Chapter 7

Conclusions

'There are more things in heaven and earth, Horatio,
than are dreamt of in your philosophy.'

(William Shakespeare, *Hamlet*)

This final chapter will consolidate all of the material presented thus far and focus the aims and achievements of this treatise. This will involve presenting conclusions on our work in the area of perceptually modulated LOD, and explicitly highlighting the original contribution to knowledge which has been made. In addition to this, we will offer a number of topics for further research. This will conclude the main body of our thesis.

7.1 Assessment of Perceptually Modulated LOD

The single driving ambition of this dissertation has been to reduce lag in VR systems by optimising the graphical content of objects without affecting the quality of images perceived by the user. This has been achieved. Specifically, let us note the following summaries:

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 and subsequently improve user performance, e.g. in our test application, we gained up to a five-fold improvement in frame rate, a three-fold improvement in the accuracy of users to perform the task, and an improvement in user response time of $\times\ 1\frac{2}{3}$.

With the results that were obtained from our implementation of a perceptually modulated LOD system, and the subsequent discussion that was performed, we can present the chief findings of this thesis as follows:

- We submit that a complete implementation of perceptually modulated LOD is more applicable to an immersive VR system than a desktop one. This is due to the following reasons:

  1. The display can only be optimised for a single individual because 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 and so we can never know where the user is actually looking in a desktop scenario. The best we can do in the latter case is to instantiate the focus point to a fixed point on the display.
  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 possible under desktop conditions.
  4. The visual environment within a head-mounted display is much more predictable and controllable than for a computer monitor sitting on a desk in an office. We are therefore less concerned about effects such as background illumination, glare, reflections, etc. affecting the user's percept.

- From the examples that were provided in Section 6.2.2, it is apparent that employing either eccentricity or velocity LOD on their own will prove unprofitable (this is particularly so for immersive systems using an LCD based HMD). However, if these two components are combined, then the product will be significantly greater than the sum of the individual contributions (refer to Figure 6.3). 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 $\sim$10% for our test application: see Figure 5.16). 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.

- It is important to appreciate the modeling issues associated with perceptual LOD. Our goal is to change the geometrical complexity of objects based upon perceptual judgements. The LOD generation software should reflect this fact by incorporating some mechanism to specify the scale of a simplification. We therefore submit that, to be useful, all LODs should exhibit a range of visual complexities as well as a range of geometrical complexities, i.e. as the LOD of an object decreases, there should be a corresponding drop in the spatial frequency content of each model.

7.2 Thesis Contribution

Now that we have summarised the conclusions of this thesis, we will explicitly highlight how our work makes an original contribution to knowledge. Accordingly, the following seven points are offered as the principal scholarly contributions that have been made by this thesis:

1. A thorough literature review of the field of level of detail, including the process of polygon simplification.

2. The identification of a relevant metric from the visual psychology literature to describe and model the perception of spatial detail.

3. A demonstration of the applicability of this metric to the field of real-time, full-colour computer graphics.
4. The development of a computer algorithm to assess the perceptual content of a computer-generated scene in terms of this metric.
5. The collation and extension of mathematical models of human contrast sensitivity to form an efficient, machine-computable estimate of the visibility of detail; incorporating effects of angular velocity and retinal eccentricity.
6. The development of a polygon simplification framework to support the use of level of detail in a perceptually predictable manner.
7. The description, implementation, and evaluation of a perceptually modulated level of detail system based upon the above work.

Perhaps the most valuable contribution this thesis makes is with regard to the cross-disciplinary application of theory from the field of visual perception to the field of computer science. This in itself is no minor feat because the visual perception literature is voluminous and, as Ware and Knight (1995) comment, often couched in arcane jargon. It is therefore hoped that this work has helped to make the results of vision research more accessible to computer scientists, and to illustrate how knowledge from perceptual psychology may usefully be applied to our domain.

7.3 Further Work

To complete the thesis we will present a few suggestions for further research. We have covered a lot of material in these pages; spanning topics such us image segmentation, colour vision, Fourier analysis, physiology, polygon simplification, threshold vision, system latency, psychophysics, and level of detail. There is therefore much that could be investigated furthered. However, to contain the discussion, we have selected the following four topics which we feel offer the most potential for expansion.

7.3.1 Improved Metrics for Assessing Perceived Detail

If there is only one problem which has delayed the development of a perceptually modulated LOD system before now, then it is the fact that no complete theory currently exists to explain how we perceive and differentiate arbitrary stimuli. In this thesis we have proposed and implemented a simplistic scheme using colour difference theory to segment an image. This appears to give reasonable results, but we suggest that more work could be done to extend or replace this particular technique, and hence to produce a more perceptually robust solution. For example, even if the notion of using colour differences is assumed valid, much work needs to be done to improve the accuracy and generality of current colour difference theory.

Given the nature of such an activity, this research would be best furthered by a body with cross-disciplinary skills in visual perception and computer graphics. Of course, the development of a general model of perceptual image segmentation might be too complicated at the current time; however, if we restrict ourselves to the genus of computer-generated images, then a suitable solution may present itself.

It is important to note that the development of such a metric would have value transcending this single application. For example, we could use this technique to compare the perceptual impact of two different rendering algorithms, or analyse the effect of introducing various antialiasing techniques. Alternatively, such a tool could be employed by perceptual psychologists to further analyse the intricacies of our visual system using the power and flexibility of current computer graphics technology.

7.3.2 Transparent LOD Generation

One of the achievements of this work has been to remove all subjectivity from the application of LOD techniques by basing these on sound principles of visual perception. The designer of a VE would therefore not need to worry about which values to supply for the switching thresholds between successive models: this is done automatically. However, on the modeling side, there is still an element of subjectivity because the designer has to create some number of LODs by supplying certain parameters to a polygon simplification algorithm. The ideal situation would be if the designer had a complex model that they wanted to use in a VE, and then the system would automatically produce a good suite of LOD models, with no direction required on the designer's behalf.

Associated with this is the need for polygon simplification techniques which complement a VR system. We noted in Section 2.2.4.2 (Page 49) that many current techniques do not address all of the issues inherent in a VR system. We require techniques which are developed to be useful in our domain, rather than blindly using methods from other fields. Furthermore, these techniques should incorporate a mechanism to specify the scale of a simplification in order to constrain the perceptual effect of reducing an object's detail.

One recent development in this area is worth particular comment. Luebke (1996) presents a system in which the entire scene is stored in an octree data structure; so that volumes further down the structure represent smaller volumes of the world space (and hence finer detail). When the system decides that a particular volume level in the octree is below threshold (e.g. projects to less than a pixel), then all of the vertices in that volume are collapsed to a single point, discarding any degenerate polygons that arise. The reader will recognise the conceptual parallels to the work of Rossignac and Borrel (1992). This global, real-time view to optimising detail offers many benefits. Firstly, it requires no intervention on the part of the VE designer. Secondly, it has the ability to merge groups of adjacent objects into a single entity which could not be done under an LOD paradigm.

7.3.3 An Eye Tracking Feasibility Study

On a number of occasions we have acknowledged that eye tracking is an essential component of any true perceptually modulated LOD system. However the remit of our work has not enabled us sufficiently to investigate the practicalities of this conjunction.

Our prototype evaluation was performed in such a manner that the observers' gaze was restricted to a known spatial location. Our results should therefore be applicable to any system where the observer's gaze position is known. However, there are a number of issues associated with the integration of an eye tracking facility into a perceptually modulated LOD system which should be explicitly investigated. We therefore suggest that the following would be a purposive direction of research activities.

- We require more information about the availability and performance of plausible eye tracking technologies for use in immersive VR systems. These devices have invariably been used for clinical applications in the past: we require an appreciation of the field from our perspective. Also, we need to know how active the field is, and how the technology might mature over a 5–10 year time frame.

- It would be illuminating to investigate our perceptually modulated LOD techniques in a system which supports eye tracking. For example, how do factors such as the tracker's dynamic response, resolution, accuracy, and drift affect the systems ability to modulate detail with no perceptual defect? And how is this likely to change as eye tracking technology matures?

- It is important to know the extent to which eye tracking is actually required. We have already suggested that for some applications, head tracking might be sufficient. Work needs to be done to investigate the various cost, performance, and benefit implications of utilising eye tracking technologies.

7.3.4 A General Sensory Paradigm

Throughout this thesis we have focussed purely on the issues concerning our visual sensorium. However, it is interesting to note certain parallels with other senses. For example, in the field of auditory perception there are curves known as **audibility functions** (AFs) which describe the range of tone frequencies that a normal young adult can hear. Similar descriptions exist in the field of tactile perception to describe the touch sensitivity and acuity of individuals. In fact, in the latter case, devices known as tactile gratings are used which contain a series of grooves of a particular depth and spatial frequency. The analogy to contrast gratings and contrast sensitivity functions should be obvious.

We therefore raise the question of whether the perceptual optimisations we have developed here could be extrapolated to apply to other senses. This is a purely speculative suggestion: the author has not investigated the feasibility or value of such an effort. However, it is offered here as an interesting extension to our work, and one which may warrant further investigation.