CS448B Critique:
Analytical Methods for Dynamic Simulation of Non-penetrating Rigid Bodies

Walter Luh
March 1, 2000

1 Paper Citation

Baraff, David. Analytical Methods for Dynamic Simulation of Non-Penetrating Rigid Bodies, *Computer Graphics (Proc. SIGGRAPH)*, volume 23, July 1989, ACM, pp. 223-232.

2 Synopsis

This paper presents a method for calculating the resting contact forces that would naturally arise to prevent bodies from interpenetrating. It deals with systems of rigid polyhedral bodies in motion or in static equilibrium that may touch at multiple contact points. A system of constraints guarantees that contact forces will prevent interpenetration and satisfy the laws of Newtonian dynamics. Linear programming techniques and heuristics are used to obtain a solution.

3 Background

Analytical techniques for simulating the motion of rigid bodies that are completely *unconstrained* in their allowable motion—simulations that don't care about collisions—are well known to engineers and physicists. Given external forces, simulation of the body's response can be implemented using numerical techniques for solving coupled differential equations.

The more demanding component of the simulation process occurs when rigid body motion is *constrained*, i.e. the bodies are solid and interpenetration between them is not allowed. Non-penetration is enforced by computing appropriate contact forces between contacting bodies. Given these forces, this problem reduces to solving the unconstrained motion treating these contact forces as additional external forces.

When calculating forces, it is important to distinguish between collision and contact forces. Collision forces are impulsive forces that exist only for an instant; they cause discontinuities in a body's velocity. Resting contact forces are continuous over some time interval; they prevent rigid body inter-penetration.

4 Summary

Analytical and penalty methods offer two ways to prevent object inter-penetration. Though they are simple to implement and easily extendible to non-rigid bodies, penalty methods can be computationally expensive for rigid bodies. Their correctness is difficult to verify.

Analytical methods in contrast provide exact answers and require fewer time steps to maintain accuracy. Previous work has focused on analytical methods for calculating the forces between *colliding* rigid bodies. This model prevents inter-penetration by modeling contact as a series of frequently occurring collisions.

Baraff's main contribution was to propose a method for calculating resting contact forces, yielding more dynamically correct simulations.

Assumptions on Contact Bodies in contact are assumed to touch each other at some number of contact points. The relative motion between these points indicates whether two bodies are colliding, resting, or separating. In the real world, a contact force is distributed evenly across a region of contact. For polygonal contact regions, the problem is simplified by finding equivalent forces acting only at the vertices of the contact region boundary.

Constraints on Contact Forces A set of contact forces must satisfy four conditions:

- The contact forces do not allow the bodies to inter-penetrate.
- The contact forces can "push" but not "pull."
- The contact forces occur only at contact points; once two bodies have separated at a contact point, there is no force between them at that contact point.
- Viewed as a function of time, contact forces are continuous.

These constraints can be formulated as a series of constraints one of which is quadratic in $\vec{f}$, a vector of contact forces to be solved. Further, there are inequalities involved, so this system requires quadratic programming to be solved directly, an NP-hard problem in general. Thus, the author presents a heuristic algorithm to determine a correct $\vec{f}$.

Heuristic Solution Methods

Once a correct set of vanishing points is found, linear programming can be used to determine a correct set of contact forces (which has O(n) expected time for this system). If the relative acceleration between two objects at a contact point is positive, this point is considered vanishing; otherwise, it is non-vanishing.

In practice, this set must be guessed. The basic idea is to guess an approximate solution of contact forces which, in turn, will indicate whether a contact point is vanishing. The author claims that if an incorrect set is found, reasonable results still follow because there are usually few such vanishing contact points and the effects of incorrect forces are small.

Simultaneous Collisions

The author briefly discusses how the linear programming formulation for resting contact can improve the performance of existing collision methods for certain configurations. In essence, collision impulses are calculating by drawing parallels from the resting contact problem.

5 Comments

This paper gives a thorough and well-written description of the physical considerations involved in realistic simulations, especially in the preventing object inter-penetration. The beginning of the paper does a good, but succint, job of explaining why previous analytical work and penalty techniques may yield incorrect results.

Unfortunately for the author and others investigating the problem of realistic dynamic simulation, any analysis using Newton's equations of motion will be inherently cumbersome. For this paper, this resulted in a system of inequalities (not easily solvable) instead of a system of equalities (easily solvable). Because the author has resorted to linear programming and heuristics, the author's method requires a good understanding of numerical analysis and optimization to implement.

The author touts the correctness of his analytical methods because they are based on Newton's equations themselves; however, his heuristic algorithm solves for the contact forces by making predictions about the locations of vanishing contact points. If the prediction is incorrect, the simulation ceases to become physically correct. The author states that satisfactory results are still produced, but this assumption relies on configurations where vanishing contact points occur infrequently. It is not clear what limitations exist on the complexity of the system given their assumptions. For example, how would a crumbling house of cards (in which multiple resting contact points vanish) fare? Would energy be satisfactorily conserved?

Ultimately, evaluating methods for dynamic simulations must depend on which is more important: correctness or ease of implementation. It appears that if a large premium is placed on correctness, an appropriate method will be more difficult to implement, (at least based on a superficial glance of the evidence presented in this paper). The correctness of penalty methods are difficult to verify but easy to implement and extendible to non-rigid bodies; analytical methods are harder to implement, but provide more correct results.

6 Discussion Questions

My discussion questions are framed from the perspective that it would be nice to solve a system of linear equations rather than a system of inequalities (which necessitate linear or even quadratic programming as the case may be):

- Physicists like to use alternative analytical techniques to circumvent the practical difficulties that arise in attempts to apply Newton's equations to particular problems. Lagrange's equations, for example, constitute an alternative formulation of mechanics yielding a prescription of dynamics equivalent to Newton's equations. The advantage is being able to obtain a system of linear equations without explicitly calculating forces. They can even be used to account for the effects of nonconservative forces such as friction.

Thus, a natural question would be whether it is possible to use these tools in a general simulator. If so, would there be any advantages both in computational expense and in simulated realism to use these tools? I can naively point out two obstacles. First, these techniques usually work well in strongly coupled systems of rigid bodies and often take for granted notions of resting contact and non-penetration; therefore, the problem of preventing object interpenetration must depend on collision detection. Second, the success of this technique depends on defining the coupling between objects. For linked objects, Lagrangian dynamics is a natural choice; for a group of disjoint objects such as a set of bricks stacked in a staircase fashion, the coupling is inexorably related to resolving contact points—possibly even collision detection.

- Is it possible to reduce the problem of resting contact (and, implicitly, of object inter-penetration) to an equivalent collision detection problem? That is, whenever objects are in resting contact or come into resting contact, can we impose resting contact conditions to prevent inter-penetration rather than calculate contact forces? What I envision is something similar to penalty methods, except that we look at the conditions of the collision. If the objects are moving so slowly and the collision dissipates sufficient energy, the objects will come into resting contact, in which case we impose resting contact conditions (rather than base resting contact on the contact force). If the objects collide with sufficient energy, we simply calculate the resulting trajectories as usual, i.e. conserving momentum and energy.