Sebastien J. Blind
Microstrategy
Vienna, VA 22182, USA

Christopher C. McCullough
BAE Systems
San Diego, CA 92127, USA

Srinivas Akella
Department of Computer Science
Rensselaer Polytechnic Institute
Troy, NY 12180, USA
sakella@cs.rpi.edu

Jean Ponce
Beckman Institute
University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA

Manipulating Parts with an Array of Pins: A Method and a Machine

Abstract

This article investigates the manipulation of polygonal parts using a simple device consisting of a grid of retractable pins mounted on a vertical plate. This "Pachinko machine" is intended as a reconfigurable parts feeder for flexible assembly. A part dropped on this device may come to rest on the actuated pins, bounce out, or fall through. The authors propose a novel algorithm for part reorientation. Its input consists of the shape of a part, its initial position and orientation, and a goal configuration, and its output is a sequence of pin actuations that will bring the part to the goal configuration. The proposed approach does not attempt to predict the part motion between the equilibria associated with the active pins in the output sequence; instead, it constructs the capture region of each equilibrium (i.e., the maximal subset of the part's configuration space such that any motion starting within it is guaranteed to end at the equilibrium). Assuming frictionless contacts and dissipative dynamics, reorienting a part reduces to finding a path from initial to goal states in a directed graph whose nodes are the equilibria and whose arcs link pairs of nodes such that the first equilibrium lies in the capture region of the second one. The proposed approach has been implemented on a prototype of the Pachinko machine, and initial experiments are presented.

KEY WORDS—parts feeding, flexible assembly, automation, capture region, manipulation

The International Journal of Robotics Research
Vol. 20, No. 10, October 2001, pp. 808-818,

1. Introduction

This article addresses the problem of manipulating polygonal parts using a simple device consisting of a grid of retractable pins mounted on a vertical plate (Fig. 1). When a part is dropped on this device, it may come to rest on the actuated pins, bounce out, or fall through. By carefully selecting the set of actuated pins for a given part shape, we can control the set of equilibrium part configurations that are captured. The objective is to compute a sequence of pin actuations and retractions that brings the part to some desired configuration by a sequence of transitions from one set of pins to the next. We refer to the device as the "Pachinko machine" since it is reminiscent of the vertical pinball machines popular in Japan.

Our approach is based on the construction of the *capture region* (Brost 1991; Kriegman 1997) of each part equilibrium (i.e., the maximal subset of the part's configuration space such that any motion starting within this subset is guaranteed to end up in the equilibrium configuration). The capture regions are constructed from the associated kinematic and dynamic constraints in the part's configuration space. In this framework, reorienting a part reduces to constructing a directed graph whose nodes consist of equilibrium states and whose arcs link pairs of nodes, such that the first equilibrium lies in the capture region of the second one, and then exploring this graph to find paths from initial to goal nodes. We have implemented an algorithm to compute the capture regions and generate these paths and have built the prototype Pachinko machine shown in Figure 1 to illustrate the method.

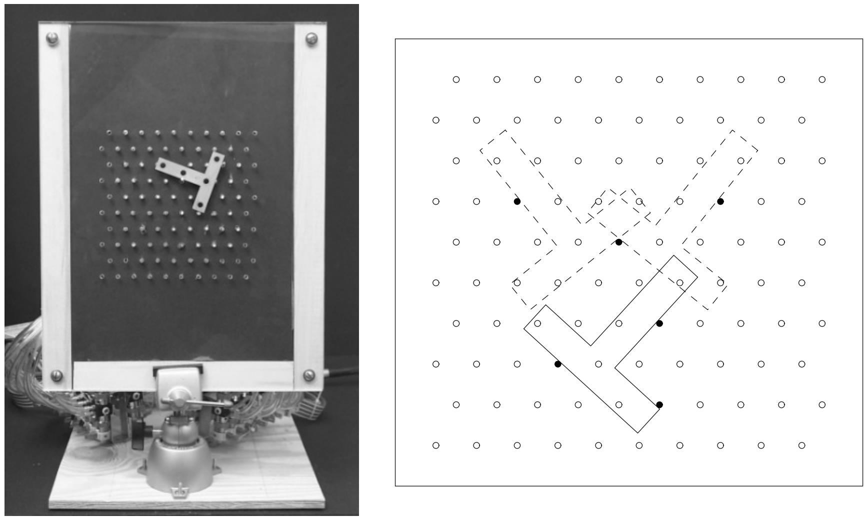

(a) (b)

Fig. 1. The Pachinko machine: (a) the actual prototype, and; (b) the schematic depiction of its operation. The capture region of the goal configuration at the bottom of the plate is used to determine whether the T-shaped part falling from one of the initial equilibrium configurations at the top will be guaranteed to end up in the goal configuration. Actuated pins are indicated by solid circles and retracted pins by hollow circles.

We do not assume that contact is maintained during the execution of a reorientation task, nor do we attempt to predict the object motion between equilibria. Indeed, such predictions are difficult and may be unreliable, even when detailed models of friction and contact dynamics are available. Instead, we assume rigid polygonal parts that make frictionless contact with the pins and support surface and assume that the mechanical energy of the part dissipates until an equilibrium is reached. Since the energy of a part starting at rest cannot exceed its initial potential energy, the part can reach and come to rest only in configurations where its center of mass does not lie above its initial height. We propose an approach that guarantees under these assumptions that a part falling on a pair or triplet of selected pins will end up being supported by these pins in a known configuration, despite the fact that it may slide, rotate, break contact, or bounce during its fall. (Note that a part may come to rest when in contact with one, two, or three pins. We use only stable equilibrium configurations in which the part makes simultaneous contact with two or three pins.) This reduces the problem of planning the overall motion of an object to constructing a sequence of deterministic transitions between its equilibrium states.

The Pachinko machine is intended for use as a parts-feeding device for flexible assembly. It uses simple binary actuators, has a modular design, and can be automatically reconfigured for different polygonal part shapes. This type of mechanism can potentially replace fixed automation devices such as bowl feeders. The Pachinko machine can be used as a parts nest to capture parts in a single configuration, with parts that are not captured being recirculated. More interestingly, it can be used to bring a part dropped from an unknown or uncertain initial configuration to a desired configuration, a process we dub reorienting the part, despite the fact that changes in both orientation and position may be involved. In principle, the Pachinko machine could also be used to distinguish and sort parts by using the same set of pins to induce different motions for different parts. In all instances, parts that are not captured or that are not in appropriate configurations can be fed back to the top of the device by a recirculating mechanism.

A preliminary version of this article appeared in Blind et al. (2000a).

2. Background and Approach

2.1. Related Work

There is a rich literature on object manipulation using a small number of contacts. For example, Fearing (1986) describes grasping and reorienting strategies for polygonal objects manipulated by two fingers of a dexterous robot hand. Abell and Erdmann (1995) use two frictionless fingers to stably support and rotate a polygon in a gravitational field. Leveroni and Salisbury (1996) use finger gaits to manipulate a polygon in the plane by sequencing the grasping fingers to move within their workspaces while maintaining a stable grasp of the object. Farahat, Stiller, and Trinkle (1995) derive solutions for the configuration of a polygon in sliding and rolling contact with two or three independently controlled polygons. Our use of binary actuators is an example of the minimalist approach advocated by Canny and Goldberg (1994) for industrial automation tasks. Lynch (1999) analyzes toppling as a manipulation primitive to change the resting face of a part on a surface.

Arrays of actuators can also be used to manipulate objects by creating vector force fields. Böhringer et al. (1994) demonstrate sensorless manipulation of objects using motions of an array of microelectromechanical actuators to create a desired vector field. Böhringer et al. (2000) present vector fields to orient and position objects in unique equilibrium configurations. Luntz, Messner, and Choset (1997) model the dynamics of a modular array of macroscale actuators to transport and manipulate objects in the plane. Force fields can also be created by a single vibrating plate. Böhringer, Bhatt, and Goldberg (1995) demonstrate sensorless manipulation of objects toward the nodes of a transversely vibrating plate, and Reznik and Canny (1998) consider orienting parts in parallel by choosing motions of a horizontally vibrating plate.

An important characteristic of any robust manipulation plan is that it should tolerate errors in part positioning. It is possible to dramatically reduce position uncertainty by actively exploiting the task mechanics. Work in this area was pioneered by Inoue (1974) and Whitney (1982), and algorithms for reorienting 2-D parts through pushing and grasping operations have been developed by several authors—notably, Mason (1986), Mani and Wilson (1985), Brost (1988), Erdmann and Mason (1988), Peshkin and Sanderson (1988), Goldberg (1993), Rao and Goldberg (1995), and Akella et al. (1996). These techniques assume some predictive model of the contact mechanics. The approach advocated in this paper does not attempt such prediction and does not assume that contact is maintained during manipulation.

Parts feeding and fixturing are industrial applications that have motivated much work in manipulation. The SONY APOS system (Shirai and Saito 1989) uses a vibrating pallet with an array of nest-like depressions designed to capture parts in a single orientation. Krishnasamy, Jakiela, and Whitney (1996) analyze the effect of shape and vibration parameters on the energy of parts and hence their efficiency for entrapment in an APOS-like vibration system. Caine (1994) considers the design of interacting part shapes to constrain motion and applies it to a vibratory bowl feeder feature design. Wiegley et al. (1997) have developed a complete algorithm to find the shortest sequence of frictionless curved fences to orient a polygonal part. Berretty et al. (1999) give algorithms to design traps—polygonal cutouts that act as filters—for vibratory bowl feeder tracks. Wallack and Canny (1997) and Brost and Goldberg (1996) have independently developed algorithms to automatically generate the set of feasible modular fixture designs for prismatic parts. Wentink, van der Stappen, and Overmars (1997) have characterized parts that can be fixtured using both point and edge contacts.

Lozano-Pérez, Mason, and Taylor (1984) introduced an approach to fine motion planning in the presence of control and sensing uncertainty, called *preimage backchaining* (see Canny 1989, Donald 1988; Erdmann 1984 for related work and Latombe 1991 for an overview). This technique shares some features with our approach: complex manipulation tasks such as peg insertion are planned as sequences of (uncertain) motions between recognizable states. Control errors are explicitly modeled in Canny (1989), Donald (1988), Erdmann (1984), and Lozano-Pérez, Mason, and Taylor (1984), but the range of part motions under physical constraints is not determined. Techniques for characterizing these motions under kinematic constraints have been proposed independently by Rimon and Blake (1996) and Sudsang, Ponce, and Srinivasa (1998, 2000). In this context, the object is constrained to lie in a pocket of configuration space, called a *cage* (Rimon and Blake 1996) or an *inescapable configuration space (ICS) region* (Sudsang, Ponce, and Srinivasa 1998, 2000).

Taking dynamics into account, Brost (1991, 1992) and Kriegman (1997) have constructed the *capture regions* of equilibria of objects dropped under the action of gravity. Their work is aimed at characterizing the regions of configuration space from which the object is guaranteed to reach a stable equilibrium, and it is the direct ancestor of the approach presented in this paper. Brost determines the kinematic and dynamic behavior of a moving polygon in contact with a fixed polygon by explicitly computing a boundary representation of the configuration space obstacle and uses this to design nests that capture dropped polygonal parts in a unique orientation. Kriegman has developed an algorithm based on Morse theory to compute the capture regions of stable poses of 3-D parts in contact with a support plane.

2.2. Proposed Approach

We consider polygonal parts and frictionless contacts and assume dissipative dynamics, meaning that the mechanical energy of the system decreases with time until the part comes to rest. Under these assumptions, the potential energy of an object starting at rest can never exceed its initial value. The configuration of an object is specified by the position of its center of mass (x, y) and its orientation θ in a reference coordinate system where the y axis is the upward vertical direction. In our case, gravity provides a downward force field, and the potential energy corresponds to the height y of the center of mass. Equilibrium configurations correspond to the local minima of the potential energy field.

Our approach relies on the concept of capture regions, as proposed by Brost (1991) and Kriegman (1997). The capture region of an equilibrium configuration is the largest region of free configuration space such that (1) kinematic and dynamic constraints prevent the object from escaping it, and (2) the object eventually comes to rest at the equilibrium configuration. In this framework, manipulation planning amounts to constructing a directed *transition graph* whose nodes consist of equilibrium states and whose arcs link pairs of nodes, such that the first equilibrium lies in the capture region of the second one, and then exploring this graph to find paths from initial to goal states.

Equilibrium configurations are identified as local minima of the potential energy in configuration space. An equilibrium configuration may be associated with one or more pins that simultaneously contact the part. As described in the next section, we will consider stable equilibrium configurations associated with a pair or a triplet of pins. To construct the associated capture regions, we use a local representation of the configuration space obstacle (or C-obstacle) in the neighborhood of every equilibrium. This local model consists of a finite collection of polygonal sections representing ranges of orientations for which the topology of the polygonal structure does not change, as well as an explicit parameterization of the actual positions of the vertices of these polygonal sections as a function of object orientation.

A capture region is constructed by sweeping a horizontal plane upwards along the vertical axis from the associated equilibrium until an escape configuration is reached. The local boundary representation of the capture region is constructed and updated during the sweep. As predicted by stratified Morse theory (Goresky and Macpherson 1980), the escape configurations correspond to certain types of saddle points on the surface of the C-obstacle. The construction of the capture region stops when the sweep plane reaches such a point. The transition graph is then constructed by an efficient algorithm for determining all pairs of equilibria such that the first element in the pair belongs to the capture region associated with the second one. The next two sections present the details of our approach. An implemented prototype of the Pachinko machine is presented, along with experiments, in Section 5.

3. Configuration Space Analysis

Our goal is to bring a part from an initial configuration at the top of the Pachinko machine to a goal configuration at its bottom by selectively retracting and activating pins as the part descends due to gravity and moves from one stable configuration to the next. In this section, we characterize the capture regions associated with these equilibria. As noted earlier, we consider polygonal parts in a downward gravity field and assume frictionless contacts and dissipative dynamics. Under these assumptions, the potential energy of an object is proportional to its height y, and its equilibria correspond to local minima of that energy under the kinematic constraints imposed by the active pins. A falling part starting at rest may bounce but never higher than where it started from, and it will eventually come to rest in a new equilibrium. The capture regions of the equilibria are wells of the height field in configuration space, and their extent is determined by critical values of the variable y.

3.1. Morse Theory

Let us assume for an instant that the configuration space obstacles are bounded by a compact smooth surface such as the one shown in Figure 2. Morse theory (Milnor 1963) shows that the level sets of such a surface are in general smooth, nonintersecting curves. They may only be singular at a finite number of critical values of the height function: minima, maxima, and saddle points. There are in fact two types of saddles: those in which two connected level curves merge into a single one as the height increases (type I saddles; see points C and E in the figure) and those in which a connected curve splits into two components (type II saddles; see point D). The qualitative shape of the level curves does not change between critical points.

In this context, the capture region associated with a minimum is the subset of the free configuration space bounded below by the obstacle and above by the plane that contains a type I saddle point with minimal height such that there exists a path from the saddle to the minimum with monotonically decreasing height (Fig. 2). In reality, of course, the surfaces of configuration space obstacles are only piecewise smooth. Stratified Morse theory (Goresky and Macpherson 1980) provides a complete classification of the level curves and critical points of piecewise-smooth surfaces whose height functions

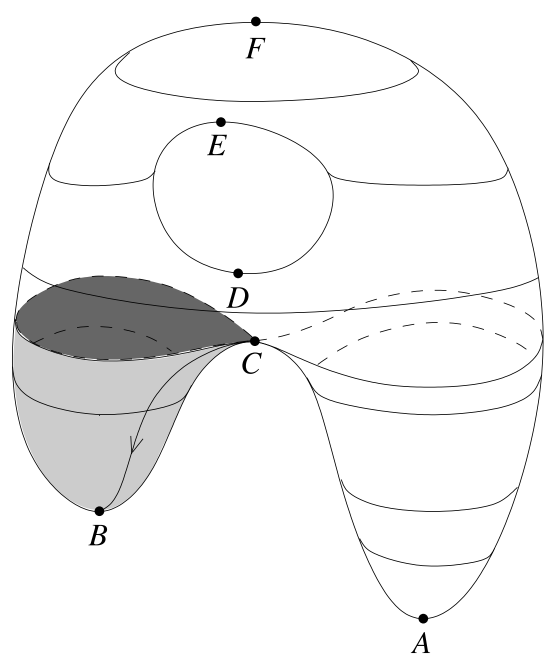

Fig. 2. Morse theory shows there are three types of critical points: minima (points A and B), maxima (point F), and type I saddles (points C and E) and type II saddles (point D). The capture region of the minimum B is shaded.

satisfy certain nondegeneracy conditions (Morse functions). It has been used by Kriegman (1997) to construct the capture regions of the stable poses of solids bounded by algebraic and polyhedral surfaces (see Rao, Kriegman, and Goldberg 1996 for applications to part reorienting using a pivoting gripper) and by Mason, Rimon, and Burdick (1995) to analyze the stable configurations of heavy objects held by multiple contacts. The rest of our discussion will not explicitly use the classification from stratified Morse theory; instead, it will exploit the existence of a finite number of (smooth or not) critical points and characterize the capture region of a minimum by finding the lowest type I saddle linked to it by a path with monotonically decreasing height.

3.2. Critical Points and Critical Curves

In general, a polygonal part may rest on one, two, or three pins of the Pachinko machine. In this presentation, we will restrict our attention to equilibria in which a part touches either two or three pins.[1] The corresponding obstacles in the x, y, θ configuration space of the part are bounded by a patchwork of ruled surfaces whose rulings run parallel to the x, y plane. When two contacts occur simultaneously, the corresponding ruled surfaces intersect along a *double-contact* (or *DC*) curve. Three contacts occur when the associated ruled surfaces and DC curves meet at a *triple-contact* (or *TC*) vertex. Consider a DC point with orientation θ_0 and define the *section* associated with this point as the intersection of the configuration space obstacle and the plane $\theta = \theta_0$. A local minimum M of a DC curve Γ corresponds to an equilibrium configuration when the edges incident to M in the associated section point upward (Fig. 3a). Likewise, a TC vertex T yields an equilibrium when the three DC curves Γ_1, Γ_2, and Γ_3 and polygonal edges incident to it point upward (Fig. 3b). A part simultaneously touching two pins such that it can translate horizontally can have a range of equilibrium configurations (Fig. 3c), which are conceptually treated as a single equilibrium.

There are four main type I saddle configurations, as shown in Figures 4a through 4d: a saddle point may occur at a local maximum (or minimum) M of the y variable along a DC curve when the corresponding polygonal section has a minimum (or maximum) at this point (Figs. 4a-4b). A saddle may also occur at a TC vertex T, when two of the incident DC curves point down (or up) and the polygonal section has a minimum (or maximum) at this point (Figs. 4c-4d). There are also four degenerate instances of these saddle point configurations corresponding to horizontal edges of the polygonal section (Figs. 4e-4f). See Blind (1999) for details.

1. Single-contact equilibria are degenerate double-contact equilibria that occur at concave vertices. They do not exist for pins with nonzero radius. A generic polygon cannot touch four (or more) pins simultaneously.

3.3. Capture Regions

Our representation of the C-obstacle consists of a finite collection of polygonal sections of the C-obstacle representing ranges of orientations for which the topology of the polygonal structure does not change, an explicit parameterization of the curves swept by the vertices of these polygonal sections as the object orientation changes, and the adjacency relationships between these polygons and curves. We determine the capture region associated with an equilibrium configuration by iteratively constructing a local boundary representation of the configuration space obstacle's surface in the vicinity of the corresponding minimum.

The local model defining the capture region is constructed by sweeping a horizontal plane upwards along the y-axis from the associated equilibrium until an escape configuration is reached. The model is initialized by using the position and orientation of the equilibrium and the adjacency information afforded by the contact edges at equilibrium. It is updated when the sweep plane reaches the height corresponding to a critical point (minimum, maximum, saddle, or vertex of the C-obstacle surface). The sweep terminates when a type I saddle is found or when the capture region vanishes (this may happen when the object is captured by the pins independent of gravity). Figures 5a through 5c show examples of capture regions associated with two or three pins. Figure 5d shows an example where the part may only touch two or three pins in any equilibrium configuration, but a fourth pin bounds its possible translations and the corresponding capture region.

During the construction of the capture region, we must also ensure that a different stable equilibrium associated with a subset of the activated pins is not in the interior of the capture region. If it is, the sweep must be terminated at the height corresponding to the equilibrium, so the capture region does not include this equilibrium.

4. Construction of the Transition Graph

The nodes of the transition graph consist of equilibria of the object as it rests on a pair or a triplet of pins, together with the associated capture regions. The directed arcs of the graph link pairs of nodes such that the equilibrium associated with the first node lies inside the capture region associated with the second one. As shown in this section, the graph can be constructed efficiently by taking advantage of the fact that the pin positions are arranged on a grid.

4.1. Enumeration of the Graph Nodes

To compute the nodes of the transition graph, we first enumerate all the pin configurations that may achieve simultaneous contact with a pair or triple of object edges. This enumeration can be done as proposed by Wallack and Canny (1997) and Brost and Goldberg (1996) in the context of modular fixturing:

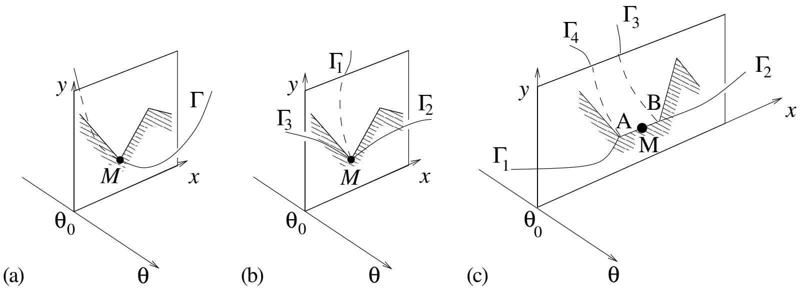

Fig. 3. Configuration space analysis of the Pachinko machine. (a)-(c): equilibria (minima) in configuration space. See text for details.

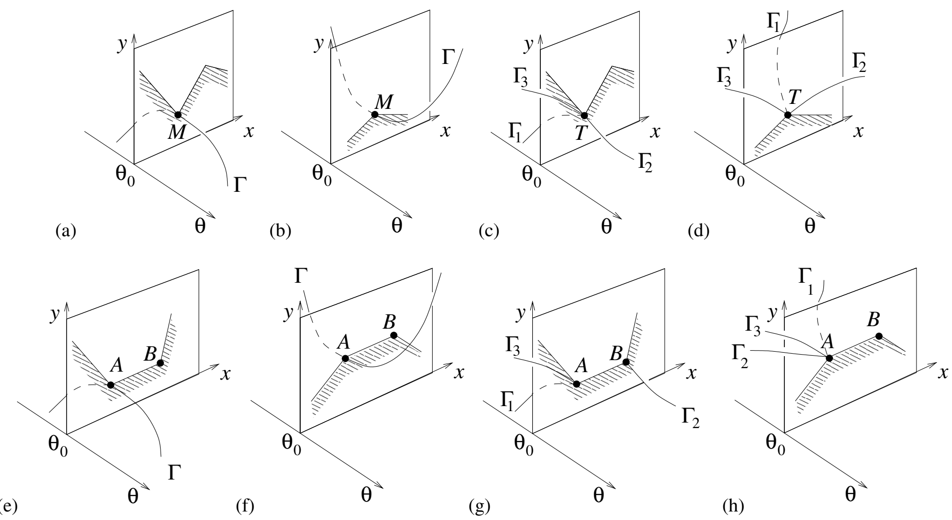

Fig. 4. Saddle point configurations for the Pachinko machine that correspond to escape configurations. See text for details.

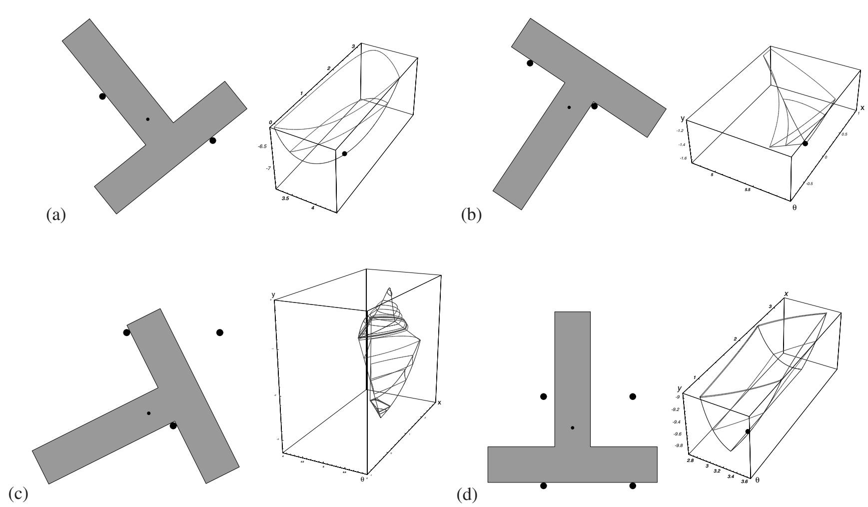

Fig. 5. Capture regions associated with two pins (a-b), three pins (c), and four pins (d). Note that the T-shaped object is caged by the pins in (c), independently of gravity.

if we consider a pair of pins in contact with two object edges, the minimum and maximum distances between the edges constrain the second pin to lie inside a circular shell centered at the first pin. All integer positions inside this shell can be enumerated in optimal time proportional to their number by a scan-line conversion algorithm. Likewise, when three pins are used, the position of the third pin is constrained to lie in the intersection of two circular shells centered at the first and second pins (Fig. 6). Once again, all integer positions in this intersection can be enumerated in optimal time by a scan-line conversion algorithm. Once candidate pairs and triplets of pins have been identified together with the associated edges, the set of actual equilibria is computed.

4.2. Construction of the Graph Arcs

Let us assume we have constructed the nodes E_i ($i = 1, \ldots, n$) of the transition graphs. We construct all the arcs pointing to the node E_i in two steps:

1. Identify every equilibrium configuration E_j whose orientation belongs to the θ range of the capture region of E_i. Equilibrium configurations are stored in a binary search tree according to their orientation and can be rapidly identified.
2. For each selected equilibrium E_j, find the θ-section of the capture region of E_i at the orientation of E_j. Apply a scan-line conversion algorithm to find all the translations of E_j by $\lambda V_1 + \mu V_2$, where V_1 and V_2 are the pin-spacing vectors and λ and μ are integers, which lie in the capture region (Fig. 7). An arc is created from E_j to E_i if a valid translation exists; all valid translations are stored along with the arc. If no valid translations exist, the nodes are not linked.

4.3. Implementation

A naive implementation of the algorithm described above has been performed in C++ (Blind 1999). The computation of the transition graph of the T-shaped object shown in Figure 5 currently takes 3 hours and 45 minutes on a 233 MHz Pentium II. The graph has 1013 nodes and 14,889 arcs and is shown in Figure 8. The connectivity of the graph depends on the object shape and the pin spacing. Equilibrium configurations where the pins are in contact with convex vertices of the part are not used as nodes of the graph.

5. Experiments

We have built a prototype Pachinko machine with 68 pins. A front view is shown in Figure 1, and a side view is shown in Figure 9. The pins are the metal shafts of pneumatic cylinders and are mounted on a Plexiglas plate slightly tilted away from the vertical. The pneumatic cylinders are actuated by solenoid valves that can be individually controlled by a PC.

A sample 10-step manipulation sequence is shown in Figure 10. See also Blind et al. (2000) for a video of an earlier implementation with fewer pins. Our experiments indicate that the capture region computation is conservative; that is, the computed capture regions are smaller than experimentally observed capture regions. Occasionally, plans failed when a part got stuck in an incorrect configuration due to friction between the part and the support plate. This can be avoided by

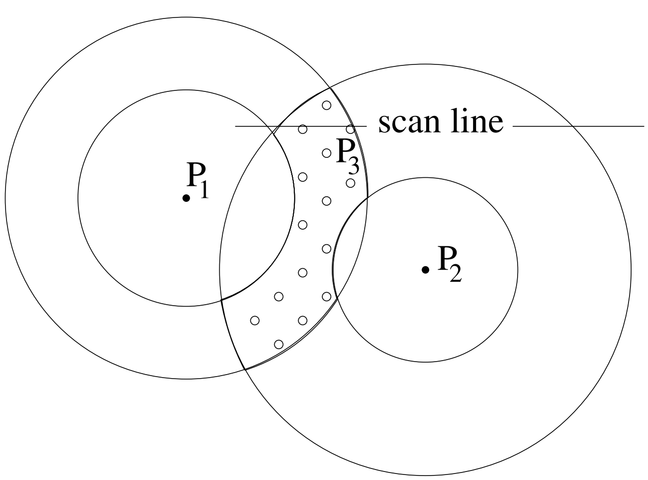

Fig. 6. The positions P_1 and P_2 of two pins on a grid constrain the position P_3 of the third pin to lie at integer points in the intersection of two circular shells.

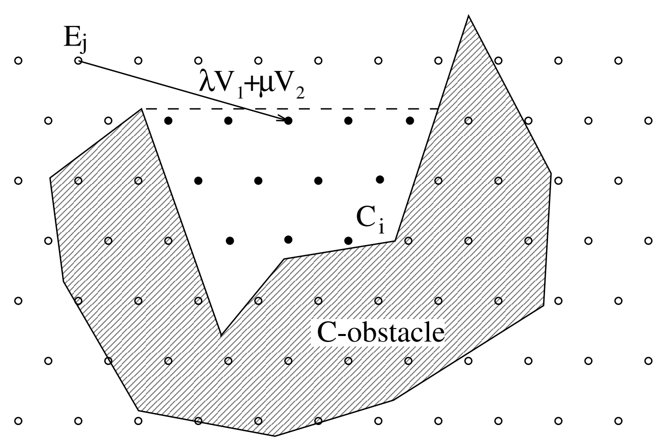

Fig. 7. Finding the translations that bring the equilibrium E_j inside the capture region C_i of E_i. Here C_i is delimited by the (shaded) C-obstacle below and a maximal height (dashed line) above.

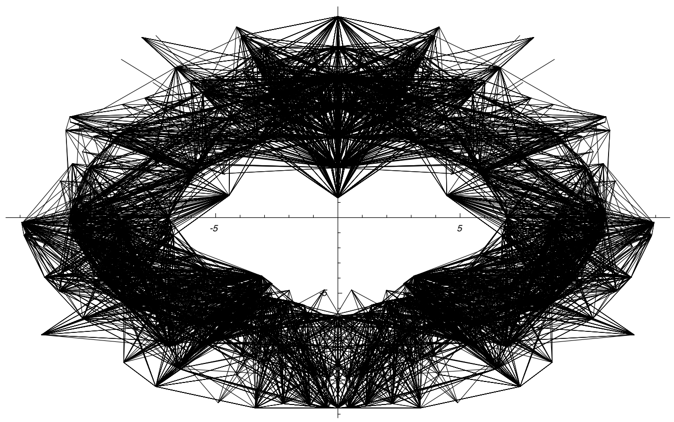

Fig. 8. Transition graph of the T-shaped object. The equilibria are depicted in polar coordinates using y and θ.

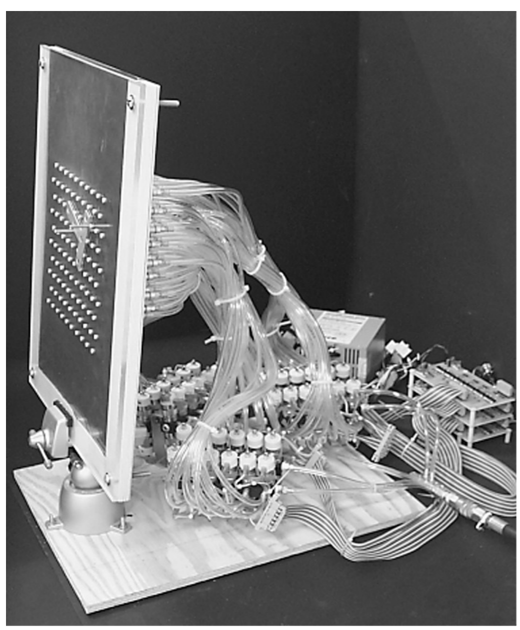

Fig. 9. A side view of the Pachinko machine, showing the solenoid valves that activate the cylinders.

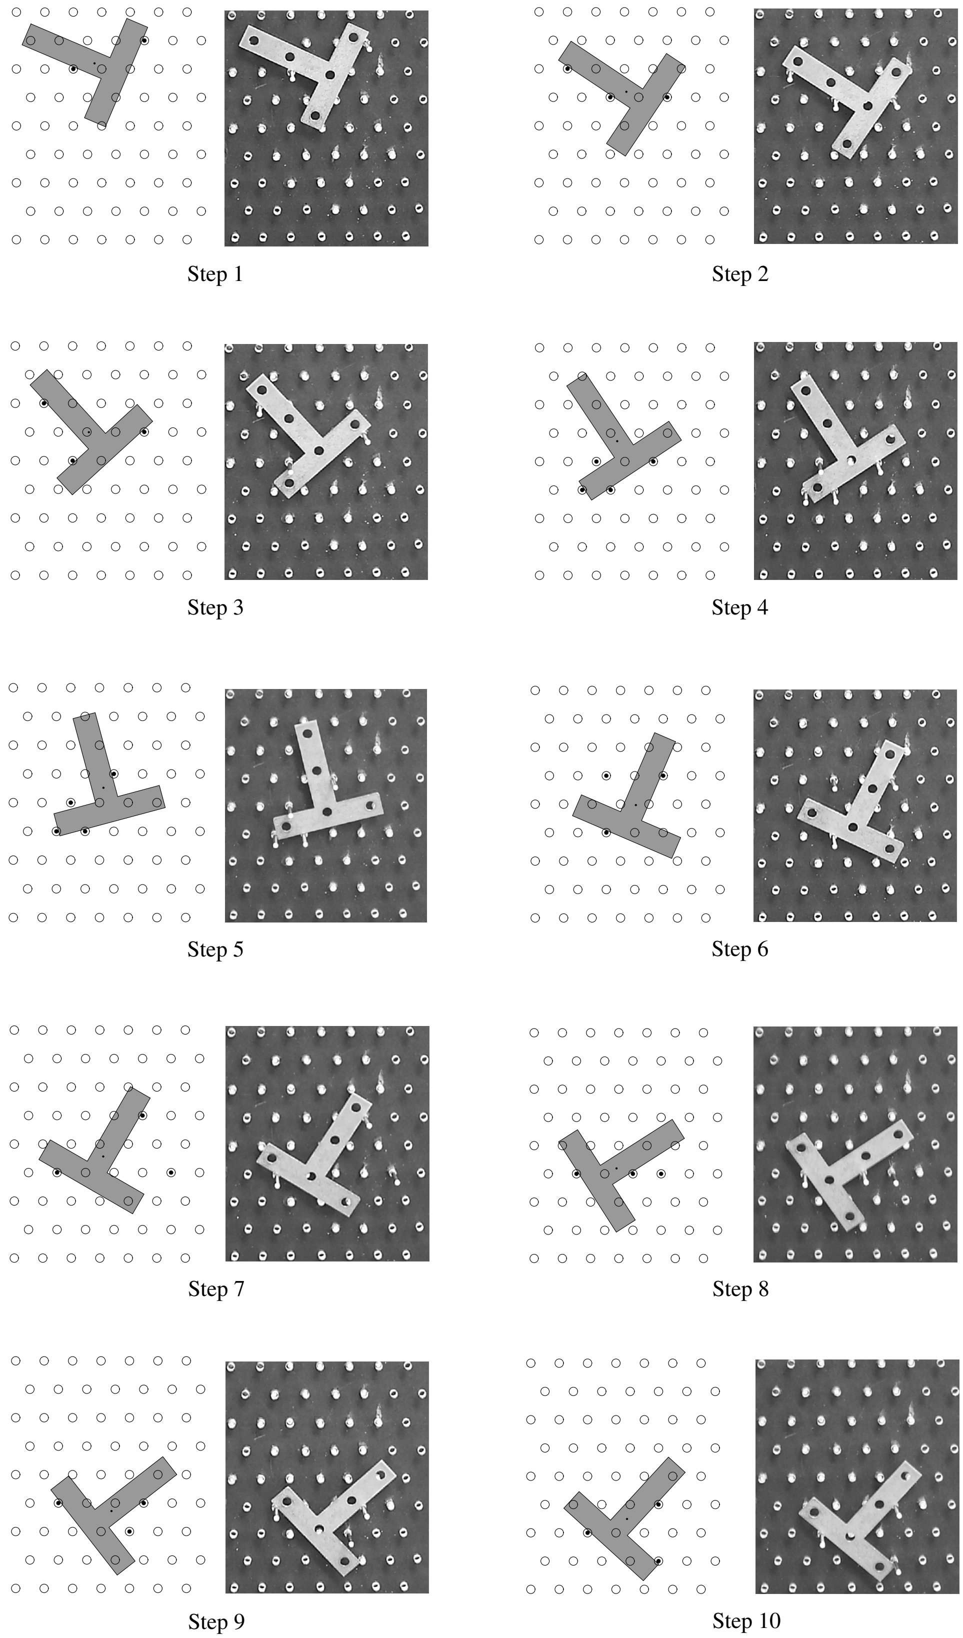

Fig. 10. A 10-step manipulation sequence for the T-shaped object

selecting low-friction support plate material or by introducing small vibrations perpendicular to the support plate.

6. Future Work

Our long-term goal is to develop a parts feeder that can sort parts of different shapes automatically and handle complex 3-D part geometries. There are several directions for future work. We currently assume that the Pachinko machine knows the configuration of the part when initially caught. We plan to add an array of optical proximity sensors to the device to detect the set of possible initial part configurations (as in Jia and Erdmann 1996, for example) and to generate plans that work despite ambiguity in part configuration. We would also like to generate completely sensorless pin activation sequences that can transform a set of initial part configurations to the same goal configuration. Important practical challenges include improving graph connectivity by less conservative computation of the capture regions and incorporating the effects of friction between the part and support plate. The Pachinko machine can potentially be used for part recognition and sorting and may even enable product assembly by coordinated motions of contacting parts. Developing techniques to plan trajectories that permit fully dynamic part motions, without intermediate equilibrium configurations, is another interesting direction for future work.

Acknowledgments

This research was conducted while Sebastien Blind, Christopher McCullough, and Srinivas Akella were with the Beckman Institute for Advanced Science and Technology. It was supported in part by the National Science Foundation under grants IRI-9634393 and IRI-9907009, an NSF REU grant, a UIUC Research Board grant, and the Beckman Institute.

References

Abell, T., and Erdmann, M. A. 1995. Stably supported rotations of a planar polygon with two frictionless contacts. *IEEE/RSJ International Conference on Robots and Systems* 3:411–418.

Akella, S., Huang, W. H., Lynch, K. M., and Mason, M. T. 1996. Planar manipulation on a conveyor with a one joint robot. In *Robotics Research: The Seventh International Symposium*, ed. G. Giralt and G. Hirzinger, 265–276. London: Springer.

Berretty, R.-P., Goldberg, K., Cheung, L., Overmars, M. H., Smith, G., and van der Stappen, A. F. 1999. Trap design for vibratory bowl feeders. *IEEE International Conference on Robotics and Automation*, pp. 2558–2563, Detroit, MI, May.

Blind, S. J. 1999. Automated programming of a reconfigurable feeder. Master's thesis, Department of Computer Science, University of Illinois at Urbana–Champaign.

Blind, S. J., McCullough, C. C., Akella, S., and Ponce, J. 2000a. Manipulating parts with an array of pins: A method and a machine. In *Robotics Research: The Ninth International Symposium*, ed. J. M. Hollerbach and D. E. Koditschek, pp. 123–130. London: Springer-Verlag.

Blind, S. J., McCullough, C. C., Akella, S., and Ponce, J. 2000b. A reconfigurable parts feeder with an array of pins. *Video Proceedings, IEEE International Conference on Robotics and Automation*, San Francisco, April.

Böhringer, K. F., Bhatt, V., and Goldberg, K. Y. 1995. Sensorless manipulation using transverse vibrations of a plate. *IEEE International Conference on Robotics and Automation*, pp. 1989–1996, Nagoya, Japan, May.

Böhringer, K. F., Donald, B. R., Kavraki, L. E., and Lamiraux, F. 2000. A single universal force field can uniquely pose any part up to symmetry. In *Robotics Research: The Ninth International Symposium*, ed. J. M. Hollerbach and D. E. Koditschek, pp. 395–402. London: Springer-Verlag.

Böhringer, K. F., Donald, B. R., Mihailovich, R., and MacDonald, N. C. 1994. Sensorless manipulation using massively parallel microfabricated actuator arrays. *IEEE International Conference on Robotics and Automation*, pp. 826–833, San Diego, CA, May.

Brost, R. C. 1988. Automatic grasp planning in the presence of uncertainty. *International Journal of Robotics Research* 7(1):3–17.

Brost, R. C. 1991. Analysis and planning of planar manipulation tasks. Ph.D. thesis, Carnegie Mellon University School of Computer Science. Technical Report CMU-CS-91-149.

Brost, R. C. 1992. Dynamic analysis of planar manipulation tasks. *IEEE International Conference on Robotics and Automation*, pp. 2247–2254, Nice, France, May.

Brost, R. C., and Goldberg, K. Y. 1996. A complete algorithm for designing planar fixtures using modular components. *IEEE Transactions on Robotics and Automation* 12(1):31–46.

Caine, M. 1994. The design of shape interactions using motion constraints. *IEEE International Conference on Robotics and Automation*, pp. 366–371, San Diego, CA, May.

Canny, J. F. 1989. On computability of fine motion plans. *IEEE International Conference on Robotics and Automation*, pp. 177–182, Scottsdale, AZ.

Canny, J. F., and Goldberg, K. Y. 1994. RISC for industrial robotics: Recent results and open problems. *IEEE International Conference on Robotics and Automation*, pp. 1951–1958, San Diego, CA, May.

Donald, B. R. 1988. A geometric approach to error detection and recovery for robot motion planning with uncertainty. *Artificial Intelligence* 37(1-3):223–271.

Erdmann, M. 1984. On motion planning with uncertainty. Artificial Intelligence Laboratory AI-TR-810, MIT.

Erdmann, M., and Mason, M. T. 1988. An exploration of sensorless manipulation. *IEEE Journal of Robotics and Automation* 4(4):369–379.

Farahat, A. O., Stiller, P. F., and Trinkle, J. C. 1995. On the geometry of contact formation cells for systems of polygons. *IEEE Transactions on Robotics and Automation* 11(4):522–536.

Fearing, R. S. 1986. Simplified grasping and manipulation with dextrous robot hands. *IEEE Journal of Robotics and Automation* RA-2(4):188–195.

Goldberg, K. Y. 1993. Orienting polygonal parts without sensors. *Algorithmica* 10(2):201–225.

Goresky, M., and Macpherson, R. 1980. *Stratified Morse Theory*. New York: Springer-Verlag.

Inoue, H. 1974. Force feedback in precise assembly tasks. AI Lab. AIM-308, MIT, 1974.

Jia, Y.-B., and Erdmann, M. 1996. Geometric sensing of known planar shapes. *International Journal of Robotics Research* 15(4):365–392.

Kriegman, D. J. 1997. Let them fall where they may: Capture regions of curved objects and polyhedra. *International Journal of Robotics Research* 16(4):448–472.

Krishnasamy, J., Jakiela, M. J., and Whitney, D. E. 1996. Mechanics of vibration-assisted entrapment with application to design. *IEEE International Conference on Robotics and Automation*, pp. 838–845, Minneapolis, MN, April.

Latombe, J.-C. 1991. *Robot Motion Planning*. Norwell, MA: Kluwer Academic.

Leveroni, S., and Salisbury, K. 1996. Reorienting objects with a robot hand using grasp gaits. In *Robotics Research: The Seventh International Symposium*, ed. G. Giralt and G. Hirzinger, 39–51. London: Springer.

Lozano-Pérez, T., Mason, M., and Taylor, R. 1984. Automatic synthesis of fine motion strategies for robots. *International Journal of Robotics Research* 3(1):3–24.

Luntz, J. E., Messner, W., and Choset, H. 1997. Parcel manipulation and dynamics with a distributed actuator array: The virtual vehicle. *IEEE International Conference on Robotics and Automation*, pp. 1541–1546, Albuquerque, NM, April.

Lynch, K. M. 1999. Toppling manipulation. *IEEE International Conference on Robotics and Automation*, pp. 2551–2557, Detroit, MI, May.

Mani, M., and Wilson, W.R.D. 1985. A programmable orienting system for flat parts. *North American Manufacturing Research Conference XIII*.

Mason, M. T. 1986. Mechanics and planning of manipulator pushing operations. *International Journal of Robotics Research* 5(3):53–71.

Mason, R., Rimon, E., and Burdick, J. 1995. The stability of heavy objects with multiple contacts. *IEEE International Conference on Robotics and Automation*, pp. 439–445, Nagoya, Japan, May.

Milnor, J. 1963. *Morse Theory*. Princeton, NJ: Princeton University Press.

Peshkin, M. A., and Sanderson, A. C. 1988. Planning robotic manipulation strategies for workpieces that slide. *IEEE Journal of Robotics and Automation* 4(5):524–531.

Rao, A., and Goldberg, K. 1995. Manipulating algebraic parts in the plane. *IEEE Transactions on Robotics and Automation* 11(4):598–602.

Rao, A., Kriegman, D. J., and Goldberg, K. 1996. Complete algorithms for reorienting polyhedral parts using a pivoting gripper. *IEEE Transactions on Robotics and Automation* 12(2):331–342.

Reznik, D., and Canny, J. 1998. A flat plate is a universal planar manipulator. *IEEE International Conference on Robotics and Automation*, pp. 1471–1477, Leuven, Belgium, May.

Rimon, E., and Blake, A. 1996. Caging 2D bodies by one-parameter two-fingered gripping systems. *IEEE International Conference on Robotics and Automation*, pp. 1458–1464, Minneapolis, MN, April.

Shirai, M., and Saito, A. 1989. Parts supply in Sony's general-purpose assembly system "SMART." *Japanese Journal of Advanced Automation Technology* 1:108–111.

Sudsang, A., Ponce, J., and Srinivasa, N. 1998. Grasping and in-hand manipulation: Experiments with a reconfigurable gripper. *Advanced Robotics* 12(5):509–533.

Sudsang, A., Ponce, J., and Srinivasa, N. 2000. Grasping and in-hand manipulation: Geometry and algorithms. *Algorithmica* 26:466–493.

Wallack, A. S., and Canny, J. F. 1997. Planning for modular and hybrid fixtures. *Algorithmica* 19:40–60.

Wentink, C., van der Stappen, A. F., and Overmars, M. 1997. Algorithms for fixture design. In *Algorithms for Robotic Motion and Manipulation*, ed. J.-P. Laumond and M. Overmars, pp. 321–346. Wellesley, MA: A. K. Peters.

Whitney, D. E. 1982. Quasi-static assembly of compliantly supported rigid parts. *ASME Journal of Dynamic Systems, Measurement, and Control* 104:65–77.

Wiegley, J., Goldberg, K. Y., Peshkin, M., and Brokowski, M. 1997. A complete algorithm for designing passive fences to orient parts. *Assembly Automation* 17(2):129–136.