

Collisions and Perception

Carol O'Sullivan and John Dingliana
Image Synthesis Group, Trinity College Dublin

Level of Detail (LOD) techniques for real-time rendering and related perceptual issues have received a lot of attention in recent years. Researchers have also begun to look at the issue of perceptually adaptive techniques for plausible physical simulations. In this paper, we are particularly interested in the problem of realistic collision simulation in scenes where large numbers of objects are colliding and processing must occur in real-time. An interruptible and therefore degradable collision handling mechanism is used and the perceptual impact of this degradation is explored. We look for ways in which we can optimise the realism of such simulations and describe a series of psychophysical experiments that investigated different factors affecting collision perception, including eccentricity, separation, distractors, causality and accuracy of physical response. Finally, strategies for incorporating these factors into a perceptually adaptive real-time simulation of large numbers of visually similar objects are presented.

Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*Animation*; H.1.2 [Models and Principles]: User Interfaces—*Human Factors*

General Terms: Algorithms, Experimentation, Human factors, Measurement

Additional Key Words and Phrases: Animation, Collision Handling, Graphics and Perception, Simulation Levels of Detail, Time-critical computing

1. INTRODUCTION

To express solidity in a simulated virtual environment, objects need to respond at the right time and in the right manner to collisions with other objects in a scene. An efficient and realistic collision handling mechanism is fundamental to any physically plausible animation system. For example, consider the problem of a large number of rocks tumbling down the side of a mountain, each crashing against the side and with the other rocks, as implemented by Mirtich [2000]. A large amount of contact information and dynamics needs to be processed to accurately simulate this scenario. Unfortunately, the problems of detecting when objects collide, modelling the contacts between them and subsequently determining an appropriate response, are all inherently expensive operations. With increasing numbers and complexity of objects in the scene, collision handling can quickly become a major bottleneck. Mirtich found that it took 97 seconds on average to compute each frame of his avalanche simulation on an SGI Onyx (200 MHz R10K CPU), not because of any shortcomings in his simulation algorithm, but because of the high number and

Author's address: Carol O'Sullivan and John Dingliana, Image Synthesis Group, Trinity College Dublin, Ireland.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 200 ACM 0730-0301/200/0100-0111 \$5.00

complexity of contact groups formed. His implementation favoured robustness over efficiency and for some applications this may be necessary. However, for real-time animations such a sacrifice is not an option and it is obvious that a trade-off between detection accuracy and speed is necessary to achieve a high and constant frame-rate. We now ask: what effect will this have on the viewer's perception of the resulting physics?

Having established that it will become necessary to sacrifice detail in order to meet real-time demands for most scenes of significant complexity, a number of strategies can be adopted to adaptively modulate levels of detail in different parts of the simulation process. These range from culling dynamics in non-visible parts of the scene: [Chenney and Forsyth 1997], using simplified rules for less important parts: [Carlson and Hodgins 1997], considering the concept of physical plausibility: [Barzel et al. 1996; Chenney and Forsyth 2000], or varying the complexity or resolution of different regions in terms of geometric or visual detail: [Reddy 1997; Funkhouser and Séquin 1993]. Such approaches allow the speed-accuracy trade-off to be optimised by exploiting uncertainty or the user's inability to distinguish simplifications in less obvious parts of the scene. Adaptive detail modulation wherever possible is preferable to the limiting alternative of pre-emptive simplification, i.e. reducing the complexity of the whole scene to achieve a target frame rate, regardless of any other factors. However, techniques need to be established for achieving this in an efficient manner for dynamical simulations such as the avalanche scenario described above.

2. COLLISION PERCEPTION

The aim of Classical, or Newtonian, mechanics is to describe the motion and interactions of objects within the physical world. There is also evidence to show that the human perceptual system relies heavily on representations of dynamic properties of the world. For example, Baillargeon et al. [1985] showed that even young babies of 5 months old could already tell that one solid object cannot occupy the space of another. However, these human naïve representations are not always compatible with the accurate physical models of Newton and are in fact closer to medieval concepts of force and impetus. In a study of physics students taking a university course, Clement [1982] reported that most had intuitive preconceptions concerning mechanical events that, although incorrect according to Newtonian mechanics, were highly stable and widespread. Proffitt and Gildea [1989] showed that people use only one dimension of information when making dynamical judgements. Therefore, when a dynamic event involves more than one dimension of information such as velocity and rotation, i.e. an extended body motion as opposed to a particle that has only one dimension of information, humans are less able to correctly identify anomalous physical behaviour. Can we exploit this imprecision of the human brain for the purpose of producing plausible real-time physics? Gildea and Proffitt [1989] also conducted an experiment that investigated people's ability to judge the relative mass of two objects involved in a collision. They discovered that judgements were made based on heuristics and that people used kinematic data, such as velocity after impact and the way that the colliding objects ricocheted, to make their decisions. Knowing this, we must ask ourselves if it is necessary, or even wise, to



A moving ball hits a stationary one that moves immediately or after a delay. Participants perceived causality in the first but not the second condition.

Two balls move from left to right or vice versa. The objects are perceived to be animate, i.e. participants stated such things as “The big ball is chasing the little one.” or “The little ball is following the big one.”

Fig. 1. Michotte’s Causality Experiments

expend a large amount of processing time on trying to replicate real-world physics.

Causality refers to the ability to detect whether one event causes another. For example, a collision of a moving object with a stationary one will cause the second object to move, whereas a stationary object that starts to move by itself is perceived to be autonomous. The human brain appears to have very low-level processes that allow people to distinguish between animate and inanimate objects. Therefore, the concept of causality occupied philosophers such as Aristotle, Galileo and later Hume, but modern research on the perception of causality was sparked by the work of Albert Michotte [1963]. He showed his participants very simple primitives moving on a screen and asked them to explain what they perceived to be happening (See Figure 1). His work clearly showed that, even with such abstract objects, the impression that participants had of one event causing another or not was extremely strong. In a related study, Leslie and Keeble [1987] illustrated that six-month-old children could distinguish between causal and non-causal scenarios, showing that this ability is established very early on in life. It seems that prompt processing of collisions is therefore necessary in order to maintain the perception of causality for collision events.

As we will see in the following section, one of the possible consequences of reducing the accuracy of collision detection will be objects that bounce off each other at a distance. The extent to which this separation, or gap, between two colliding objects is perceivable will be an important factor in determining the ability of humans to detect an anomalous collision. This is because there is a topology-preserving mapping from the cells in the retina to the cells in the primary visual cortex, called a retinotopic mapping, that is quite precise and enables spatial location information to be efficiently processed: [Tootell et al. 1982]. Research into the phenomenon of apparent motion may be relevant when trying to establish such thresholds: Two spatially separated stimuli, flashed in a temporal sequence, are perceived to appear either simultaneously or successively, or to be actually only one stimulus in motion, depending on the time delay and the distance between them. The metrics

D_{max} and T_{max} are used to describe the maximum spatial and temporal spacing for which such apparent motion is perceived: [Braddick 1974].

The fact has long been established in vision literature that many visual processing tasks deteriorate at increasing eccentricities from the fixation point: [Aubert and Förster 1857; Weymouth 1958]. A summary of the physiological reasons for decreased spatial resolution in the periphery appears in [DeValois and DeValois 1988]. In particular, there is more representation for the fovea than for the periphery in the visual cortex, allowing for more acute visual processing. A cortical magnification factor (M) can be used to quantify the rate of decline in the periphery, or indeed to eliminate it by magnifying stimuli, a process called M -scaling. See [Rovamo and Virsu 1979] and [Carrasco and Frieder 1997] for further details.

This eccentricity effect can be exploited in a real-time application by tracking the user's fixation position. Gaze-directed adaptive rendering has been investigated: [Ohshima et al. 1996; Watson et al. 1997; Luebke et al. 2000] and similar strategies can also be applied to simulation. When the viewer is looking directly at a collision, it would be given a higher priority than a collision occurring at a slight eccentricity, which itself would receive a higher priority than other collisions presented more peripherally. However, it may be that an eccentricity metric alone is not sufficient to guide adaptive processing in all cases. Aubert and Förster [1857] showed that M -scaling equalized performance of two-point separation in the near periphery but, in the same study, the theory failed completely for two-point separation in the far periphery. It also fails in many other cases, as illustrated in [Strasburger et al. 1994], where both size and contrast had to be increased to neutralise the eccentricity effect.

The presence of distractors is also likely to affect the ability to accurately detect collision or non-collision, as is the type of distractor. These issues arise in the area of Visual Search and have been investigated by researchers such as Saarinen [1994] and Treisman [1982]. In such studies, a display usually includes a stimulus that participants must look for, called the target, while other stimuli are also added to attempt to distract them from their task, i.e. distractors. If the distractors are in a clearly distinguishable perceptual grouping from the target, the identification of this grouping occurs automatically, without any attention or search being necessary. Such a grouping may be of similar colour, orientation or common movement, which differs from the target. This is referred to as a preattentive pop-out task and increasing the number of distractors present does not affect the outcome of the search. If such an obvious grouping is not immediately apparent, it is necessary to focus attention on each item in turn and this is called a serial search task. In such a task, performance is significantly worse than in the pop-out tasks and increasing the number of distractors has a negative effect. In the tasks that we are considering, i.e. simulations of large numbers of visually homogeneous interacting entities, there usually would not be obvious and disparate perceptual groupings of objects. Therefore, the ability of viewers to detect a collision anomaly in such a scenario is of major interest to us.




Fig. 2. Different levels of Collision Detail using a hierarchical sphere tree

3. INTERRUPTIBLE COLLISION DETECTION

We use a collision detection mechanism based on the approach by Hubbard [1995] which checks for intersections between nodes in hierarchical data structures, in this case sphere-trees, that are used to approximate the objects' volumes at different levels of detail. This is described in more detail in [O'Sullivan et al. 1999] and [Dingliana and O'Sullivan 2000]. Processing begins when the volume representations at the coarsest level of detail have been found to be colliding. Based on this information, an initial approximation of the contact points can be made. The next phase achieves progressive refinement of collision data by traversing the sphere hierarchy of each colliding object to get a better approximation of the contact points. The children of colliding nodes that have been found to be colliding at the coarser level of detail are checked for collisions (See Figure 2). This process is repeated until we reach the leaf nodes, at which point we can return the highest resolution approximation of the contact points based on our volume representation. In a time-critical approach this progressive refinement phase is halted when the system runs out of time allocated for the collision processing task and the approximation of contact points is made based on the highest level of detail that has been processed.

Contact data are used to calculate collision response based on an instantaneous impulse model: [Witkin et al. 2001; Mirtich 1996]. As a result, the calculated collision response accuracy is directly dependent on the level of collision detection detail. Therefore, a lower level of detail will result in less physically accurate physics and objects not touching when they bounce, thus leaving a possibly perceivable gap. Just how perceptible these anomalies are is the subject of the remainder of this paper.

Even at this very fundamental level, the approach has the advantage of being able to guarantee real-time performance in cases where it would not normally be possible without pre-emptive simplification of the scene. However, by careful study of the different factors that cause degradation in the believability of the simulation, we can make strategic simplifications in less noticeable areas to better optimise the speed-accuracy trade-off. To achieve this, we assign some measure of priority to different collisions in the scene. More important collisions are then allocated more processing, i.e. sphere-tree traversal is applied to a deeper level, while the accuracy of less important collisions is sacrificed in order to save on processing time and meet the target frame-rate. This approach is similar to that taken by Funkhouser

and Séquin [1993], who used factors such as size and speed of objects to choose the levels of detail at which to render objects in a scene.

If we can give each collision some measure of importance, the next logical step would be to sort all collisions and apply increasing levels of refinement to each in order. However, in practice it was found that the overhead of a full sorting process was computationally very expensive, thus reducing the gains that were being made from not using traditional exact collision detection techniques. A more fruitful approach is to partition the set of collisions into discrete subsets based on the prioritisation criteria. A similar approach was taken by Duchaineau et al. [1997], who maintained priority queues to select which regions of a mesh to simplify while rendering terrain in real-time. In our system, collision processing is first applied at a low attainable level to all collisions and then collisions in the higher priority subsets are refined to whatever level is possible in the allocated time remaining. Even with just two partitions this approach has yielded significant gains. In our implementation, this was achieved by storing collision data in two separate priority lists. All collision events are represented as collision data structures in either one of the two lists based on some importance criteria. Collisions in the high-priority list are allocated more processing time so that the contact model and resulting response is more believable.

4. PSYCHOPHYSICAL EXPERIMENTS

In computer graphics research, it is often necessary to design and execute psychophysical experiments in order to investigate some of the specific problems raised: [Hodgins et al. 1998; Watson et al. 2001]. We carried out several sets of experiments to determine the extent to which the effects reported in the literature are applicable to the particular collision handling scenarios being considered. These were designed to examine the effects of factors such as eccentricity, separation, presence and number of similar and dissimilar distractors, causality and physics on participants' perception of collisions.

4.1 Eccentricity

Twenty participants took part in an experiment to determine whether the ability to detect anomalous collisions, in this case objects that do not touch each other but leave a gap, decreases with increasing eccentricity of the collision point. Green spheres served as stimuli with a red sphere displayed in the centre of the screen for fixation purposes (see Figure 3). Collisions were presented at 5 eccentricities of 1.4, 2.9, 4.3, 5.7 and 7.2 degrees of visual angle. There were 8 replications at each eccentricity, each in a different screen location. The order of presentation was also randomised. The spheres appeared and moved towards each other for 1 second, and then moved apart for 1 second after either touching, leaving a small gap of 4 millimetres (mm), or a larger gap of 8 mm. The orientation of the objects with respect to each other was randomised to be horizontal, vertical or diagonal (as shown).

Participants were required to indicate after each collision whether the objects had touched or not. If they were unsure, they were told to choose the touch option. They were instructed to fixate on the red sphere at all times and not to look directly at the green stimuli. Although we did not validate whether participants actually




Fig. 3. Stimuli for the eccentricity experiments

maintained this fixation or not, from the results we can backwards-infer that they did, as this is the only physically plausible explanation for the observed behaviour. They first performed a trial run of the experiments to familiarise themselves with the methods and stimuli. We can see from Figure 4 that there was a fall-off in detection accuracy with eccentricity that was statistically significant: > 99% in both cases. It is also evident that separation of the stimuli affected performance, with the smaller gap eliciting fewer correct responses and a faster fall-off with eccentricity than the larger.

4.2 Distractors

To evaluate the effect of distractors on people's ability to detect anomalous collision events, we ran two experiments with 12 participants. This time, the stimuli were pairs of white disks presented on a black background that collided at 3 different eccentricities: 2.7, 5.4, and 8.1 degrees of visual angle. At each eccentricity there were runs at 8 locations on the screen and every combination of variables was replicated 3 times. Direction of motion was randomised for each run and the order of runs was also random. Again, participants were instructed to fixate on a target in the centre of the screen. This time we ensured that fixation was maintained by only allowing the collisions to occur at a time when it was impossible for the subjects




Fig. 4. Average performance for gapsizes vs eccentricity

to have generated an eye-movement, which takes at least 250 milliseconds (msec) to execute. At each run, two stimuli appeared and moved towards each other at a velocity of 40 mm per second. After 150 msec they reversed direction after either touching or leaving a gap of 2 mm or 5 mm. There was a 2-second delay between successive presentations of each set of stimuli, during which a blank screen was displayed for 1 second, then the fixation target only for 1 second. Participants were instructed to hit the left mouse button if they felt that the stimuli had touched, or if they were not sure. They should only hit the right mouse button if they were sure that they had perceived a gap between the two stimuli. There was a trial run of each experiment, to familiarise them with the methods and stimuli.

In the first experiment, 1, 5 or 9 bright red circles appeared simultaneously with the white stimuli, each moving along a random linear trajectory, with a random velocity (see Figure 5). These distractors were not allowed to occlude the colliding stimuli. The second experiment was similar, except that in this case randomly moving white circles that were identical in appearance to the colliding stimuli were added (see Figure 6). This task was particularly difficult, because it was not always easy to detect which were the colliding stimuli. Participants were therefore given the added instruction that if they did not identify any circles as moving towards each other and then separating, they should respond as in the colliding or not-sure cases, i.e. by clicking the left mouse button. Only if they were certain that they saw two circles approach and separate without touching, should they indicate a gap by clicking the right mouse button. In this experiment we also collected information at the point of fixation. These data were not considered in comparisons with the first task and were only included to see what effect the complexity of this task had on performance at the location of fixation. Participants were told as before to fixate on a target in the centre of the screen, which was hidden when the colliding stimuli and distractors were present and re-appeared immediately afterwards.




Fig. 5. Experiments with visually dissimilar distractors




Fig. 6. Experiments with visually homogeneous distractors

We can see from Figure 7 that the addition of the visually homogeneous distractors did impact more severely on participants' ability to correctly detect a gap than visually different ones. The effect of separation was again significant. The slight increase in detection ability in the presence of different distractors, although interesting, is not statistically significant: $< 50\%$ and $< 30\%$ for the 5mm and 2mm gaps respectively. However, the decrease in accuracy in the presence of visually homogeneous distractors is significant, more so in the case of the larger than the smaller gap: $> 95\%$ and 70% respectively. These findings are consistent with those from the field of Visual Search, discussed in Section 2. Performance was worse in the presence of similar distractors and increasing the number of such distractors had a negative effect, whereas increasing the number of different distractors did not cause any deterioration. We also examined the effect of eccentricity and found it to be statistically significant in both cases: $> 99\%$. Figure 8 shows the eccentricity effect by gap size for the experiment with similar distractors, with performance at the fixation point included. This is the situation that is closest to the real-world scenario we are interested in, i.e. simulations involving large numbers of similar objects. The data could quite easily be approximated by a mathematical function of gap size and eccentricity that exhibits an exponential fall-off, as discussed in [O'Sullivan et al. 1999].

4.3 Causality

Sudden drops in frame rate during collision events can have serious effects on the perception of an animation. As Michotte [1963] demonstrated in his experiments, objects moving apart after a certain delay are no longer perceived to be doing so as a result of the collision. Unfortunately, for non-interruptible collision handling systems, delays are one of the more likely side effects when there is a sudden increase in computational workload, such as when a large number of objects happen to

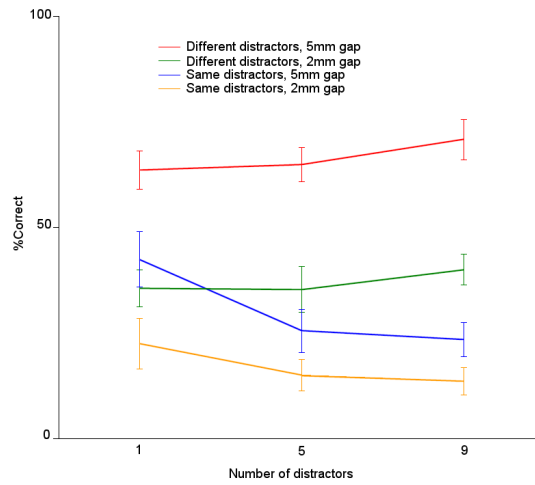


Fig. 7. Effect of same vs different distractors




Fig. 8. Performance in the presence of similar distractors.

collide simultaneously.

To evaluate the effect of such delays, an experiment was set up to extend Michotte's studies for animated 3D scenes. Twenty participants were shown animations of collisions involving very simple objects: two spheres of equal volume and mass colliding and moving apart (see Figure 9). The collisions all occurred in the centre of the screen. In each simulation run the relative initial and final velocities of the objects were equal and opposite but delays of 0 msec, 100 msec and 300 msec were artificially introduced at the instant the objects collided. For each delay




Fig. 9. Causality Experiment in 3D




Fig. 10. Effect of an induced delay on the rating of collisions (low scores are better)

time the simulations were run for 3 different initial velocities and from 3 different viewpoints, with 3 replications of each condition. Participants were asked to rate each collision on a simple integer scale of 1-3, where 1 was very believable and 3 was very unbelievable. They were subsequently questioned about their strategy.

We found that the effect of delay upon the rating given to collisions was highly significant: > 99%, as can be seen in Figure 10. Please note that low scores are better. We examined the data collapsed over viewpoint and velocity and found no significant effect. The delay effect was also significant when examined for each velocity separately. In addition, almost all participants complained about the collisions that seemed to “stick together” or were “less bouncy”. These results are consistent with Michotte’s, in that the addition of a delay reduced the perception of causality, thus impacting negatively upon collision realism. We can therefore conclude that the longer a real-time system spends processing collisions and the longer the delay that is thus generated, the less believable the resulting collisions will be.




Fig. 11. L-shaped bodies used for dynamics/kinematics experiments

4.4 Dynamics and Kinematics

Directly following the causality experiments, two experiments were performed with the same 20 participants to test their sensitivity to variations in the levels of collision detail. The first experiment tested the impact of varying levels of the volume model used for collision detection. As described in Section 3, lower levels of sphere tree detail result in larger gap-sizes in collisions. Participants were shown collisions involving three-dimensional L-shapes (see Figure 11). These simple concave extended bodies were chosen as they are useful for testing the features of the collision detection system. The starting velocities and orientations were set so that the resulting collision at the centre of the screen caused both objects to be repelled back in directions directly opposite to their initial velocities. Users were again asked to rate the collisions on a simple scale of 1 to 3 and were questioned on their strategy afterwards. The simulations were run at varying levels of sphere-tree detail from 3 different points of view and with the objects' velocities being scaled by factors of 50%, 100% and 150%, with 3 replications of each condition.

The second experiment tested participants' responses to the model for collision response used in the system. As discussed in Section 3, varying the levels of collision detection detail has a direct effect on the objects' final velocities after collision. Participants were shown collisions involving the same L-shaped bodies used in the previous experiment. This time however the objects started off with different initial conditions with the resulting collision causing both a change in trajectory and angular momentum. The main goal of this test was to evaluate how adversely the collision levels of detail affected the perceived collision response. The lowest level of collision detail involves approximating the objects as spheres. As a consequence of this, the calculated collision impulses tend not to impart changes to the angular velocity of the objects. Because such response looks so unbelievable, as one expects rotations in collisions involving such non-symmetric objects to result in spin (an observation later verified by user feedback), a random change was added to the objects' angular velocities at the lowest collision level. This change was uniformly distributed between two limits that were chosen by inspecting the values generated for a series of accurate collisions.

One hypothesis we tested was whether the viewpoint affected the perception of the collision anomalies. There were three cases, the first being the 0° case, where objects approached each other parallel to the viewing plane. The other cases involved angles of 45° and just less than 90° respectively. In the latter condition,

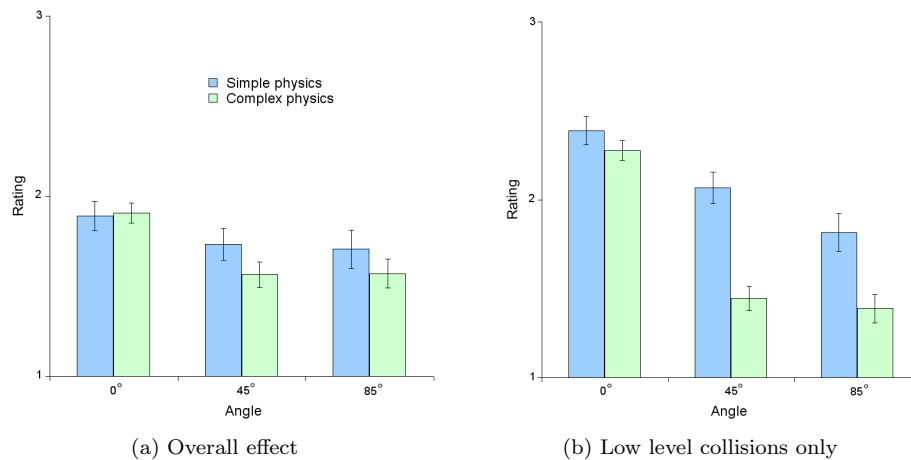


Fig. 12. Effect of viewpoint on rating of collisions

although the viewers could not see the actual points of impact, they could see the object as it approached and bounced off, thus providing them with enough information to choose a rating. In the simple response case, there was a weakly significant effect: $> 75\%$, with a more significant effect in the second experiment with more complex physics: $> 99\%$ (see Figure 12a). In particular we were interested if the obscured viewpoints reduced the number of gaps perceived. We found that the main effect was with the largest gap size only, i.e. lowest level collision handling, where the effect was highly significant for both experiments but particularly strong for the more complex physics (see Figure 12b).

We also predicted that velocity would reduce a viewer's ability to detect anomalous collisions. We looked at the overall effect for both experiments and discovered that there was only a weak effect in the simple case, where perception was worst for the slower collisions, but there was an opposite and more significant effect in the complex physics case, where faster collisions were actually rated worse. These unexpected results led us to examine the effect of velocity at each level of detail separately. We found that at the slowest velocity, the effect of collision resolution was most significant in both cases, with the effect being most significant in the experiment with simple physics. In this case, performance was more as expected (see Figure 13a), with low resolution collisions being less detectable with increasing velocity.

The results in the experiment with complex physics were more surprising (see Figure 13b). When velocity was slow, the low-resolution collisions were more obvious and hence the worse rating. However, as velocity increased, participants actually rated the more accurate collisions worse than the medium level and those with randomised low-level responses. Although this may seem contradictory, it actually backs up the argument that we make that humans are not very good at accurately detecting anomalous collisions when there is more than one dimension of information, as Profitt and Gilden [1989] also showed. Instead, they used arbitrary and subjective strategies when presented with the more complex collisions, such as:

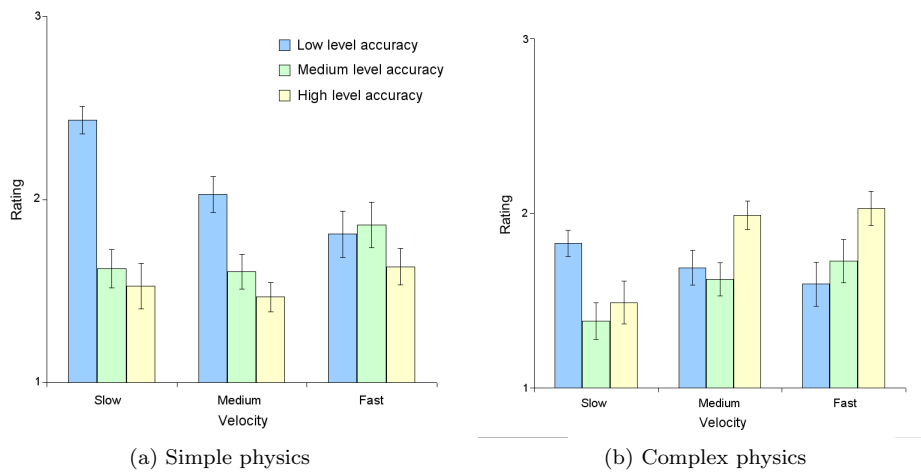


Fig. 13. Effect of velocity on rating of collisions

“I liked them better when they were head on” or, contradictorily: “I liked them better when they were side by side”; “They looked more bouncy when they were moving at medium speed”; “I couldn’t really tell with the fast ones, so I just rated them badly”.

This is further illustrated in Figure 14, which shows the effect of the different collision resolutions for both experiments. In the simple physics case, the effect of low resolution detection is highly significant: > 99%, whereas in the complex case, the effect is less significant: 90% and we can see that we have quite effectively masked the negative impact of reduced physics in the lowest resolution case by introducing a completely random rotation after contact. This was confirmed by the comments of almost all participants who reported that the most realistic collisions were those that spun a lot after colliding.

5. CONCLUSIONS AND FUTURE WORK

The philosophy behind a time-critical approach is to ensure real-time performance and preserve user perception of causality. This is achieved by strategic simplifications in different aspects of the simulation and in many cases we can afford to cull a significant degree of accuracy without degrading overall believability. By perceptually adaptive manipulation of simulation levels of detail we come closer to optimising the speed-accuracy trade-off and delivering a truly real-time system. In this paper we have presented several robust factors that can significantly affect a viewer’s perception of a collision.

Firstly, we saw that erroneous collisions in the periphery of a viewer’s point of fixation are less likely to be detected, while collision anomalies that occur in the presence of increasing numbers of visually homogeneous distractors are also less noticeable. We have seen that adding a time delay between object contact and collision response reduced the perception of causality and thereby the plausibility of the collision event itself, as in [Michotte 1963]. Therefore, interruption is imperative in any real-time system if the physics are to be believable, but we can reduce the

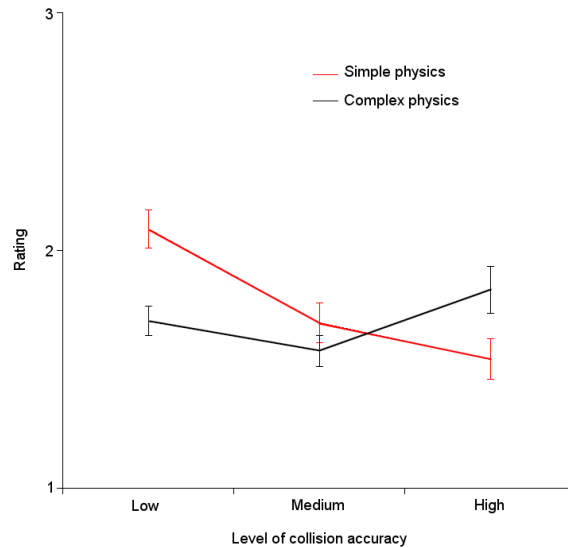


Fig. 14. Overall effect of collision detection accuracy

negative impact of the resulting loss of resolution by applying some heuristics for perceptually guided scheduling of collision processing.

For example, if we use eccentricity as our heuristic, we can have one high-priority list for collisions that project inside a region of interest on the screen, while less important collisions are relegated to a low priority list (see Figure 15). The size of this region of interest may be adapted according to the number of visually homogeneous distractors present, i.e. the more distractors, the smaller the region. Alternatively, if eye-tracking is not feasible, the viewpoint with respect to the point of collision has also been shown to be a viable heuristic for scheduling collisions, as is object velocity. We should beware of using too many variables, as the calculation of the priorities becomes a more expensive task, thus leaving even less time available for collision processing, as shown in [O’Sullivan et al. 1999]. However, shadows and interreflections: [Madison et al. 2001] and sounds: [van den Doel et al. 2001] provide strong contact cues and we intend to investigate these effects further.

We have also seen that, despite reduced collision detection resolution, it is sometimes possible to produce a random collision response that is as believable as the more accurate ones, thus further masking collision anomalies. The reason for this is that people seem to be capable of correctly perceiving errors in collision response only when there is one salient feature (such as gap size), whereas when the simulation becomes more complex they rely on their own naïve or common-sense judgements of dynamics, which are often inaccurate.

This research to date has been more qualitative than quantitative, exploring trends and relationships rather than establishing thresholds that practitioners could use in their real-time systems. These results may now be used to guide further investigations with the aim of developing psychometric functions that model human perception of collisions. One particular problem that we have encountered is the



Fig. 15. Collisions inside the region of interest receive more processing time.

difficulty in evaluating the perception of animations, due to the high number of variables involved. Perception of motion can be affected by many factors other than those we have just discussed, e.g. colour, luminance, contrast, texture, orientation, direction of motion and many more. So far, we have investigated the most robust factors that affect collision perception in relative isolation and under somewhat unrealistic conditions. Many useful insights were gained in this way, which will guide more complex experiments that examine the perception of large numbers of animated objects. We are trying to design more natural tasks, which will evaluate people's perception of collisions without actually directing their attention to them. One strategy we are pursuing is to compare the perception of real dynamical events with simulations. To this end, we are building a scaled model of a real mountain down which we will throw real rocks, while in parallel creating an identical computer model. This will provide a baseline with which to evaluate our simulations and allow us to perform real-vs-simulated psychophysical investigations. Another strand of this research involves the analysis of eye-movements while viewing such simulations, using an eye-tracking device.

REFERENCES

- AUBERT, H. AND FÖRSTER, O. 1857. Beiträge zur kenntniss des indirekten sehens (i): Untersuchungen über den raumsinn der retina. *Arch. Ophthalmology* 3, 1–37.
- BAILLARGEON, R., SPELKE, E., AND WASSERMAN, S. 1985. Object permanence in five-month-old infants. *Cognition* 20, 191–208.
- BARZEL, R., HUGHES, J., AND WOOD, D. 1996. Plausible motion simulation for computer graphics animation. In *Computer Animation and Simulation'96*. Springer-Wien, 183–197.
- BRADDICK, O. 1974. A short-range process in apparent motion. *Vision Research* 14, 519–527.
- ACM Transactions on Graphics, Vol. 0, No. 0, xyz 200.

- CARLSON, D. AND HODGINS, J. 1997. Simulation levels of detail for real-time animation. In *Proceedings Graphics Interface*. 1–8.
- CARRASCO, M. AND FRIEDER, K. 1997. Cortical magnification neutralizes the eccentricity effect in visual search. *Vision Research* 37, 1, 63–82.
- CHENNEY, S. AND FORSYTH, D. 1997. View-dependent culling of dynamic systems in virtual environments. In *ACM Symposium on Interactive 3D Graphics*. 55–58.
- CHENNEY, S. AND FORSYTH, D. 2000. Sampling plausible solutions to multi-body constraint problems. In *Proceedings Siggraph 2000*. 219–228.
- CLEMENT, J. 1982. Students' preconceptions in introductory mechanics. *American Journal of Physics* 50, 1, 66–71.
- DEVALOIS, R. AND DEVALOIS, K. 1988. *Spatial Vision*. Oxford University, New York.
- DINGLIANA, J. AND O'SULLIVAN, C. 2000. Graceful degradation of collision handling in physically based animation. *Computer Graphics Forum (Eurographics 2000 Proceedings)* 19, 3, 239–247.
- DUCHAINEAU, M., WOLINSKY, M., SIGETI, D., MILLER, M., ALDRICH, C., AND MINEEV-WEINSTEIN, M. 1997. Roaming terrain: Real-time optimally adapting meshes. In *Proceedings of IEEE Visualization '97*. 81–88.
- FUNKHOUSER, T. AND SÉQUIN, C. 1993. Adaptive display algorithm for interactive frame rates during visualization of complex virtual environments. In *Proceedings SIGGRAPH '93*. 247–254.
- GILDEN, D. AND PROFITT, D. 1989. Understanding collision dynamics. *Journal of Experimental Psychology: Human Perception and Performance*. 15, 2, 372–383.
- HODGINS, J., O'BRIEN, J., AND TUMBLIN, J. 1998. Perception of human motion with different geometric models. *IEEE Transactions on Visualization and Computer Graphics*. 4, 4, 307–316.
- HUBBARD, P. 1995. Collision detection for interactive graphics applications. *IEEE Transactions on Visualization and Computer Graphics*. 1, 3, 218–230.
- LESLIE, A. AND KEEBLE, S. 1987. Do six-month-old infants perceive causality? *Cognition* 25, 265–288.
- LUEBKE, D., HALLEN, B., NEWFIELD, D., AND WATSON, B. 2000. Perceptually driven simplification using gaze-directed rendering. Tech. rep., CS-2000-04 University of Virginia.
- MADISON, C., THOMPSON, W., KERSTEN, D., SHIRLEY, P., AND SMITS, B. 2001. Use of interreflection and shadow for surface contact. *Perception and Psychophysics* 63, 2, 187–194.
- MICHOTTE, A. 1963. *The Perception of Causality*. Basic Books., New York.
- MIRTICH, B. 1996. Impulse based dynamic simulation of rigid body systems. Ph.D. thesis, University of California, Berkeley.
- MIRTICH, B. 2000. Timewarp rigid body simulation. In *Proceedings SIGGRAPH 2000*. 193–200.
- OHSHIMA, YAMAMOTO, AND TAMURA. 1996. Gaze-directed adaptive rendering for interacting with virtual space. In *Proceedings IEEE VRAIS'96*.
- O'SULLIVAN, C., RADACH, R., AND COLLINS, S. 1999. A model of collision perception for real-time animation. In *Computer Animation and Simulation'99*, N. Magnenat-Thalmann and D. Thalmann, Eds. Springer-Wien, 67–76.
- PROFITT, D. AND GILDEN, D. 1989. Understanding natural dynamics. *Journal of Experimental Psychology: Human Perception and Performance*. 15, 2, 384–393.
- REDDY, M. 1997. Perceptually modulated level of detail for virtual environments. Ph.D. thesis, University of Edinburgh.
- ROVAMO, J. AND VIRSU, V. 1979. An estimation and application of the human cortical magnification factor. *Experimental Brain Research* 37, 495–510.
- SAARINEN, J. 1994. Visual search for global and local stimulus features. *Perception* 23, 237–243.
- STRASBURGER, H., RENTSCHLER, I., AND JR., L. H. 1994. Cortical magnification theory fails to predict visual recognition. *European Journal of Neuroscience* 6, 1583–1588.
- TOOTELL, R., SILVERMAN, M., SWITKES, E., AND VALOIS, R. D. 1982. Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science* 218, 902–904.
- TREISMAN, A. 1982. Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance* 8, 194–214.

- VAN DEN DOEL, K., KRY, P. G., AND PAI, D. K. 2001. Foleyautomatic: Physically-based sound effects for interactive simulation and animation. In *Proceedings SIGGRAPH 2001*. 537–544.
- WATSON, B., FRIEDMAN, A., AND MCGAFFEY, A. 2001. Measuring and predicting visual fidelity. In *Proceedings SIGGRAPH 2001*. 213–220.
- WATSON, B., WALKER, N., HODGES, L., AND WORDEN, A. 1997. Managing level of detail through peripheral degradation: effects on search performance in a head mounted display. *ACM Trans. Computer-Human Interaction* 4, 4, 323–346.
- WEYMOUTH, R. 1958. Visual sensory units and the minimal angle of resolution. *American Journal of Ophthalmology* 46, 102–113.
- WITKIN, A., BARAFF, D., AND KASS, M. 2001. Physically based modelling. In *Siggraph 2001 Course Notes* 25.