

FINAL - 01 JUN 92 TO 31 JUL 95

GEOMETRIC ALGORITHMS FOR MANUFACTURING, MACHINING AND DESIGN

AFOSR-91-0328
3484/S5 61102F

5. AUTHOR(S)
PAUL CHEW, BRUCE DONALD AND DANIEL HUTTENLOCHER

CORNELL UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE
4130 UPSON HALL
ITHACA, NEW YORK 14853-7501

AFOSR-TR-96

0083

AFOSR/NM
110 DUNCAN AVE, SUITE B115
BOLLING AFB DC 20332-0001

AFOSR-91-0328

12. DISTRIBUTION AVAILABILITY STATEMENT

Approved for public release;
distribution unlimited.

Methods have been developed for task-level assembly and parametric design. Efficient methods for determining the shape of parts given specifications of their functions and parameterized geometry have been determined.

19960320 057

14. SUBJECT TERMS

15. NUMBER OF PAGES

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT

UNCLASSIFIED

18. SECURITY CLASSIFICATION OF THIS PAGE

UNCLASSIFIED

19. SECURITY CLASSIFICATION OF ABSTRACT

UNCLASSIFIED

20. LIMITATION OF ABSTRACT

UNCLASSIFIED

1 Introduction

Efficient and practical algorithms for solving geometric problems are important for applications ranging from computer aided design to automatic parts inspection and assembly. This research program focused on the development and implementation of geometric algorithms, for solving problems in design and manufacturing. The methods that we have developed combine techniques from combinatorial geometry and algebra with methods from engineering and computer science, in order to produce systems that have a firm underlying theoretical basis and that also work well in practice. Geometric algorithms underly many of the computational tools for design and manufacturing: geometric modelers are widely used, and provide speed and expressive power far beyond paper-based design methods; differential equation modelers and visualization systems use geometric algorithms for subproblems such as meshing and volume rendering; visual inspection systems are beginning to use geometric model-based matching algorithms to increase speed, accuracy and flexibility.

One of the central tenets of our approach to research in this area has been to develop fundamental geometric algorithms that have a wide range of applicability to problems in manufacturing, design and beyond. Under this contract we have developed theoretical algorithms and implemented systems for,

1. Comparing shapes based on the Hausdorff distance. These methods have not only been used in a number of research projects (including visual inspection, visual control, and video segmentation) but have also been applied by Hughes Aircraft Company for target recognition, by Xerox Corporation for document processing, and General Electric Company for satellite image interpretation.
2. Task-level assembly and parametric design. We have developed efficient methods for determining the shape of parts such as snap fasteners, given specifications of their parameterized geometry and the functions that they should perform.
3. Producing guaranteed-quality triangular meshes in the plane and on surfaces. These methods have been applied in the computer graphics community (e.g., by Welch et al) to develop systems for the manipulation of free-form surfaces.
4. Model-based visual inspection and servoing. Model-based geometric matching techniques have recently supplanted traditional statistical and correlation-based methods for problems such as automated target recognition (ATR). We have been developing efficient and robust geometric matching algorithms, and applying them to matching and recognition problems.

5. Parametric design of nano-structures. We have been developing algorithms for determining the shape of microelectromechanical (MEM) structures given specifications of their parameterized geometry and functionality. MEM structures are particularly suited for automated parametric design because their production process allows only a limited number of elementary structures, the geometry and functionality of the devices is relatively simple, and a given array may involve a large number of devices (analogous to VLSI where manual design is impossible).

2 Technical Accomplishments

Computational geometry provides powerful tools and techniques for use in many problem domains. While practical applications have long been a motivation for developments in computational geometry, the resulting algorithms have often been more of theoretical than of practical interest. Our group at Cornell, on the other hand, has a history of applying tools of computational geometry in order to solve practical problems, in areas such as robotic assembly, model-based object recognition, automatic target recognition and mesh generation.

2.1 Geometric Algorithms for Visual Matching

A central problem in pattern recognition and computer vision is determining the extent to which one image is like another. Operations such as correlation, template matching and model-based matching are all techniques for determining the difference between images. We have been investigating new measures for determining the degree to which two shapes differ from one another. These measures have two critical properties: (i) they tolerate outliers, or are robust in a statistical sense, and (ii) they can be computed efficiently. The measures that work best in practice are based on the Hausdorff distance.

Given two point sets \mathcal{P} and \mathcal{Q} , with m and n points respectively, and a fraction, $0 \leq f \leq 1$, the generalized Hausdorff measure is defined as

$$h_f(\mathcal{P}, \mathcal{Q}) = f^{\text{th}} \min_{p \in \mathcal{P}} \max_{q \in \mathcal{Q}} \|p - q\|, \quad (1)$$

where $f^{\text{th}} \min_{p \in \mathcal{P}} g(p)$ denotes the f -th quantile of $g(p)$ over the set \mathcal{P} . For example, the 1-th quantile is the maximum (the largest element), and the $\frac{1}{2}$ -th quantile is the median. This generalizes the classical Hausdorff distance, which *maximizes* over $p \in \mathcal{P}$. Hausdorff based measures are asymmetric, for example $h_f(\mathcal{P}, \mathcal{Q})$ and $h_f(\mathcal{Q}, \mathcal{P})$ can attain very different values as there may be points of \mathcal{P} that are not near any points of \mathcal{Q} , or vice versa. This asymmetry is useful in recognition problems, where a hypothesize and test paradigm is often employed.

The generalized Hausdorff measure has been used for a number of matching and recognition problems. There are two complementary ways in which the measure has been employed:

1. Specify a fixed fraction, f , and then determine the distance, $d = h_f(\mathcal{P}, \mathcal{Q})$. In other words, find the smallest distance, d , such that $k = \lceil fm \rceil$ of the points of \mathcal{P} are within d of points of \mathcal{Q} . This has been termed "finding the distance for a given fraction". Intuitively, it measures how well the best subset of size $k = \lceil fm \rceil$ of \mathcal{P} matches \mathcal{Q} , with smaller distances being better matches.

2. Specify a fixed distance, d , and then determine the resulting fraction of points that are within that distance. In other words, find the largest f such that $h_f(\mathcal{P}, \mathcal{Q}) \leq d$. Intuitively, this measures what portion of \mathcal{P} is near \mathcal{Q} , for some fixed neighborhood size, d . This has been termed "finding the fraction for a given distance." It measures how well two sets match, with larger fractions being better matches.

Most applications of the measure are based on the second of these interpretations, computing the *Hausdorff fraction*, because in most visual matching problems there is a reasonable prior estimate of the uncertainty in the positional location of image features. For example, a positional error of one pixel is generally introduced by the digitization process. If the feature points are edge features, then there is an uncertainty based on the degree of smoothing of the image. Efficient methods for finding the transformations of one point set such that the Hausdorff fraction is above some threshold (and the distance below some threshold) have been developed for affine transformations of the plane. When the transformations are restricted to translations the fastest methods use dilation and correlation, whereas for full affine transformations the fastest methods use a hierarchical decomposition of the parameter space. The initial methods for computing Hausdorff distances were combinatorial algorithms using techniques from computational geometry, but current practice does not use these combinatorial techniques.

The Hausdorff matching methods have been used in a number of applications, including visually guided navigation, motion tracking, automatic target recognition and document processing. The transformations from model-to-image that can be computed efficiently include translation (illustrated above), translation with separate scaling along the two image axes, and full affine transformations (translation, rotation, scaling and shearing). These practical methods are derived from theoretical investigations reported in papers authored under this contract (see below). These works have been a major influence in the recent increased research activity in geometric algorithms for pattern matching problems. One of the central reasons for the success of these methods is the fact that they are insensitive to outliers, unlike linear methods such as correlation (that are based on measures such as the mean).

In recent work with D. Dorit, A. Efrat, and K. Kedem, Paul Chew has shown that, using the L_∞ metric, the minimum Hausdorff distance under translation between two point sets of cardinality n in d -dimensional space can be computed in time $O(n^{(4d-2)/3} \log^2 n)$ for $d > 3$. This improves the previous time bound of $O(n^{2d-2} \log^2 n)$ due to Chew and Kedem. For $d = 3$ we obtain a better result of $O(n^3 \log^2 n)$ time by exploiting the fact that the union of n axis-parallel unit cubes can be decomposed into $O(n)$ disjoint axis-parallel boxes. We prove that the number of different translations that achieve the minimum Hausdorff distance in d -space is $\Theta(n^{\lfloor 3d/2 \rfloor})$; thus, it can take significantly longer to report all solutions than it does to find a single solution. The proof for the improved time bound is based on the use of Orthogonal Partition Trees (OPTs), a geometric data structure originally developed by Overmars and Yap. OPTs cannot be used directly for this problem; thus, it was necessary to develop a novel two-level version of the OPT data structure.

2.2 Robotic Design, Manipulation and Assembly

Minimalism pursues the following agenda: For a given robotics task, find the minimal configuration of resources required to solve the task. Thus, minimalism attempts to reduce the resource signature for a task, in the same way that (say) Stealth technology decreases the radar signature of an aircraft. Minimalism is interesting because doing task A without resource B proves that B is somehow inessential to the information structure of the task. We present experimental demonstrations and show how they relate to our theoretical proofs of minimalist systems.

In robotics, minimalism has become increasingly influential. Marc Raibert showed that walking and running machines could be built without static stability. Erdmann and Mason showed how to do dextrous manipulation without sensing. Tad McGeer built a biped, kneed walker without sensors, computers, or actuators. Rod Brooks has developed on-line algorithms that rely less extensively on planning and world-models. We have taken a minimalist approach to distributed manipulation. First, we described how we built distributed systems in which a team of mobots cooperate in manipulation tasks without explicit communication.¹ Second, we are now building arrays of micromanipulators to perform sensorless micromanipulation. We will soon describe how well our experimental design worked, and how manipulation experiments using it mirrored the theory.

Some Details

This report describes our experience in building distributed systems of robots that perform manipulation tasks. We have worked at both the macroscopic and the microscopic scale. First, we described a team of small autonomous mobile robots that cooperate to move large objects (such as couches). The robots run SPMD² and MPMD² manipulation protocols with no explicit communication. We developed these protocols by distributing off-line, sequential algorithms requiring geometric models and planning. The resulting parallel protocols are more on-line, have reduced dependence on a priori geometric models, and are typically robust (resistant to uncertainty in control, sensing, and initial conditions).

We now discuss our work on sensorless manipulation using massively parallel arrays of microfabricated actuators. The single-crystal silicon fabrication process opens the door to building monolithic microelectromechanical systems (MEMS) with microactuators and control circuitry integrated on the same chip. Our actuators are servoed to uniquely orient (up to symmetries) an object lying on top, and require no sensing. We can also program the array as a sensorless geometric filter—to sort parts based on shape or size.

We developed both the macroscopic and the microscopic systems by distributing and parallelizing sequential manipulation algorithms with global control, to obtain distributed algorithms running on independent physical agents. Our MEMS control algorithms for micromanipulation are SPMD; for the macroscopic (furniture-moving) task, we described implementations and experiments with both SPMD and MPMD control.

We have implemented and extensively tested our macroscopic distributed manipulation strategies—at ISER, we'll show videos of the robot systems we built and discuss how well they performed. We have built MEMS prototypes, and we are now fabricating and testing

¹No RF or IR messages are sent between the robots.

²SPMD (MPMD) = *Single (Multiple) Program, Multiple Data*.

our biggest micro-array yet (the entire wafer can be tiled with 11,000 microactuators in one square inch). In March and April we performed manipulation experiments with this array, and we'll report on the results at ISER. Our macroscopic algorithms use no direct communication between the agents, but do employ sensing. Our microscopic algorithms are sensorless, but require a small amount of local communication to initialize and synchronize. Our theory predicts a trade-off between communication and sensing when we parallelize a manipulation strategy. We'll discuss experiments we have performed to experimentally observe and validate these trade-offs.

Sample Experimental Results

We only detail one of our macroscopic MPMD results; the papers below detail our SPMD protocols in both the micro- and the macro- cases. Here is an example of our results. At ISER we will show videos demonstrating our implementation on a physical team of mobile robots.

1. *Information invariants theory indicates that we should be able to distribute a manipulation task across multiple robots with essentially no additional resource cost.*

This idea is borne out by an examination of the program code that we run on our *Pusher/Steerer* MPMD distributed manipulation system system: The *steerer* has only the code that pertains to navigation, and the *pusher* has only the code that pertains to pushing the object. We use no explicit communication, so no program code needs to be added for this purpose.

2. *A mechanics analysis of large-scale manipulation indicates that distributing a manipulation task across multiple robots using a Pusher/Steerer model will allow robots with limited control and sensing to perform that manipulation task in a manner equivalent to a single-robot system with much more sophisticated control and sensing.*

Using robots with simple 8-pushbutton bumpers and simple SCHEME control-loops, we were able to manipulate a variety of objects through complex paths; we were able to change and manipulate new objects without changing the programs, and with no re-tuning.

(Our MEMS work is joint with Noel MacDonald and Rob Mihailovich in the Cornell National Nanofabrication Laboratory).

2.3 Meshing and Voronoi Diagrams

Mesh generators are useful for the numerical solution of partial differential equations (PDEs). Such numerical solutions are needed for a wide variety of engineering problems including, for example, fluid dynamics, structural analysis, thermal analysis, and electromagnetics. For several commonly-used PDE solution techniques (the Finite Element Method and the Boundary Element Method, for instance), the first step is to divide the problem region into simply-shaped elements creating a *mesh*. For two dimensional problems the elements are usually triangles or quadrilaterals; for three dimensional problems the elements are usually tetrahedra or hexahedra.

Chew has developed an algorithm for creating high-quality triangular meshes for regions on curved surfaces. For both flat and curved surfaces, the resulting meshes are guaranteed to exhibit the following properties:

- internal and external boundaries are respected,
- element shapes are guaranteed — all elements are triangles with angles between 30 and 120 degrees (with the exception of badly shaped elements that may be required by the specified boundary), and
- element density can be controlled, producing small elements in “interesting” areas and large elements elsewhere.

This work also included the development of a practical generalization of *Delaunay triangulation* to curved surfaces. The Delaunay triangulation and its geometric dual the *Voronoi diagram* have proved to be among the most useful data structures in all of Computational Geometry. See, for instance, the text by Preparata and Shamos for more information on Delaunay triangulations and Voronoi diagrams. Basically, for a set S of points in the plane, the *Delaunay triangulation* of S is triangulation of S with the Empty Circle Property: for each triangle Δ of the triangulation, the circumcircle of Δ contains no points of S in its interior.

Chew’s generalization of the Delaunay triangulation to curved surfaces has been used in graphics to enable freeform surface manipulation. In this work, a curved-surface Delaunay triangulation is maintained on the freeform surface. Manipulation is accomplished by moving the vertices and edges of the Delaunay triangulation. As vertices are moved and the surface is stretched, the Delaunay triangulation is maintained via edge flipping and the introduction of new vertices on the surface.

An important point about Chew’s work on mesh generation is the existence of a quality guarantee for the resulting mesh. This guarantee is in the form of a mathematical proof that the mesh has certain desirable properties. This type of guarantee is useful but not absolutely necessary for many 2D problems — if a heuristic technique is used, one without a mathematical guarantee, then a user can inspect the mesh, correcting and improving the mesh as necessary. But for 3D problems, a user can no longer inspect the mesh in an effective way — it is just too difficult to understand a picture of a 3D tetrahedral mesh. Thus mathematical guarantees are likely to be necessary for truly useful mesh generation in 3D.

2.3.1 3D Voronoi Diagrams

Recent work by Paul Chew (in collaboration with Klara Kedem, Micha Sharir, Boaz Tagansky, and Emo Welzl) has led to a breakthrough in the understanding of 3D Voronoi diagrams. We have developed the first subcubic bound on the size of Voronoi diagrams that use sites more complex than simple points. Our results show that, for a polygonal distance function (such distance functions include the commonly-used L_1 and L_∞ metrics), a 3D Voronoi diagram of lines has near-quadratic complexity. This size-bound indicates that such diagrams are likely to be useful for practical applications such as placement and motion planning.

Voronoi diagrams and their generalizations have proved to be among the most useful data structures in all of Computational Geometry, but until this recent work, only the simplest type of 3D Voronoi diagrams (i.e., Voronoi diagrams of sites that are simple points) were well-understood. Voronoi diagram generalizations, the generalizations needed to solve practical problems, were understood only for 2D. For instance in 2D, Voronoi diagrams of sites that are polygons are useful for motion planning — the Voronoi boundaries correspond to paths that are “far” from the polygonal obstacles. If we were limited to Voronoi diagrams of point sites then we could use them only for motion planning problems in which the obstacles are represented as points, an impractical restriction for most motion planning problems. This is the kind of situation that has existed for 3D Voronoi diagrams — useful Voronoi variations using more than point sites were not understood well enough to be used.

Papers published under the contract:

1. M. Bern, L. Paul Chew, D. Eppstein, and J. Ruppert, Dihedral Bounds for Mesh Generation in High Dimensions, *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms* (1995), 189–196.
2. K.-F. Böhringer, B.R. Donald, R. Mihailovich, and Noel C. MacDonald, “A Theory of Manipulation and Control for Microfabricated Actuator Arrays,” 7th IEEE Workshop on Micro Electro Mechanical Systems (MEMS’94). Kanagawa, Japan (January, 1994).
3. Böhringer, K.-F. *et al.* “Sensorless Manipulation Using Transverse Vibrations of a Plate,” IEEE International Conference on Robotics and Automation, Nagoya, Japan 1995.
4. K.-F. Böhringer, B. R. Donald, R. Mihailovich, and Noel C. MacDonald, “Sensorless Manipulation Using Massively Parallel Micro-fabricated Actuator Arrays,” *Proc. IEEE International Conference on Robotics and Automation*, San Diego, CA (1994).
5. A. J. Briggs and B.R. Donald, “Automatic Sensor Configuration for Task-Directed Planning,” *Proc. IEEE International Conference on Robotics and Automation*, San Diego, CA (May, 1994).
6. Brown, R. and Jennings, J., *A Pusher/Steerer Model for Cooperative Mobile Robot Cooperation*, Proc. IEEE/Robotics Society of Japan International Workshop on Intelligent Robots and Systems, Pittsburgh, PA (1995).
7. L. Paul Chew, K. Kedem, M. Sharir, B. Tagansky, and E. Welzl, Voronoi Diagrams of Lines in 3-Space Under Polyhedral Convex Distance Functions, *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms* (1995), 197–204.
8. L. Paul Chew, M. T. Goodrich, D. P. Huttenlocher, K. Kedem, J. M. Kleinberg, and D. Kravets, Geometric Pattern Matching under Euclidean Motion, *Computational Geometry: Theory and Applications*, to appear.
9. L. Paul Chew, Near-Quadratic Bounds for the L_1 Voronoi Diagram of Moving Points, *Computational Geometry: Theory and Applications*, to appear.
10. L. Paul Chew, D. Dor, A. Efrat, and K. Kedem, Geometric Pattern Matching in d -Dimensional Space, *Proceedings of the European Symposium on Algorithms* (1995).
11. B. R. Donald, J. Jennings and D. Rus, “Information Invariants for Cooperating Autonomous Mobile Robots,” International Symposium on Robotics Research (ISRR). Hidden Valley, PA (1993).
12. B.R. Donald “Information Invariants for Distributed Manipulation”, *International Journal of Robotics Research*, (Accepted, to appear) 1995.
13. B.R. Donald, P. Xavier, J. Canny, and J. Reif, “Kinodynamic Motion Planning”, *Journal of the ACM*, Vol. 40, No. 5, Nov., 1993. pp. 1048-1066.

14. B.R. Donald, "Information Invariants in Robotics," *Artificial Intelligence*, Vol. 72, Nos. 1-2, pp. 217-304 (Jan. 1995).
15. B.R. Donald, J. Jennings and D. Rus, "Analyzing Teams of Cooperating Mobile Robots," *Proc. IEEE International Conference on Robotics and Automation*, San Diego, CA (May, 1994).
16. B.R. Donald, J. Jennings and D. Rus, "Moving Furniture with Teams of Autonomous Mobile Robots," *Proc. IEEE/Robotics Society of Japan International Workshop on Intelligent Robots and Systems*, Pittsburgh, PA (1995).
17. B.R. Donald, "Distributed Robotic Manipulation: Experiments in Minimalism," *International Symposium on Experimental Robotics*, Stanford, CA (1995).
18. B.R. Donald and P. Xavier, "Provably Good Approximation Algorithms for Optimal Kinodynamic Planning for Cartesian Robots and Open Chain Manipulators" To appear in *Algorithmica*. (In Press). 1995
19. W.E.L. Grimson, D.P. Huttenlocher and D.W. Jacobs. A Study of Affine Matching with Bounded Sensor Error, *Intl. Journal of Computer Vision*, vol. 13, no. 1, pp. 7-32, 1994.
20. D.P. Huttenlocher G.A. Klanderman and W.J. Rucklidge. Comparing Images Using the Hausdorff Distance, *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 15, no. 9, pp. 850-863, 1993
21. D.P. Huttenlocher, K. Kedem and M. Sharir. The Upper Envelope of Voronoi Surfaces and its Applications, *Discrete and Computational Geometry*, vol. 9, no. 3, pp. 267-291, 1993.
22. D.P. Huttenlocher. Using Two-Dimensional Models to Interact with the Three-Dimensional World, in *3D Object Representation for Computer Vision*, edited by M. Hebert, J. Ponce and T. Boult, Springer-Verlag, to appear.
23. D.P. Huttenlocher, M.E. Leventon and W.J. Rucklidge. Visually Guided Navigation by Comparing Edge Images, in *Algorithmic Foundations of Robotics*, edited by K. Goldberg et. al., pp. 85-96, A.K. Peters, 1995.
24. D.P. Huttenlocher and E.W. Jaquith. Computing Visual Correspondence: Incorporating the Probability of a False Match, *Proceedings of the Fifth International Conference on Computer Vision*, pp. 515-522, 1995.
25. D.P. Huttenlocher, M.E. Leventon and W.J. Rucklidge. Visually Guided Navigation by Comparing Two-Dimensional Edge Images *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference*, 1994, pp. 842-847.
26. D.P. Huttenlocher and J.M. Kleinberg. Comparing Point Sets Under Projection *Proceedings of the Fifth ACM-SIAM Symposium on Discrete Algorithms*, 1994, pp. 1-7.

27. D.P. Huttenlocher and W.J. Rucklidge. A Multi-Resolution Technique for Comparing Images Using the Hausdorff Distance, *Proceedings of the IEEE Computer Vision and Pattern Recognition Conference*, 1993, pp. 705-706.
28. D.P. Huttenlocher, J.J. Noh and W.J. Rucklidge. Tracking Nonrigid Objects in Complex Scenes, *Proceedings of the Fourth International Conference on Computer Vision*, 1993, pp. 93-101.
29. Prasad, R., Böhringer, K., and MacDonald, N., *Design, Fabrication, and Characterization of Single Crystal Silicon Latching Snap-fasteners for Micro Assembly*, Proceedings ASME International Congress and Exposition (IMECE'95), San Francisco, CA (Nov. 1995).