

REACT: REal-time Adaptive Collision Testing

An Interactive Vision approach

Carol A. O'Sullivan Ronan G. Reilly

University College Dublin
Trinity College Dublin
email: cosulliv@cs.tcd.ie, Ronan.Reilly@ucd.ie

Abstract. As the demand for high levels of interaction in computer systems increases, so too does the need for real-time, interactive animation. Detecting collisions between geometrically modelled objects remains a major bottleneck in areas such as Virtual Reality (VR). In order to maintain a constant frame-rate, a trade-off between speed and accuracy is necessary. This is possible if, at each frame, potential collisions are graded by their importance to the viewer's perception. An appropriate Level Of Detail (LOD) at which to test each object may then be chosen, based on the importance of the collision in which it is involved. We adopt some ideas from an emerging area of research, Interactive Vision, and propose a scheme which uses an eye-tracking device to locate the position of the user's gaze. This, along with other perceptual criteria, may be used to choose an appropriate LOD for each colliding object at each frame, allowing the application to degrade detection accuracy where it is least likely to affect the user's perception of the collision.

1 Introduction

The demand for highly interactive computer systems is increasing rapidly, as people wish to communicate with computers in a more natural fashion. This need for interaction adds to the pressure on developers of systems to produce realistic, real-time animations. However, this high degree of interaction with the user may also be exploited, in order to adaptively improve the realism of real-time animations. In order to maintain a user's immersion in the animation, a high and constant frame rate is necessary. Realistic rendering, motion synthesis, and collision handling all place an impossibly high computational load on graphics workstations. In a scene with many moving objects, maintaining the target frame rate, while rendering each frame to the highest level of image realism, and controlling motion and interaction of objects to full accuracy is virtually impossible with single-processor work-stations.

In some applications, such as scientific simulation, fully accurate physically-based collision response is necessary. In these cases, solutions must be considered such as degrading the image realism to maintain physically-based response, or increasing the number of processors and developing parallel algorithms for tasks such as rendering and collision detection. However, in a growing number of applications, what is

important is that the viewer perceive events, such as collision response, to be accurate. Testing all potentially colliding objects at full resolution is not always necessary to achieve this goal. In fact, in some cases this testing may be fully redundant, for example, if the viewer is not actually looking at the objects when they collide. The application should adapt at each frame to the perceptual capabilities of the viewer while maintaining a high and constant frame rate (e.g. 10-20 frames per second). In order to achieve this type of adaptation, some form of eye tracking will be an important component of the system.

In Section 2 of this paper, an overview of the proposed approach is given. Section 3 outlines the heuristics to be used, and Section 4 discusses collision testing issues. Section 5 offers conclusions and plans for future work.

2 Overview of approach

An approach is proposed that will maximise the realism of the collision response achieved in a target frame time (e.g. 50 milliseconds), by choosing an LOD representation for each object based on an evaluation of the perceptual importance of the collision in which that object is possibly involved. We propose to build on previous work [O'Sullivan 1996] by taking the full scene and motion complexity into account when grading collisions. Factors which affect the viewer's perception of potential collisions are also allowed for in this approach, as is the current attention focus of the viewer, by tracking their fixation location.

2.1 Previous Work

One solution to the problems of bottlenecks in real-time animation and rendering has been to throw more processing power at the problem. In [Van Reeth et al 1993], networks of transputers are used to implement parallel rendering algorithms. They claim that animation systems can be linearly expanded in performance by adding new processors. However, [Casciola and Morigi 1995] found that, although advantages are gained in the computation phase by parallel processing, display can be a bottleneck. This can sometimes degrade performance to such an extent as to make parallelization a waste of time.

Another possible solution would be to incorporate some intelligence into the system, to enable it to use processing resources more efficiently. The main research into Artificial Intelligence (AI) techniques for collision detection comes from the robotics community, where either neural nets have been trained to predict and/or avoid collisions, or heuristics were developed to do this. (See [Chande et al 1993], [Hiraga et al 1993], [Payeur et al 1994], [Tseng and Wu 1995], [Yuan 1995]).

Adaptive techniques have been proposed, where accuracy of processing is traded against speed. In [Funkhouser and Sequin 1993], a heuristic approach is used to determine the importance of objects, and to choose a suitable level of detail (LOD) to render each object at. A cost heuristic is based on polygon and pixel capacity, and the

benefit is a weighted average of factors such as object size, accuracy and importance. The optimization is a version of the knapsack problem, and is incrementally and approximately solved at each frame.

[Hubbard 1995] also proposes the use of LODs, this time for collision detection. An object is approximated by various levels of sphere-trees, the lowest being a single sphere, which is subsequently refined in higher levels to provide successively tighter approximations to the object's surface. When two objects are deemed to be close to colliding, the application tests for collisions between these LODs, starting at the lowest level, and continuing to test at higher and higher resolutions until the time allotted for collision testing has expired. At this point, the testing process is terminated, and collision response is handled. What is not catered for in this approach is the fact that some collisions are less 'important' than others. No grading of potential collisions takes place. It is, however, recommended that the broad phase should be able to "*selectively ignore objects that the application temporarily designated less important*" (page 229).

By adopting a heuristic-based approach similar to that in [Funkhouser and Sequin 1993], and combining it with an adaptive refinement approach to narrow-phase collision testing, as in [Hubbard 1995], it is possible to grade potential collisions based on their importance, and then to choose a suitable LOD for each object involved in a potential collision, or to trivially ignore or accept 'unimportant' collisions. This will allow for degradation of accuracy where it least affects the perception of the user, allowing the system to put more processing effort into more 'noticeable' collisions.

2.2 Interactive Vision

The objective of all real-time animation applications is to achieve as constant a frame rate as possible, while maintaining high visual realism. In many applications, such as VR, the accuracy of the collision detection is only important insofar as it allows believable collision response.

As an alternative to the orthodoxy they call 'Pure Vision', [Churchland et al 1995] present the idea of 'Interactive Vision'. Pure Vision theorists claim that the human visual system creates a detailed replica of the visual world, works in a strictly hierarchical fashion, and operates independently of other senses. Interactive Vision researchers draw on experimental results that suggest that other sensory systems do play a significant role in what is seen; that the brain is only approximately hierarchical; and most particularly, that the brain does not maintain a fully developed model of the world, but rather assumes the world is constant, and uses it as a type of 'external memory'. They introduce the idea of 'Visual Semiworlds', and claim:

"What we see at any given moment is a partially elaborated representation of the visual scene; only immediately relevant information is explicitly represented. ... Although unattended objects may be represented in some minimal fashion (sufficient

to guide attentional shifts and eye movements, for example) they are not literally seen in the sense of 'visually experienced'." (Page 25).

We use the ideas from this area of research to identify and use some factors which influence a viewer's perception of a collision to some degree, i.e:

- a) Eye movements and direction of gaze
- b) Number of moving/colliding objects in scene
- c) Speed at which objects are travelling
- d) Size of objects
- e) Visibility of intersection points
- f) Input from other sensory organs, e.g. sound, touch.
- g) Semantics, i.e. meaning of objects

A heuristic approach to grading collisions is proposed, similar to that described in [Funkhouser and Sequin 1993], with the benefit heuristic being a weighted combination of the above factors, and the cost heuristic being the time needed to process intersections between multi-level LOD representations. This may be expressed as a version of the Knapsack problem, and approximately optimised at each frame. LODs may in this way be chosen for each potentially colliding object to maximise benefit for a fixed cost (i.e. the target frame rate). It is not necessary to find a global minimum, as local minima are usually very good for this type of problem. As we do not have the problem of smooth progression between successive LODs at each frame which occurs with adaptive rendering techniques, the number of LODs available per object may be kept reasonably low, so that the time to find a minimum will be acceptable. However, it may be the case that the overhead of solving for a minimum may be excessive, in which case a simpler, sub-optimal solution must be considered.

2.3 Eye tracking

When viewing a still picture (or reading text), the eye picks up information in a series of fixations lasting on average about 220 msec and interspersed with rapid jumps called saccades lasting around 20-40 msec. During saccade, it appears that no useful information is acquired by the eye. The issue is more complicated when the eye is viewing one or more moving objects. In this case, it is the time that the eye concentrates on a particular moving object that is important, rather than the actual fixation time. The 'danger' time, from our perspective, is when the eye is distracted from one moving object and moves to view a different part of the screen. If this happens within a frame, after the eye-position has been determined and before rendering occurs, it is possible that the viewer will perceive a collision response anomaly. How likely this is to occur, and how significant its effect would be, is one of the main questions which arises. This can only be answered after fully implementing and testing the proposed system. However, certain speculations and assumptions may be made at this stage.

Let us consider an animation consisting of 1000 frames, produced at a rate of 20 per second, (i.e. 50 milliseconds per frame). We need to predict the average time the eye spends concentrating on one moving object in the animation. It seems reasonable to assume this time to be at least equal to the average fixation time while viewing a still image i.e. 220 milliseconds. We know that saccades last between 20 and 40 milliseconds. Therefore, each fixation-saccade cycle can then be estimated to last 250 milliseconds, i.e. 5 frames. In an animation consisting of 1000 frames, the 'danger time' will then occur 200 times.

On the basis of these tentative calculations, we may predict that there is a 20% likelihood at each frame that the eye will fixate in a new position. The chances that this fixation will occur where it is not expected, i.e. after position has been recorded and before rendering, are less again, while the probability that a collision is occurring at the new point of fixation exactly at that time is significantly smaller. These speculations indicate that the probability of a viewer perceiving a significant collision response anomaly is very low.

Since the turn of the century researchers have been developing systems for determining the exact location of a person's eye during picture viewing or reading. Eye tracking technology has currently reached a level where reasonably non-intrusive instruments can be used to measure fixation location with a useful level of accuracy (e.g., one degree of visual arc). There is, however, a tradeoff between accurate registration of fixation location and intrusiveness of the instrumentation required. For resolutions finer than one degree one usually needs to employ head restraints and bite bars to reduce head and body movement.

In the application being considered in this paper, a high degree of accuracy in measuring eye position will be quite important while we explore the sensitivity of viewers to degradations in collision handling accuracy. The frequency at which the eye-tracking processor is polled to receive the fixation location is also crucial. Ideally, this should occur once at the beginning of each frame, e.g. every 50 milliseconds. The extent to which this affects overall efficiency must be assessed and compared to the benefit gained in collision detection performance.

3 Heuristic approach

We propose adopting an approach similar to that described in [Funkhouser and Sequin 1993], and define a collision tuple $(O1, L1, O2, L2)$ as being a collision of object 1 at LOD 1 and object 2 at LOD 2. We can define two heuristics for these collision tuples:

Cost($O1, L1, O2, L2$) which estimates the time it takes to test for collisions between the objects at those levels of detail, and

Benefit($O1, L1, O2, L2$), which estimates the benefit to the overall perceived realism of the animation of that collision at those LODs.

If we define C to be the set of all collision tuples occurring in a frame, and T to be the target frame time, then we can choose a level of detail for each colliding object in that frame by:

$$\text{Maximising: } \sum_C \mathbf{Benefit}(O1,L1,O2,L2)$$

$$\text{subject to: } \sum_C \mathbf{Cost}(O1,L1,O2,L2) \leq T$$

The cost heuristic is an estimate of the time required to test both objects at those LODs for intersection. This depends on the intersection testing algorithm used, but is usually directly proportional to the number of vertices, and should be a conservative estimate. A collision testing algorithm is needed for which the processing time can be accurately estimated in advance, i.e. no worst or best time. The benefit heuristic should be based on Interactive Vision ideas, and should be a weighted sum of the factors listed above. An attempt to quantify the effect of these factors follows:

3.1 Eye movements and direction of gaze

In [McConkie 1990], an experiment is described, where viewers were asked to gaze at an image on a screen. As they allowed their eyes to roam over the display, foveating it to absorb all its detail, their eye movements were tracked. Major display changes were made during saccades which went completely unnoticed. Objects were added or removed, or their colours altered, while the viewer retained the impression that they were viewing a completely unchanging picture.

As yet, research on human visual collision detection has not involved eye movement recording. A key question is the degree to which motion discontinuities are detected at different regions of the retina, partially with what range of motion parameters are acceptable as indicating normal response from collisions, and partially with the degree to which motion discontinuities at different retinal regions grab attention. Observations have shown that collisions occurring on the periphery of subject's foveal region are perceptible at a very low resolution. Objects could be repulsed at quite a significant distance, but still be perceived by the subjects as having collided.

[Funkhouser and Sequin 1993] do not track eye movements in their adaptive display algorithm, but simply reduce the benefit of each object by an amount proportional to its distance from the middle of the screen. If an eye-tracker is used, it may detect if the viewer's eye is moving, the (x,y) location of the point of focus if the eye has rested, and the distance of the eye from the point of focus. As the above experiment has shown, as the eye is moving, major changes will go unnoticed, so the benefit of each collision may be calculated as being equal. Otherwise, the midpoint between the centres of two potentially colliding objects is calculated, and the benefit of this

collision will be reduced by an amount proportional to the distance of this midpoint from the point of focus.

3.2 Number of moving/colliding objects

In [Verghese and Pelli 1992], an experiment called “find the dead fly” was devised in which subjects were first presented with a large number of moving spots, and one stationary one, and secondly with a large number of stationary spots, and one moving one. They found that increasing the number of moving objects in the former case reduced the ability of the observer to detect the stationary one by an inverse proportion, whereas in the latter case, the ability to pick out the moving object was independent of the number of stationary ones. Hence, although the time available to process each collision (and hence accuracy of response) is reduced depending on the number of colliding objects, the attention of the viewer to individual collisions is also reduced, leading to an overall decreased ability to notice response anomalies.

3.3 Speed at which objects are moving

Objects which approach each other and bounce off at high speed will appear blurred, as the viewer’s eye and brain cannot process the motion in time. Hence, the benefit of a collision can be reduced by an amount proportional to the sum of the perceived relative speed of both objects involved in it.

3.4 Size of objects

The size on screen of colliding objects, measured by pixel coverage, also contributes to how noticeable a collision is. Hence, the benefit of a collision can be increased by an amount proportional to the sum of the screen sizes of the objects involved in it.

3.5 Visibility of intersection points

Further observations have shown that even if a subject is looking directly at colliding objects, if the potential points of collision are obscured, they cannot tell if the colliding objects have actually touched or not. This is particularly true for objects which approach each other along paths almost coinciding with the users direction of gaze. This is the case for quite significant gaps of repulsion. However, if the collision points are on the visible silhouette of both objects, and the objects are approaching each other along paths almost perpendicular to the users direction of gaze, the subject’s perception of the collision is extremely acute. This indicates that the benefit of a collision should be reduced by an amount proportional to the angle that the vector between the centres of the objects makes with the view-plane. It also implies the need for very accurate collision tests on the visible silhouettes of objects, less accurate tests on visible interior faces, and very approximate tests on occluded back faces.

3.6 Input from other sensory organs, e.g. sound, touch

[Churchland et al 1995] describe subjective motion experiments which show that adding auditory stimuli affects the way in which subjects 'see' events. First they showed a blinking dot and a shaded square, and subjects perceived no motion. Then they added a tone to the left ear whenever the dot blinked on, and to the right ear when it blinked off. This led to subjects perceiving the dot as moving in and out behind the shaded square.

Undergraduate students in our laboratory have incorporated sound and colour changes (e.g. colliding objects flashing bright red and beeping) into a very basic collision response handling system for their animation projects. This increases the realism of collisions, convincing the viewer that a collision has occurred, even when the collision detection is not completely accurate. This need not, however, be taken into account as part of the benefit heuristic, but simply noted as a useful device to increase the realism of a collision

3.7 Semantics

[Churchland et al 1995] introduce the concept that animals and humans have a 'relevant to my lifestyle' model of the world, as opposed to a 'world with all its perceptual possibilities'. They claim that the purpose of visual perception is to: "*facilitate the organisms thriving at the four Fs: feeding, fighting, fleeing and reproduction*" (Page 25). Hence, it could be claimed that people only 'see' things in detail which are relevant to the task that they are currently involved in.

Some objects in the animation may have a more important "meaning" to the user than others. For example, in an interactive football game, the ball will be the most important object, and collisions between it, the players, and the goal-posts will be much more attended than say collisions between the players themselves. These 'importance ratings' can be set by the application, and the benefit of each collision should be increased by an amount proportional to the sum of the importance of both colliding objects.

4 Collision testing

Hybrid collision detection is a term used by [Kitamura et al 1994] to refer to collision detection methods which first perform approximate tests to identify interfering objects in the entire workspace and then perform more accurate tests to identify the object parts causing interference. [Hubbard 1995] and [Cohen et al 1995] both propose hybrid algorithms for collision detection. Hubbard calls the two phases of the algorithm the "broad phase", where approximate intersections are detected, and the "narrow phase", where exact collision detection is performed. [Palmer and Grimsdale 1995] proposes an algorithm with three stages, the first stage using bounding

volumes, the second stage using sphere-trees to approximate an object's geometry, and the third stage performing very accurate polygon intersection tests. As in previous work [O'Sullivan 1996] a Hybrid Algorithm is proposed, which first runs a Broad Phase, testing for approximate collisions between all pairs of objects, grades them by importance, and then triggers a Narrow Phase, performing more accurate tests only on pairs of objects between which approximate collisions have been detected.

The strategy of projecting higher dimensional objects onto a lower dimension, and testing for intersections between these projections has been used by [Cohen et al 1995] , [Shinya and Fergie 1995], [Thalmann and Volino 1996] and [O'Sullivan 1996] among others. This approach is advantageous for the following reasons:

- Dimension reduction: the dimension in which manipulations are performed is reduced by 1, making it simpler to implement
- Robustness: because of the limited use of topological information, i.e. not many different special cases.

The down-side of using dimension-reduction techniques is loss of information, especially if used in the narrow phase. However, the purpose of this research is not to develop a highly accurate collision detection routine, which always produces the actual points of collision. If two objects are so close to each other as to look to the viewer as if they were colliding, this can usually be taken to be a collision. A dimension-reduction approach is proposed for both phases of the collision detection algorithm: A Broad Phase which culminates in the "grading" of the approximate collisions detected, and the selection of LODs for each object involved in these collisions, and a narrow phase which tests for intersection between these objects at the chosen LODs. The narrow phase should be such that it is possible to accurately predict the time needed to perform testing.

Trading speed for accuracy will give rise to anomalies which will be of one of the following types: 1) Missed collisions, i.e. collisions which occur and go undetected, or 2) non-existent collisions, i.e. collisions which have not occurred but are nevertheless detected. Either one or the other anomaly must be accepted - but which one? Which is most noticeable, or most disturbing for the user?

The choice of accepted anomaly determines the type of LOD representation to be used for objects. If anomalies of type 1 are to be avoided, i.e. no collisions should go undetected, a superset of all collisions must be detected. In this case, each LOD representation of an object should totally enclose the object. Hence LOD generation should start with the lowest LOD being a loose boundary of the object, which is refined to produce closer and closer approximations to the boundary of the object. The higher the level of detail, the less likely that non-existent collisions will be detected. Sphere-trees [Hubbard 1995] provide this type of conservative approximation, as do Oriented Bounding Box (OBB) trees, described by [Gottschalk et al 1996].

If, however, collisions which have not occurred should never be mistakenly detected, i.e. type 2 anomalies should be avoided, only a subset of all collisions can be detected. Therefore, LOD generation should start with the most accurate representation of the object, and lose detail gradually, preserving local features, but losing more of the global nature of the object. The higher the level of detail, the lower the chance of missing a collision.

5 Conclusions and Future work

An approach has been described in this paper, which proposes an adaptive collision detection algorithm that allows degradation of detection accuracy where it least affects the perception of the viewer. In this way, maximum realism should be attained within a target frame time. The proposed approach could be applied not only to the collision detection problem, but to any kind of multiresolution problem in real-time animation systems.

The major issues involved in the proposed approach have been discussed, and some informal investigations into its feasibility have been described. We now intend to rigorously test the perceptual criteria proposed, in order to assess their impact. The eye tracker described in section 2 will be central to these investigations. This is necessary to enable us to quantify each of the factors discussed in section 3, in order to schedule the accuracy of the collision test for each object pair. Once these factors have been quantified, implementation and testing may occur, and the benefit of the proposed approach may be measured. At the moment, the proposed technique is valid for only one viewer. Whether the ideas proposed can be adapted to cater for multiple viewers is another area worthy of further investigation.

Hardware constraints must also be considered. Many forms of eye-tracking equipment are quite intrusive. The viewer's head is kept immobile with the use of supports and bite bars (not unlike some mediaeval instruments of torture). However, we are not suggesting that this be adopted as the method of eye-tracking in interactive animation systems such as VR. What we do wish to establish is that eye-movement tracking can be used to improve the perception of computer-generated animations. Once the principle has been established, more non-intrusive methods can be considered, as can issues relating to the compatibility of eye-trackers and other viewing systems such as Head Mounted Displays (HMD).

References:

[Casciola and Morigi 1995] G. Casciola, S. Morigi. "Graphics in parallel computation for rendering 3D modelled scenes." *In Parallel Computing* Vol.21 Pages 1365-82, Aug.1995.

[Chande et al 1993] P.K.Chande, M.Shrivastava, G.N.Sharma. "Neural Assisted Robot Arm Collision Avoidance." in *SICE 93 Proc. of the 32nd SICE Ann. Conf. Inter. Sess.* Kanusawa, Japan. Pages 1547-1550, Aug. 1993.

[Churchland et al 1995] P.S.Churchland, V.S.Ramachandran, T.j.Sejnowski. "A Critique of Pure Vision" in eds.C.Koch, J.L.Davis *Large-Scale Neuronal Theories of the Brain*, MIT press, 1995.

[Cohen et al 1995] J.D.Cohen, M.C.Lin, D. Manocha, M.K.Ponamgi, "I-COLLIDE: An Interactive and Exact Collision Detection System for Large-Scale Environments" in *Proc. 1995 Symp. Interactive 3D graphics*, Monterey, Calif. pages 189-196, 1995.

[Funkhouser and Sequin 1993] T.A.Funkhouser, C.H.Sequin. "Adaptive Display Algorithm for Interactive Frame Rates During Visualization of Complex Virtual Environments." in *ACM SIGGRAPH 93*, Anaheim, Calif. published as *Computer Graphics, Annual Conference Series*, 1993. Pages 247-254. Aug. 1993.

[Gottschalk et al 1996] S.Gottschalk, M.C.Lin, D.Manocha. "OBB-Tree: A Hierarchical Structure for Rapid Interference Detection." To appear in *Proc. of ACM Siggraph '96*. 1996.

[Hiraga et al 1993] I.Hiraga, T.Furuhashi, Y. Uchikawa, S.Nakayama. "Knowledge Acquisition for Collision Avoidance Using Fuzzy Neural Networks". in *IJCNN 93 Proc. of the 1993 Int. Joint Conf. on Neural Networks*, Nayoya, Japan. Vol.1. Pages 673-6. Oct. 1993.

[Hubbard 1995] Philip M.Hubbard. "Collision Detection for Interactive graphics applications." In *IEEE Transactions on Visualization and Computer Graphics*, Vol. 1, No. 3 pages 218-230

[McConkie 1990] G.W.McConkie "Where vision and cognition meet". *Paper presented at the HFSP Workshop on Object and Scene Perception*, Leuven, Belgium. 1990.

[Kitamura et al 1994] Y.Kitamura, H.Takemura, N.Ahuja, F.Kishino. "Efficient Collision Detection Among Objects in Arbitrary Motion Using Multiple Shape Representations". in *Proc.12th IAPR Int. Conf. on Pattern Recognition*, Jerusalem. Vol.1. Pages 390-6. Dec.1994.

[O'Sullivan 1996] C.A.O'Sullivan. "Real-time Collision Detection for Computer Graphics." *M.Sc dissertation*. Dublin City University. Oct.1996.

[Palmer and Grimsdale 1995] I.J.Palmer, R.L.Grimdale. "Collision Detection for Animation using Sphere-Trees." in *Computer Graphics Forum*. Vol.14. Pages 105-16. June 1995.

[Payeur et al 1994] P.Payeur, H.Le-Huy, C.Gosselin. "Robot Path Planning Using Neural Networks and Fuzzy Logic." in *IECON 94 Int. Conf. on Ind. Electronics, Control and Inst.* Bologna. Vol.2. Pages 800-5. Sept. 1994.

[Shinya and Fergie 1995] M.Shinya, M.Fergie. "Laying out objects with geometric and physical constraints." in *Visual Computer*. Vol.11, pages 188-201, 1995

[Thalmann and Volino 1996] N.M.Thalmann and P.Volino. "Sculpting, Clothing and Hairdressing our Virtual Humans" in eds N.M.Thalmann, D.Thalmann *Interactive Computer Animation* Prentice Hall, 1996.

[Tseng and Wu 1995] C.Tseng, C.Wu. "Collision Detection for Multiple Robot Manipulators by using Orthogonal Neural Networks." in *Journal of Robotic Systems*, Vol.12 Pages 479-90.1995.

[Van-Reeth et al 1993] F.Van Reeth, E. Flerackers, W. Lamotte. "Animating architectural scenes utilizing parallel processing." in *Visualization and Intelligent Design in Engineering and Architecture*, Southampton. Pages 149-64, Apr. 1996.

[Verghese and Pelli 1992] Verghese,P. and Pelli,D.G. "The information capacity of visual attention", in *Vision Research*, Vol 32, Pages 983-95, 1992.

[Yuan 1995] J.Yuan. "A Neural Network Measuring the Intersection of m-dimensional Convex Polyhedra." in *Automatica*. Vol.31. Pages 517-29. April 1995.