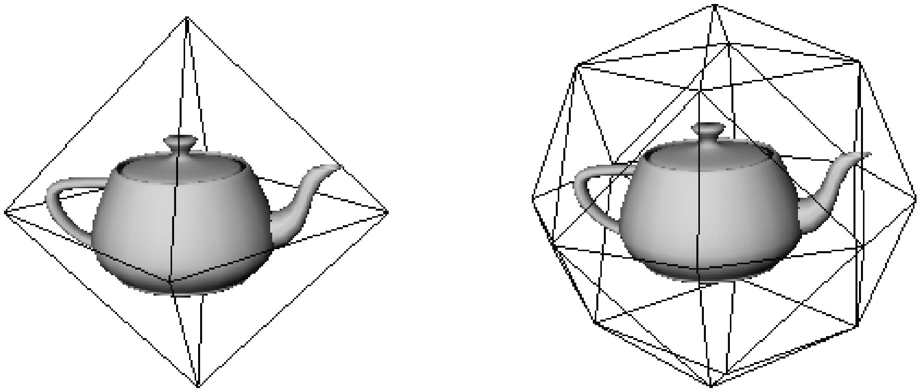

(a) **(b)**

Figure 1: Two examples of the space of possible sample viewpoints on an object's view sphere. In each case, the viewpoint locations are denoted by the vertices of the tessellated bounding sphere. All viewpoints are assumed to be oriented towards the centre of the sphere. (a) illustrates the simplest case of 6 sample viewpoints, while (b) shows a more complex example with 18 viewpoints.

Once on-line, the system must monitor the size on screen, angular velocity, and eccentricity of each object (in units of pixels, deg/s, and deg, respectively). It must then use this information to compute the highest resolvable spatial frequency under those conditions. Then, given the spatial frequency data that were computed during the off-line stage, we can estimate the instantaneous frequency content of each LOD. Finally, the system can then use this information to choose the optimal LOD for the object, e.g. the LOD whose highest unique spatial frequency lies immediately below the limit of vision for the object. That is, the least complex model such that any change evoked will be below the user's threshold of vision.

From the above, it is clear that our general model should be split into two fundamental stages: an off-line (preprocessing) stage, and an on-line (scheduling) stage. The following two sections will detail the implementation of the principal component of each of these stages: the spatial frequency analysis and the visual acuity calculation.

3.3 Evaluating an Object's Perceptual Content

We have resolved that we wish to calculate the spatial frequency content of a computer-generated image. This essentially involves finding a suitable mapping to transform a 2D function of colour values into a 2D function of spatial frequencies. The system developed to perform this task uses an image segmentation algorithm to extract all of the perceptually atomic features in an image. Fourier techniques were considered for this process but were deemed inapposite due to problems of accuracy and generality (Reddy, 1996). Details of the spatial frequency analysis outlined below are given by Reddy (1997a).

1. **Feature Extraction** : a region growing algorithm was used to find the extent of each visual feature in an image. This used a colour difference model, based upon the CIELUV perceptually uniform colour space (Carter, 1989), in order to decide whether a pixel in a colour image would be perceived as an edge to a feature, or as part of that feature.
2. **Spatial Frequency Calculation** : the spatial frequencies in each feature were calculated by finding the length of the feature at various orientations. At this stage, these frequencies are in terms of pixels only, i.e. the units of spatial frequency at this stage are cycles per pixel (c/pixel).
3. **Spatial Frequency Transformation** : the relative spatial frequency values for each feature were scaled into units of c/deg. This transformation can be performed once we know the field of view (FOV) of the display device.

3.4 Estimating a User's Visual Acuity

Reddy (1997b) presents the development of a computational model for spatiotemporal visual acuity with close reference to work in the field of visual perception. This results in an equation for spatial frequency, H, which can be defined as:

$$H(v, E) = G(v) \times M(E) \ \ \mathrm{c/deg}, \tag{1}$$

where,

$$G(v) = \begin{cases} 60.0, & \text{when } v \leq 0.825 \\ -27.78\log_{10}(v) + 57.69, & \text{when } v > 0.825 \end{cases} \tag{2}$$

$$M(E) = \begin{cases} 1.0, & \text{when } E \leq 5.79 \\ 7.49/(0.3E+1)^2, & \text{when } E > 5.79. \end{cases} \tag{3}$$

Here, v represents angular velocity (deg/s) and E represents eccentricity (deg). It is worth noting that we have now solved Ohshima *et al.* (1996)'s dilemma by showing that the product of the velocity and eccentricity scaling factors should be taken; $G(v)$ and $M(E)$ in our model, respectively.

The above model would be sufficient for evaluating the highest visible spatial frequency for objects in the real world; however the angular resolution of a computer display limits the size of detail which users can experience. We can incorporate this factor into our model by introducing the notion of a highest displayable spatial frequency, ξ. This characterises the highest frequency which can be displayed by an output device, depending upon its field of view and pixel resolution.

Specifically, a single pixel can be considered to be half a contrast cycle (Reddy, 1997a). Therefore the number of cycles that can be displayed on a device is equal to half its pixel resolution. This can be transformed into a value of spatial frequency once we know the FOV of the display using the following simple equation:

$$\xi = \max\left(\frac{horizPixels}{2 \times horizFOV}, \frac{vertPixels}{2 \times vertFOV}\right) \ \text{c/deg}, \tag{4}$$

where $horizPixels$ and $vertPixels$ are the horizontal and vertical pixel resolutions of the display, and $horizFOV$ and $vertFOV$ are the horizontal and vertical angular resolutions of the display (in degrees). Our final visual acuity model can therefore be given as:

$$\min(H(v,E), \xi). \tag{5}$$

4 Theoretical Evaluation

Now that we have a computational model for visual acuity, we can investigate some of the implications that result from this relationship. For example, we can now answer questions such as, 'How fast does an object have to travel, or how far does it have to progress into the periphery, before we can reduce its detail?'. We can also speculate whether these situations are likely to occur regularly as a user navigates through a VE.

4.1 A Desktop Example

To begin, let us take the example of a typical desktop display: 40 × 30 cm screen size, 1280 × 1024 pixel resolution, and viewed at a distance of 40 cm. This gives us a field of view of: 53.1 × 41.1 deg. Finally, we can use Equation 4 to find the angular resolution of the display in terms of spatial frequency: 12.0 × 12.5 c/deg.

We can therefore see that for this example we can never be presented with a stimulus of spatial frequency greater than 12.5 c/deg. This is about one fifth of the maximum spatial frequency that we can generally resolve (60 c/deg: Campbell and Gubisch, 1966). So what effect does this have on our system?

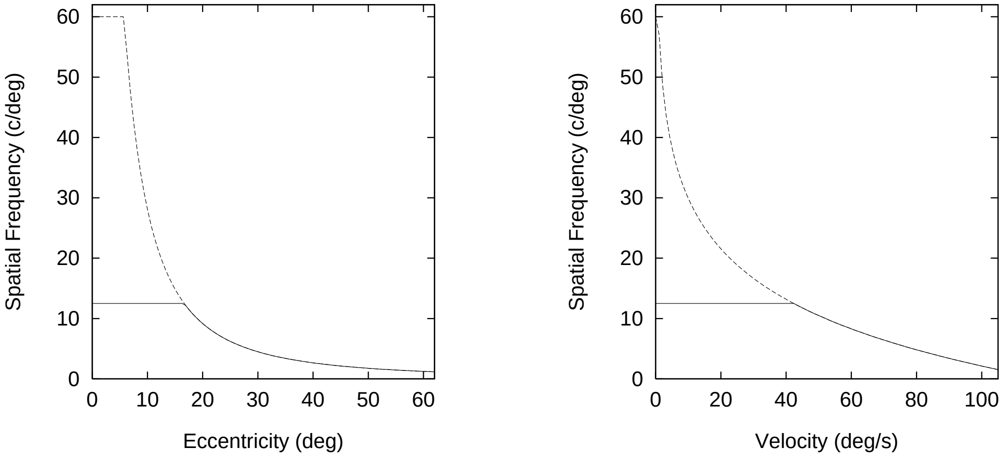

(a) **(b)**

Figure 2: The highest visible spatial frequencies for a user viewing a display where ξ = 12.5 c/deg. These are shown for (a) increasing eccentricity, and (b) increasing velocity. The eccentricity curve (broken line, a) is defined by $H(0, E)$, and the velocity curve (broken line, b) is given by $H(v, 0)$.

Figure 2 presents graphs of the eccentricity and velocity components of our model. Both of these have been thresholded using a highest displayable spatial frequency of 12.5 c/deg to illustrate the potential stimuli in our desktop example. From these two graphs we can observe firstly that we do not begin to perceive less detail in our peripheral field until an eccentricity of about 17 deg is reached. More striking than this is the result that our effective spatial perception will not degrade until an object exceeds a velocity of about 42 deg/s.

In terms of our display example (which, recall, occupies 53.1 × 41.1 deg of arc) these results mean that an object would have to be displaced horizontally from the focus point (where the user is looking) by around one third of the display—or that it would have to travel from the left edge of the display to the right edge in around 1.26 seconds—before it would be possible to select a lower LOD model.

These are extreme and isolated cases. When we combine the effect of both velocity and eccentricity, then the above figures are substantially mitigated. For example, at an eccentricity of 10 deg, then an object need only travel above roughly 13 deg/s before detail could become potentially invisible. It is therefore likely that, taken on their own, eccentricity and velocity based optimisations will provide a meager performance advantage. However, if the two are implemented together, then this will produce a synergic speedup.

4.2 An Immersive Example

We investigate the implications for immersive systems by taking one specific display device for illustration. For example, the i-glasses! head-mounted display (HMD), originally manufactured by Virtual I-O. This LCD based unit is indicative of currently popular, cheap HMD systems, both in terms of resolution and FOV. The device has a considerably smaller field of view than the desktop case (30×24 deg versus 53×41 deg). Also, given the pixel resolution of 263×230, its angular resolution is almost one third that of our desktop example (4.4×4.8 c/deg compared to 12.0×12.5 c/deg).

For ξ = 4.8 c/deg, we find that an eccentricity of 29 deg, or a velocity of 80 deg/s, must be exceeded before it would be possible to degrade objects without the user perceiving the modulation (refer to Figure 2). As the horizontal FOV of the i-glasses! is 30 deg, it is apparent that employing only eccentricity LOD will have essentially no benefit. Also, when employing only velocity LOD, then an object must travel across the entire display in under 0.375 seconds before any optimisation can occur. However, taking these two components together the situation is less drastic. For example, at an eccentricity of 15 deg, an object need only travel at 24 deg/s before we may begin to reduce LOD. The relationship between eccentricity and velocity for this example is illustrated in Figure 3.

There is greater benefit to be gained in immersive virtual reality (VR) systems because, whenever a user moves their head, the entire scene moves in relation to their viewpoint. A user's head rotation can often be in excess of around 180 deg/s, thus providing considerable opportunity for most objects to be degraded during exaggerated head movements. This is particularly desirable because many users tend to feel disoriented or nauseous when they move their heads rapidly and the system does not keep up with their movements (Regan, 1995; Frank *et al*., 1988). By reducing detail during head movements we can improve the update rate of the system to make the simulation appear more smooth and interactive. Our model automatically supports this feature because we measure the velocity of objects across the display device, relative to the user's viewpoint.

In summary, the lower angular resolution and smaller FOV of many of the current generation LCD HMDs means that there is practically no benefit in using either eccentricity or velocity optimisations on their own (in order to satisfy the requirement that the user should not be able to perceive the modulation of detail). However, by using the two optimisations in parallel, their contribution can become significant (Figure 3). This is particularly evident in situations when the user makes large head rotations. This latter optimisation could prove a major factor

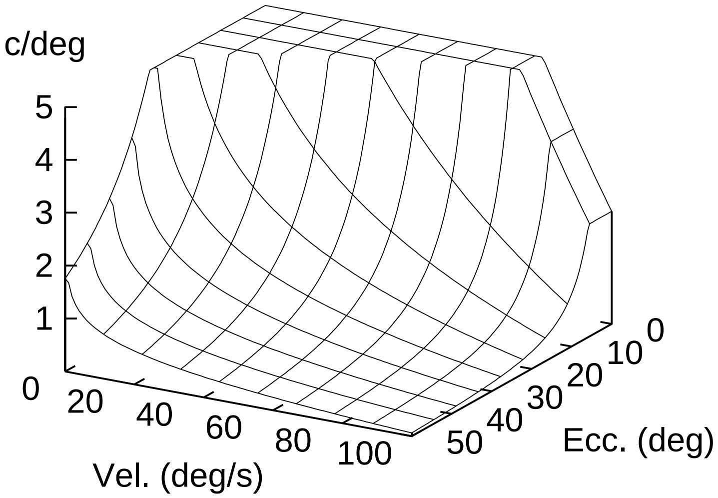

Figure 3: The effective spatiotemporal threshold surface for the i-glasses! HMD from Equation 5 with $\xi = 4.8$ c/deg. All points below the surface represent combinations of spatial frequency, velocity, and eccentricity that are potentially visible to the user.

in making immersive VR systems appear less disorienting.

5 Empirical Evaluation

In order to usefully answer questions such as, 'How much faster can my application run if I implement LOD scheme *X*?', then we should endeavour to evaluate an actual system. Subsequently, a prototype implementation was developed for the above LOD model, and a suitable experiment devised to assess the model. Note that we will not assess the visual effect of our model on the user. This has been done elsewhere using a similar experimental setup. The results of that study showed that the LOD modulation was not perceived by any of 20 subjects who each performed 64 trials (Reddy, 1997a). We therefore assert that our visual acuity model achieves the goal of eliminating LOD popping effects.

The experimental environment consisted of a number of randomly positioned objects (see Figure 4). A set of predefined paths were then navigated through this environment for a duration of five seconds in each case (i.e. no user interaction was involved). Four levels of detail were used for the test object, containing 3928, 834, 254, and 76 triangles respectively. The extent of LOD optimisation could be varied on a per-trial basis to investigate the effect of using different LOD techniques.

This particular scenario was chosen because it is inherently motion-rich, with many objects of different sizes, and a good proportion of peripherally located features. It therefore offers good opportunity to exercise all aspects of our model. Also, 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.

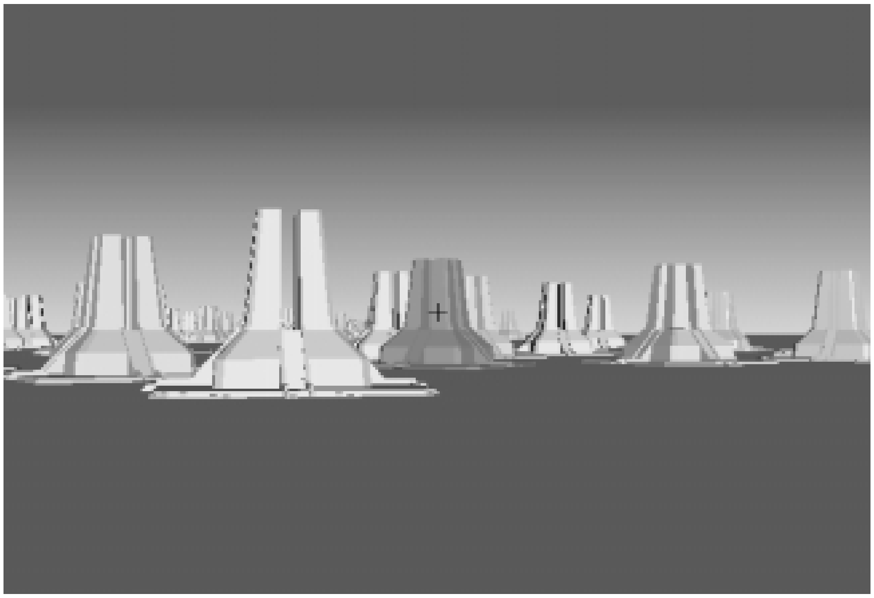

Figure 4: An example screen shot of the test environment used for all of the empirical studies.

The experiments were performed on a Silicon Graphics Inc. (SGI) Onyx RealityEngine^2 with one 200 MHz processor and 128 MB of RAM. All remote access to the machine was disabled for the duration of each study and all disk writes were made only upon completion of each trial. The display FOV was assumed to be 43.6×33.4 deg, and the simulated forward velocity was 4.36 m/s. All timings were recorded using a sub-microsecond resolution clock.

5.1 Analysis of Performance Speedup

Objective

The aim of the first study was to assess the extent to which our prototype implementation improves the frame rate of a VE. This was contrasted for a number of VEs with different numbers of objects in order to give an indication of how speedup varies with VE complexity for the test application.

Method

The average frame rate (Hz) for a trial was found by dividing the number of frames rendered by the total time for the trial. Twenty trials were performed where the number of objects in the VE were kept constant (but their initial positions randomly varied). The final frame rate figure for such a set of trials was found by averaging the results of all 20 trials. Each set was repeated for the case where no LOD optimisations were applied and when full perceptual LOD filtering was applied. Finally, a number of these sets of trials were performed for environments with different numbers of the test object (10, 20, 50, 100, 200, 500, and 1000 objects). That is, 280 trials were performed in total (7 complexities × 2 LOD cases × 20 trials).

Results and Discussion

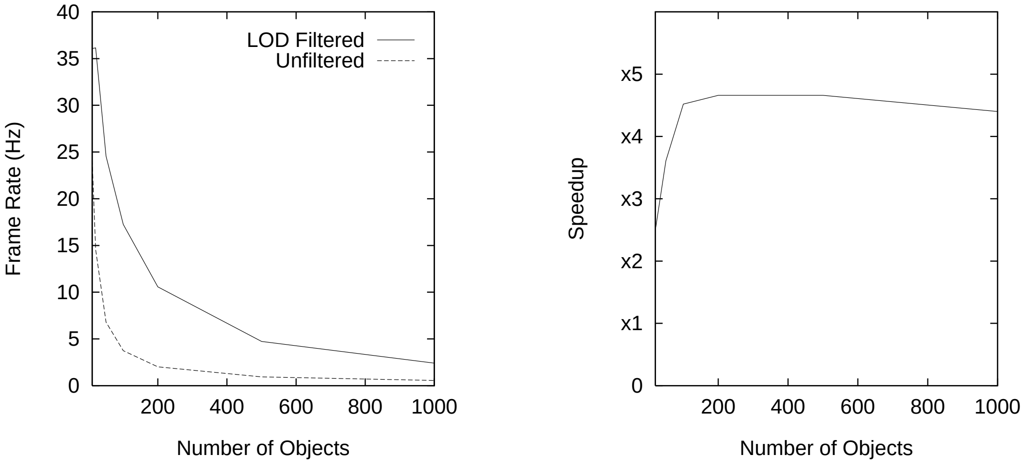

(a) **(b)**

Figure 5: Results from the performance speedup analysis. (a) contrasts the average frame rate for the test application under normal, unfiltered conditions (the broken line) and when perceptual LOD optimisations were employed (the solid line). (b) presents the same data in terms of the relative speedup which was achieved when using the perceptual LOD system over the normal, unfiltered case.

The results from this study are presented in Figure 5. From Figure 5(a) we can observe a consistent and marked increase in frame rate when our LOD optimisations were employed. To describe the degree of this increment, Figure 5(b) illustrates the relative increase in performance which occurred when our perceptual optimisations were used. From this we can see that an average speedup of over 4.5 times was quickly achieved for environments with about 50 objects, with very slight depreciation for more complex environments with over 500 objects.

Care should be taken when interpreting the initial sharp rise in speedup for environments with less than 50 objects. This does not necessarily imply that less complex environments offer fundamentally smaller speedup values. We must also consider the fact that on SGI workstations, the frame rate is constrained to be an integer multiple of the video refresh rate. Our video refresh rate was 72 Hz and so if the simulation was unable to run at 72 Hz, then it would drop down to the next integer multiple of 72, i.e. 36 Hz. Therefore, what we are most likely observing in the sharp speedup rise for less than 50 objects is the point where the optimised environment is capable of being rendered faster than 36 Hz, but not as fast as 72 Hz, and so it is restricted to only 36 Hz. The relative speedup is therefore confounded by this additional factor for simple environments.

5.2 Breakdown of Model Components

Overview

In this second computational study we will attempt to gauge the extent to which each of the three principal components of our model are utilised. That is, how much does size LOD, eccentricity LOD, and velocity LOD contribute towards the total increase in performance for our test application.

We have already discovered that analysing the eccentricity or velocity components in isolation would be fruitless. Therefore we shall compare the case where all optimisations are used, against the case where only the velocity and eccentricity based optimisations are used. In effect, this provides us with an indication of the additional benefit which can be accrued over traditional size LOD when we also incorporate optimisations based upon the velocity and eccentricity of objects. It is a valid activity to contrast these two situations because the size LOD optimisations are independent of the combined velocity and eccentricity components, and so a linear relationship exists between these two factors.

Method

The average level of detail (1–4) for a frame was found by dividing the sum of all objects' LOD by the number of objects. An average figure for an entire trial was found by dividing all of the frame averages by the total number of frames displayed. As before, 20 such trials were performed, and the average LOD figure from all of these was found. This measure was used as an assessment of the degree to which detail was optimised; with higher values indicating a greater degree of optimisation.

Two level of detail cases were compared: one where full optimisations were employed (size, velocity, and eccentricity), and the other where detail was reduced based only upon the velocity and eccentricity of objects. This relationship was analysed for environments of various complexities (10, 20, 50, 100, 200, 500, and 1000 objects). 280 trials were therefore performed in total (7 complexities $\times$ 2 LOD cases $\times$ 20 trials).

Results and Discussion

Figure 6 presents the data from this study. The most obvious result from this graph is that size optimisations appear to account for the bulk of any reduction of detail (and hence the increase in performance). The combined contribution of velocity and eccentricity optimisations rose

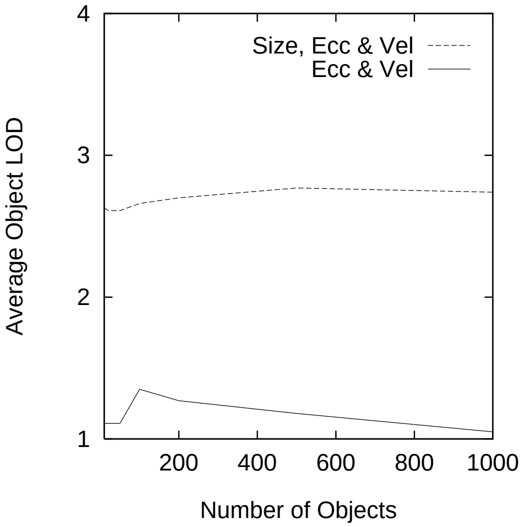

Figure 6: Comparison of average LOD (1–4) during a trial with (broken line) and without (solid line) size optimisations.

to a maximum of around 20% of the total LOD reduction for around 100 objects. This dropped down to as little as 3% of the total reduction for 1000 objects. In general, we can state that roughly 95% of all performance improvement can be attributed to size LOD only.

6 Discussion

We look first at the perceptual LOD model, and then explore some of the results from its evaluation. In the first instance, our visual acuity model is orthogonal in that each of the three components can be used individually or in any desired combination. For example, we can disable the modulation of detail with respect to eccentricity by simply instantiating all eccentricity values to 0 deg. Our model will then still function correctly and accurately, but will only modulate detail based upon the size and velocity of an object.

Although we were not primarily concerned with the issue of fixed frame rate LOD, it should be noted that our model can be easily extended to support this scheme. For example, a simple reactive fixed frame rate system could be implemented by incorporating a global scaling factor, s, into our visual acuity calculation. Then, if the time for any frame exceeds the desired frame period, we simply reduce the value of s accordingly. Alternatively, if we find that we have spare compute time after completing a frame, then we can increase the value of s until it reaches 1.0 (there is of course no need to exceed a scaling factor of 1.0). It is important to note that this enables us to develop a fixed frame rate system that will degrade resources in a perceptually linear manner, i.e. the less overloaded the system is (larger s), then the less perceivable the changes in scene complexity will be, and vice versa.

The full LOD system (size, eccentricity, and velocity) produced a 4.5 fold improvement in frame rate. This is a substantial increase, but not an exceptional one. It should be noted that various other researchers have achieved comparable performance increments using size and eccentricity based optimisations. For example, Ohshima *et al.* (1996) experienced a five-fold improvement in frame rate for their LOD system, and Levoy and Whitaker (1990) experienced an improvement in rendering performance of 4.6 times for their perceptually optimised ray tracer.

All the empirical studies were performed for environments with a fixed FOV and relatively slow forward velocity. An interesting study would be to investigate the effect of changing these parameters. Although this was not done here, it is likely that if we were to increase the forward velocity through the environment, then the velocity component of our model would be given more chance to contribute to the overall reduction in detail, and performance should improve as a result.

It would be tempting to make the same analogy for FOV, but this would be a specious assumption. Although increasing the FOV of the display would allow objects to exist at greater eccentricities (and hence theoretically allow them to have their detail reduced further), we must also consider the angular resolution of the display. If we take the same display and simply stretch it over a wider field of view, then the individual pixels will become larger and so the angular resolution—and hence the highest visible spatial frequency—will drop. So, although the field of view is larger, objects must move further into the periphery before their detail drops below threshold and a lower LOD can be selected.

7 Conclusion

The ambition of this paper has been to develop a perceptually accurate model for LOD such that the user perceives no visual change when an object switches resolution. A secondary goal has been to provide a quantitative indication towards the fundamental utility of each component of this model. We summarise the result of this work as follows:

1. We have shown how to design a perceptually modulated LOD system using sound knowledge from the field of visual perception. Furthermore, we have shown that an efficient implementation of this can be produced for use in a real-time computer graphics system.
2. We have also shown that perceptually modulated LOD can substantially reduce lag in a VR system. In our test application, a factor 4.5 improvement in frame rate was observed. Elsewhere we show that this was achieved with little or no perceivable drop in image quality (Reddy, 1997a).

With the results that were obtained from our theoretical and empirical evaluations of this perceptual LOD system, we can present the chief findings of this paper as follows:

- A complete implementation of perceptually modulated LOD (i.e. one incorporating velocity and eccentricity factors) is more applicable to an immersive VR system than a desktop one because:

    1. The display can only be optimised for a single individual as two users viewing the same scene could be looking at different regions of the screen and scanning at different speeds. Using a head-mounted display circumvents this ambiguity.
    2. Most immersive systems employ some form of head tracking. This is not normally available on the desktop, making the user's focus point unknown. The best we can do in the latter case is to instantiate the focus point to a fixed location on the display, such as the centre of the screen.
    3. In a head-tracked system, velocity LOD will be particularly beneficial for optimising the detail of a scene during periods of high user head motion. This is not generally possible under desktop conditions.
    4. The visual environment within a head-mounted display is much more predictable and controllable than for a computer monitor. We are therefore less concerned about effects such as background illumination, glare, reflections, etc. affecting the user's percept.

- From the examples provided in the theoretical evaluation, it is apparent that employing either eccentricity or velocity LOD on their own will prove unprofitable (this is particularly so for immersive systems using a low angular resolution HMD). However, if these two components are combined, then the product will be significantly greater than the sum of the individual contributions (Figure 3). We therefore submit that eccentricity and velocity optimisations, if used at all, should be used in combination.

- It appears that, even when combined, eccentricity and velocity contribute to the total speedup to a relatively minor extent, an average of around 5% for our test application (see Figure 6). We can therefore state that traditional size (or distance) based LOD methods provide the largest opportunity to improve the performance of a system. Note however that this result is only applicable under situations where the user's head is static. Under dynamic head movements the velocity component could potentially offer a far greater contribution to the reduction of detail in a scene.

References

Airey, J. M., Rohlf, J. H. and Frederick P. Brooks, J. (1990). Towards Image Realism with Interactive Update Rates in Complex Virtual Building Environments. *ACM SIGGRAPH Special Issue on 1990 Symposium on Interactive 3D Graphics* 24(2): 41–50.

Amselem, D. (1995). A Window on Shared Virtual Environments. *Presence: Teleoperators and Virtual Environments* 4(2): 130–145.

Astheimer, P. and Pöche, M.-L. (1994). Level-Of-Detail Generation and its Application in Virtual Reality. *Proceedings of the VRST '94 conference*: 299–309.

Campbell, F. W. and Gubisch, R. W. (1966). Optical Quality of the Human Eye. *Journal of Physiology* 186: 558–578.

Campbell, F. W., Hulikowski, J. J. and Levinson, J. (1966). The Effect of Orientation on the Visual Resolution of Gratings. *Journal of Physiology* 187: 427–436.

Campbell, F. W. and Robson, J. G. (1968). Application of Fourier Analysis to the Visibility of Gratings. *Journal of Physiology* 197: 551–566.

Carey, R. and Bell, G. (1997). *The Annotated VRML 2.0 Reference Manual* (Reading, MA: Addison-Wesley). ISBN 0-201-41974-2.

Carter, R. C. (1989). Calculate (Don't Guess) the Effect of Symbol Size on Usefulness of Color. *Proceedings of the Human Factors Society 33rd Annual Meeting*: 1368–1372.

Chrislip, C. A. and Ehlert Jr., J. F. (1995). *Level of Detail Models for Dismounted Infantry in NPSNET–IV.8.1*. Master's thesis, Naval Postgraduate School, Monterey, CA.

Clark, J. H. (1976). Hierarchical Geometric Models for Visible Surface Algorithms. *Communications of the ACM* 19(10): 547–554.

Frank, L. H., Casali, J. G. and Wierwille, W. W. (1988). Effects of Visual Display and Motion System Delays on Operator Performance and Uneasiness in a Driving Simulator. *Human Factors* 30(2): 201–217.

Funkhouser, T. A. and Séquin, C. H. (1993). Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments. *Computer Graphics (SIGGRAPH '93 Proceedings)* 27: 247–254.

Hitchner, L. E. and McGreevy, M. W. (1993). Methods for User-Based Reduction of Model Complexity for Virtual Planetary Exploration. *Proceedings of the SPIE – The International Society for Optical Engineering* 1913: 622–36.

Holloway, R. L. (1991). Viper: a Quasi-Real-Time Virtual-Worlds Application. *UNC Technical Report No. TR-92-004*. Deptartment of Computer Science, University of North Carolina, Chapel Hill, NC.

Kelly, D. H. (1975). Spatial Frequency Selectivity in the Retina. *Vision Research* 15: 665–672.

Kelly, D. H. (1979). Motion and Vision. II. Stabilized Spatio-Temporal Threshold Surface. *Journal of the Optical Society of America* 69(10): 1340–1349.

Kemeny, A. (1993). A Cooperative Driving Simulator. *Proceedings of the International Training Equipment Conference (ITEC)*: 67–71.

Lamming, D. (1991). Spatial Frequency Channels, *in* J. R. Cronly-Dillon (ed.), *Vision and Visual Dysfunction: Limits of Vision*. Vol. 5 chapter 8: 97–105 (MacMillan Press, Ltd.). ISBN 0-333-52713-5.

Levoy, M. and Whitaker, R. (1990). Gaze–Directed Volume Rendering. *ACM SIGGRAPH Special Issue on 1990 Symposium on Interactive 3D Graphics* 24(2): 217–223.

Maciel, P. and Shirley, P. (1995). Visual Navigation of Large Environments Using Textured Clusters. *Symposium on Interactive 3D Graphics*: 95–102.

Ohshima, T., Yamamoto, H. and Tamura, H. (1996). Gaze-Directed Adaptive Rendering for Interacting with Virtual Space. *Proceedings of the IEEE Virtual Reality Annual International Symposium (VRAIS)*: 103–110.

Reddy, M. (1995). Musings on Volumetric Level of Detail for Virtual Environments. *Virtual Reality: Research, Development and Application* 1(1): 49–56.

Reddy, M. (1996). A Measure for Perceived Detail in Computer-Generated Images. *Technical Report ECS-CSG-19-96*. Department of Computer Science, University of Edinburgh.

Reddy, M. (1997a). *Perceptually Modulated Level of Detail for Virtual Environments*. Ph.D. Thesis (CST–134–97). University of Edinburgh.

Reddy, M. (1997b). The Development and Evaluation of a Model of Visual Acuity for Computer-Generated Imagery. *Technical Report ECS-CSG-30-97*. Department of Computer Science, University of Edinburgh.

Regan, C. (1995). An Investigation into Nausea and Other Side-effects of Head-coupled Immersive Virtual Reality. *Virtual Reality: Research, Development, Applications* 1(1): 17–32.

Roehl, B. (1995). *AVRIL Technical Reference*. On-line documentation for the AVRIL graphics package. Version 2.0, March 28.

Rohlf, J. and Helman, J. (1994). IRIS Performer: a High Performance Multiprocessing Toolkit for Real-Time 3D Graphics. *Computer Graphics (SIGGRAPH '94 Proceedings)*: 381–394.

Rovamo, J. and Virsu, V. (1979). An Estimation and Application of the Human Cortical Magnification Factor. *Experimental Brain Research* 37: 495–510.

Vince, J. (1993). Virtual Reality Techniques in Flight Simulation. *in* R. A. Earnshaw, M. A. Gigante and H. Jones (eds), *Virtual Reality Systems* (Academic Press Ltd.). ISBN 0-12-227748-1.

Watson, B., Walker, N. and Hodges, L. F. (1995). A User Study Evaluating Level of Detail Degradation in the Periphery of Head-Mounted Displays. *Proceedings of the FIVE '95 Conference*: 203–212. QMW, University of London, UK.

Wernecke, J. (1993). *The Inventor Mentor: Programming Object-Oriented 3D Graphics with Open Inventor(TM), Release 2* (Reading, MA: Addison-Wesley). ISBN 0-201-62495-8.

Wloka, M. M. (1993). Incorporating Update Rates into Today's Graphics Systems. *Technical Report CS-93-56*. Department of Computer Science, Brown University, Providence, RI.