Separation-Sensitive Collision Detection for Convex Objects*

Jeff Erickson† Leonidas J. Guibas‡ Jorge Stolfi§ Li Zhang¶

Collision detection is an algorithmic problem arising in all areas of computer science dealing with the simulation of physical objects in motion. Often the problem is broken up into two parts, the so-called *broad phase*, in which we identify the pairs of objects we need to consider for possible collision, and the *narrow phase* in which we track the occurrence of collisions between a specific pair of objects [7]. For the broad phase almost all authors use some kind of simple bounding volumes for the objects so as to quickly eliminate from consideration pairs of objects that cannot possibly collide. The narrow phase is more specialized, according to the types of objects being considered.

The simplest objects to consider are convex polytopes (polygons in the plane, or polyhedra in 3-space), and this case has been extensively considered in the literature. More complex objects are approximated by their convex hulls, or broken up into convex pieces, which are tested pairwise. Since the distance between two continuously moving polytopes also changes continuously, many well-known collision detection algorithms [5, 9, 10, 11] are based upon tracking the closest pair of features of the polytopes during their motion. The efficiency of these algorithm is based on *temporal coherence*—in a sufficiently small time step, the closest pair of features will not change, or will change to some nearby features on the polytopes.

Though it is hard to imagine how one can do better than tracking the closest pair of feature when the polytopes are in close proximity, such tracking seems to be unnecessarily complicated when the polytopes start moving further from each other. Indeed most authors suggest performing first a simple bounding volume (box or sphere) test on the two polytopes, and only if that fails entering the closest feature pair tracking mode. In this paper we consider a number of general techniques to perform collision detection between two moving convex polytopes in a way that is *sensitive to the separation* between the polytopes. In order to properly quantify the separation sensitivity of our methods, we view the collision detection problem in the context of *kinetic data structures* (or *KDSs* for short), introduced in [1, 6].

In the kinetic setting we assume that the instantaneous motion laws for our polytopes are known, though they can be changed at will by appropriately notifying the KDS. Our sampling of time is not fixed, but is determined by the failure of certain conditions, called *certificates*. In our case these are *separation certificates*, which prove that the two polytopes do not intersect. The failure of a separation certificate need not mean that a collision has occurred; it can simply mean that that certificate has to be replaced by one or more others, still proving the non-intersection of the polytopes. A kinetic data structure is *compact* if it requires little space, *responsive* if it can be updated quickly after a certificate failure, *local* if it adjusts easily to changes in the motion plans of the objects, and *efficient* if the total number of events is small. Our collision data structures have all these properties.

Our key contribution is the development of a class of new kinetic data structures for collision detection between convex polytopes, where the efficiency of the structure can be analyzed in terms of natural attributes of the motion. Given two moving convex polygons in the plane, let $\mu = \min\{n, \sqrt{D/\sigma}\}$, where n is the combinatorial complexity of the polygons, D is their maximum diameter, and σ is their minimum separation during the entire motion. We maintain a single separation certificate, which changes $O(\log \mu)$ times when the relative motion of the two polygons is a linear translation, $O(\mu)$ for general translation along algebraic trajectories, and $O(\mu^2)$ during algebraic rigid motions (translation and rotation). Each certificate update is performed in $O(\log \min\{n, \sqrt{D/s}\})$ time, where s is the current separation between the polygons. In contrast to this, the closest pair of features between two polygons can change

*To appear at the 10th Annual ACM-SIAM Symposium on Discrete Algorithms, 1999. For the latest version of this paper, see http://www.uiuc.edu/ph/www/jeffe/pubs/kollide.html.

†Center for Geometric Computing, Department of Computer Science, Duke University, and Department of Computer Science, University of Illinois, Urbana-Champaign; jeffe@cs.uiuc.edu; http://www.uiuc.edu/ph/www/jeffe/. Portions of this research were done at the International Computer Science Institute, Berkeley, CA. Research partially supported by National Science Foundation grant DMS-9627683 and by U.S. Army Research Office MURI grant DAAH04-96-1-0013.

‡Computer Science Department, Stanford University; guibas@cs.stanford.edu; http://graphics.stanford.edu/~guibas/. Research partially supported by National Science Foundation grant CCR-9623851 and by U.S. Army Research Office MURI grant 5-23542-A.

§Instituto de Computação, Unversidade Estadual de Campinas, Brasil; stolfi@dcc.unicamp.br; http://www.dcc.unicamp.br/~stolfi/. Research partially supported by CNPq grant 301016/92-5.

¶Computer Science Department, Stanford University; lizhang@cs.stanford.edu; http://graphics.stanford.edu/~lizhang/

$\Omega(n^2)$ times under an algebraic rigid motion, no matter what their separation is; similar lower bounds apply to other simple separation certificates such as inner common tangent lines. Thus, our data structures interpolate smoothly between bounding volume methods, which are useful only when objects are very far apart, and local feature-based algorithms, which are efficient only when the polygons almost touch, as the separation between the objects varies. Under some mild assumptions about the nature of separation certificates, our translation bound is optimal—in the worst case, $\Omega(\mu)$ different separation certificates *must* be used during an algebraic translation.

Our collision detection algorithms are based on outer approximation hierarchies for the convex polygons involved or for their Minkowski sum—a topic which we believe is of independent interest. All these hierarchies tile the space exterior to the polygon so as to simplify the combinatorial structure of the polygon more and more as one goes away from the polygon. For a single convex polygon P, the hierarchies are obtained by starting with the bounding box of P and cutting off vertices by lines tangent to P. The resulting hierarchies are similar to the standard Dobkin-Kirkpatrick hierarchy, but careful choice of the cutting lines give us additional useful metric properties. For example, in the *Dudley hierarchy* of P (defined using results from [4]), there are only $O(\sqrt{D/s})$ vertices whose distance from P is s or greater.

Given hierarchical approximations of two convex polygons, we can combine them to obtain a decomposition of the complement of their Minkowski sum into triangles and parallelograms, called a *mixed hierarchy*. The mixed hierarchy changes in a very regular way as the polygons rotate. At certain discrete events, a triangle and a neighboring parallelogram exchange positions, much like a Delaunay flip. This regularity makes it possible to maintain a separation certificate for two moving convex polygons by separating the motion into a *translational* part, corresponding to a single point moving around in the plane, and a *rotational* part, corresponding to a deformation of the triangle or parallelogram containing that point, including possible page turns. Thus, we can maintain a separation certificate consisting of a single tile, or perhaps a small working set of tiles, around the current configuration point. Maintaining this certificate efficiently requires only the individual polygon hierarchies; we do not need to construct or maintain the entire mixed hierarchy.

The big-Oh notation in our bounds hides constants that depend on the algebraic degree of the motion. When the polytopes are in close proximity, it is clear that a 'wiggly' motion should be more costly than a smooth one. But why should it be so when they are far away? This has motivated us to develop structures that exhibit *hysteresis*—after a certificate failure has occurred, the new separation certificate cannot fail until the objects have moved by some constant fraction of their current separation. Hysteresis allows us to derive upper bounds on the number of events based on geometric properties of the path that, unlike our other results, do not require the motion to be smooth or algebraic. In particular, we can bound the number of events by the combinatorial size of a certain cover of the motion path by balls, somewhat reminiscent of [12].

The distance- and motion-sensitive bounds we give are novel and, to our knowledge, the first such to be ever presented. It was surprising to us that even for the simple setting of two moving convex polygons, there is that much novel and interesting to say; in fact, many challenging open questions remain. We anticipate few significant obstacles in extending our results to three dimensions. We expect that our kinetic structures will lead to improved practical algorithms, and we plan to undertake an implementation in the near future.

In a companion paper [2], we discuss a different set of kinetic collision techniques applicable to non-convex shapes.

References

[1] J. Basch, L. J. Guibas, and J. Hershberger. Data structures for mobile data. *Proc. 8th ACM-SIAM Sympos. Discrete Algorithms*, pp. 747–756. 1997.

[2] J. Basch, J. Erickson, L. J. Guibas, J. Hershberger, and L. Zhang. Kinetic collision detection for two simple polygons. To appear in *Proc. 10th ACM-SIAM Sympos. Discrete Algorithms*, 1999. ⟨http://www.cs.uiuc.edu/ph/www/jeffe/pubs/cdsimple.html⟩.

[3] D. P. Dobkin and D. G. Kirkpatrick. Fast detection of polyhedral intersection. *Theoret. Comput. Sci.* 27:241–253, 1983.

[4] R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *J. Approx. Theory* 10:227–236, 1974.

[5] E. G. Gilbert, D. W. Johnson, and S. S. Keerthi. A fast procedure for computing the distance between complex objects. *IEEE Journal of Robotics and Automation* 4(2), 1988.

[6] L. J. Guibas. Kinetic data structures: A state of the art report. To appear in *Proc. 3rd Workshop on Algorithmic Foundations of Robotics*, 1998.

[7] P. M. Hubbard. Collision detection for interactive graphics applications. *IEEE Trans. Visualization and Computer Graphics* 1(3):218–230, Sept. 1995.

[8] J. Klosowski, H. Held, J. S. B. Mitchell, K. Zikan, and H. Sowizral. Efficient collision detection using bounding volume hierarchies of k-DOPs. *IEEE Trans. Visualizat. Comput. Graph.* 4(1), 1998.

[9] M. C. Lin and J. F. Canny. A fast algorithm for incremental distance calculation. *Proc. IEEE Internat. Conf. Robot. Autom.*, pp. 1008–1014. vol. 2, 1991.

[10] B. Mirtich. V-clip: fast and robust polyhedral collision detection. Technical Report TR-97-05, Mitsubishi Electrical Research Laboratory, 1997. ⟨http://www.merl.com/projects/vclip/⟩.

[11] B. Mirtich and J. Canny. Impulse-based dynamic simulation. *The Algorithmic Foundations of Robotics*. A. K. Peters, 1995.

[12] J. S. B. Mitchell, D. M. Mount, and S. Suri. Query-sensitive ray shooting. *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pp. 359–368. 1994.