

Dynamic Manipulation

Matthew T. Mason

Kevin M. Lynch

The Robotics Institute
and School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

This paper describes some preliminary work on dynamic manipulation—some examples, a definition, and analysis of throwing a club.

1 Introduction

Robotic manipulators usually employ strategies that avoid or minimize the effects of task dynamics. For example, pick-and-place operations can be programmed based on a purely kinematic view of the task. The programmer assumes that a grasped object tracks the position of the hand, while all other objects remain fixed. Planning consists of finding motions that accomplish the end goal while avoiding collisions. Dynamics does not enter the process at all, except perhaps when tuning the robot's position control system.

In other cases, however, a purely kinematic approach may be inefficient or perhaps even impossible. Throwing a ball, juggling, and balancing an inverted pendulum are all examples of tasks that are dynamic by their very nature. There are many other tasks where dynamics may be greatly advantageous, though not absolutely necessary. For example, a kinematic grasping strategy is very slow, requiring the robot to stop near an object. A dynamic strategy can be much faster, by allowing the robot to snatch an object by sweeping through its position.

Dynamic manipulation and kinematic manipulation are part of a taxonomy of manipulation, which mimics the standard progression of mechanics textbooks:

- Kinematic manipulation. An operation that can be analyzed using only kinematics.
- Static manipulation. An operation that can be analyzed using only kinematics and static forces.
- Quasi-static manipulation. An operation that can be analyzed using only kinematics, static forces, and quasi-static forces (such as frictional forces at sliding contacts).
- Dynamic manipulation. An operation that can be analyzed using kinematics, static and quasi-static forces, and forces of acceleration.

Note that each class includes the preceding classes. We will refer to an operation by the earliest class that includes it.

As an example, we can describe grasping operations at every level of the taxonomy. A kinematic grasp occurs if we configure the hand so that relative motion is prevented by simple kinematic

constraint. A static grasp occurs, for example, if a cube is grasped in a frictional parallel-jaw gripper, since the frictional forces must be included to verify stability. A quasi-static grasp might correspond to pushing a sofa across the floor—the sofa retains contact with the effector to balance the forces of sliding friction. A *dynamic grasp* uses forces of acceleration (in addition to kinematic, static, and quasi-static mechanisms) to hold the object in the effector. For example, when one slaps a coin down onto a table, the downward acceleration of the palm holds the coin against the palm, a technique which we will call *dynamic closure*.

Our definition of dynamic manipulation leans heavily on earlier discussions of dynamic dexterity and dynamic robot tasks by Hodgins and Raibert [10], Koditschek [11], Sakaguchi, Masutani, and Miyazaki [17], and Schaal, Atkeson, and Botros [18]. Our definition of dynamic manipulation is an attempt to identify those tasks where the dynamics is significant during analysis and planning. Thus we might exclude robot motions employing kinematic grasps, even though the actual dynamic forces may be large. In short, does the method use the dynamics or merely tolerate the dynamics? This paper and the works cited above address tasks where the dynamics must be actively exploited.

Why should we care about dynamic manipulation?

- Kinematic manipulation is often very slow. Every object motion must be accompanied by a robot hand. To pick, place, or assemble, the hand must be brought to a stop. Dynamic manipulation can be much faster.
- Dynamic manipulation increases the repertoire of actions available to manipulators. Thus a given manipulator can do more. It can handle larger loads, cause motions outside its workspace, and so forth. By the same token, increasing the repertoire of available actions also means that a fixed task might be accomplished by a simpler manipulator.
- Dynamic manipulation may be especially important in space. Not only are dynamics more difficult to neglect, but dynamic manipulation can save in the complexity and the mass of the robot.

The rest of the paper consists of a discussion of previous work on dynamic manipulation, analysis of a simple method of throwing a club, and a concluding discussion.

2 Examples

We begin with a brief survey of dynamic operations.

Throwing. The most obvious throwing method is to plan an effector trajectory ending at the desired release position and velocity, at which point the effector releases the projectile. This requires a transition from a completely immobilizing grasp to complete freedom, simultaneous in all motion freedoms. It also requires that the effector be capable of achieving the speeds desired of the projectile at the release.

There are other options, however, which involve a sequential release of the object's degrees-of-freedom. A baseball pitcher can impart a large angular velocity to a baseball by allowing it to roll out of his hand. It is not necessary for the hand itself to attain this angular velocity. As another example, consider Claude Shannon's juggling machine, which juggles three balls using two "hands." (This machine is apparently the first juggling machine. Shannon has not described it in print, but see [18] for a description of Shannon's machine and a variant machine that juggles five balls.) Shannon's machine consists of two cups, padded with an energy-absorbing material, mounted at either end of a roughly horizontal rocker arm. The arm oscillates about its center. Each cup is mounted so that at the zenith of its travel, the ball rolls out of the cup, falls to a drumhead below, and bounces into the opposite cup, which is near its nadir. The throwing motion is simple, because the hand does not have to produce precisely the desired motion of the ball, nor is there an elaborate mechanism to release the ball at precisely the right time.

Hirofumi Miura has constructed robots that can play a ball-in-cup (*kendama*) game and spin a top (*koma*) by throwing it with a string [13]. Throwing a top might be viewed as an instance of letting the projectile roll off the effector, if we view the string as part of the effector.

Other examples of robot throwing are described in Aboaf, Atkeson, and Reinkensmeyer [1], which addresses learning to throw a ball more accurately, and Slotine [20], demonstrating throwing and catching of a ball.

Catching. In theory, if the system were energy-conserving, catching would be identical to throwing with time reversed. However, because of uncertainty in the arrival state of the projectile, energy-dissipation is crucial in catching. It is sometimes impractical to match the effector trajectory to the projectile trajectory and then just close the fingers. Instead, collisions are inevitable.

One approach is to choose an effector shape and trajectory so that a convergent sequence of collisions ends with the object motionless in contact with the effector. In general this is a very difficult problem [22]. But it becomes more practical if the collisions are inelastic, so that the sequence might consist of only one or two collisions.

Again, the Shannon juggler provides a good example. Because of the ball's symmetry, a convergent sequence of collisions can be obtained by using a simple cup-shaped effector with energy-absorbing padding. A similar strategy is used by Sakaguchi, Masutani, and Miyazaki [14, 17] to obtain robot juggling of one or two balls with a single robot hand. They have also programmed the robot to play the ball-in-cup game.

Inelastic collisions can also simplify catching of less symmetric objects. Consider the devil sticking robot described by Schaal, Atkeson, and Botros [18]. (Devil sticking, in its simplest form, consists of using two sticks, one in each hand, to pass a third stick back and forth.) The effector stick contacts the devil stick at its center of percussion, halting the effector stick and storing its energy momentarily at a springy joint. The collision is effectively inelastic, resulting in a catch. The energy is then transferred back

to the devil stick, throwing it to the other effector stick.

Dynamic closure. The devil sticking robot also uses dynamic closure: the use of acceleration forces to ensure that a contact constraint remains active. The devil stick remains in contact with the effector stick for a substantial time because of the acceleration forces. A more common example, slapping a coin down onto a table, is described above.

Dynamic closure appears to play a prominent role in SONY's APOS parts-orienting system [9, 19]. This system uses a pallet with specially shaped depressions to capture parts in the desired orientation. The pallet is tilted and vibrated so parts slide down the face of the pallet, occasionally falling into one of the depressions. The vibration is designed so that the dynamic load will hold a part that happens to be in the correct orientation, but eject a part in the wrong orientation.

Snatching. By snatching, we mean a dynamic pick operation. The effector might start in contact with an object, then accelerate through the object, controlling the object via dynamic closure. Or, even better, the hand might pass through the object's initial position without stopping, using collisions followed by dynamic closure. The operation is mechanically identical to catching, viewed from the object frame instead of the robot frame. The difference is that a much better estimate of the arrival state is possible, and the robot doesn't have to deal with a spinning object.

Batting. Batting means generating a single collision between effector and projectile in order to redirect the projectile. It combines catching and throwing in a single collision. Bühler and Koditschek [6] describe a machine using a single bar pivoting about its center to bat pucks sliding on an inclined plane. The goal is to achieve a cyclic bouncing of the puck, with stable height, at a stable location along the bar. Bühler and Koditschek found a very simple feedback law that will stabilize one or two bouncing pucks simultaneously. Aboaf, Atkeson, and Reinkensmeyer [2] and Rizzi and Koditschek [16] describe batting a ball in three dimensions.

Andersson [3] built a machine to play ping-pong, which is an adversarial form of batting. Andersson's planner incorporated a model of ball flight and impact, and used these models in a fairly conventional way to plan a nominal trajectory for the paddle. This nominal trajectory was then refined by iterated simulations, with concurrent adjustment of goals as better estimates of the ball's motion became available.

Hopping and running can also be viewed as a form of batting. Raibert and his colleagues use sophisticated dynamic models to design simple feedback control systems [15].

Unactuated freedoms. Dynamic coupling can be used to control the motion of an unactuated joint. This principle was demonstrated in early work on balancing an inverted pendulum [8], and more recently in the control of running and walking [15], control of flexible beams, and control of linkages with unactuated joints [4, 21]. In McGeer's walking machines [12], all joints are unactuated.

Placing and assembly. When we throw an object on a table, the table has to catch it. Unfortunately, the table has no sensors, motors, or computers. It is up to us to choose the object's trajectory so that the intrinsic mechanics of the task do the right thing. In some cases it may be possible to tune these intrinsic mechanics offline to achieve more effective placing or assembly. Asada and Kakumoto [5] applied this principle to analyze and

tune the Dynamic Remote Center Compliance for high speed peg insertions.

2.1 Design, control, and planning

Each example of dynamic manipulation described above required numerous design decisions, including the design of shape and mechanical properties, actuation and perception, control system design, and motion planning. In some cases the desired behavior can be obtained entirely from the intrinsic mechanics of the device, as with Shannon’s juggler and McGeer’s passive walking machines. In other cases, notably Raibert’s hoppers and Koditschek’s jugglers, the intrinsic passive mechanics is integrated with simple feedback strategies to obtain the desired system behavior. We also have examples where the desired behavior requires high-level motion planning, notably Andersson’s ping-pong player. An alternative approach to high-level motion planning is provided by Witkin and Kass [23], who formulate a hopping problem as a two-point boundary value problem and develop a relaxation method to obtain motor torque programs.

In most cases described above, decisions are guided by a model of the intended dynamic behavior. Sometimes the model is quite explicit, as with the planning process of Andersson’s ping-pong player. But even in cases where an explicit model does not exist in the robot, such a model may still guide design and control decisions.

We should not neglect the potential role of learning in addressing design and control decisions. Chris Atkeson’s group has explored learning in the context of several dynamic tasks [1, 2, 18]. Christiansen, Mitchell, and Mason [7] demonstrated learning in a dynamic parts-orienting task. Adaptive control and parameter estimation have often focused on dynamic aspects of manipulation.

With a learning robot, models play a different role in design, control, and planning. Let us view a model as a body of information about the mechanics of a task. Then learning is a mechanism for the robot to obtain such information from observations of the task domain, reducing the requirements for prior information. While it may appear that learning can reduce reliance on models, we would argue instead that it merely shifts some of the modeling burden from the human to the robot. A non-learning robot relies on human-constructed models, which generalize empirical data gathered by humans. A learning robot has to gather its own data and make its own generalizations. Implicitly or explicitly, the robot relies mostly on human-constructed models, with the potential of making some incremental refinements of these models via learning.

3 Throwing a club

As a first step, we analyzed a simple example of dynamic manipulation: club-throwing using dynamic closure. This example was inspired by Michael Kass’ ability to catch and throw a club using his foot (see Figure 1). This process consists of three distinct phases: the catch, the carry, and the release. In this section we will focus on the carry and release phases, which comprise a throw. During the throw, the foot accelerates along an arc and then decelerates, tossing the club into the juggling pattern. This process is not kinematic manipulation, because it is impossible to close the foot around the club. Neither is it static nor quasi-static manipulation, because forces of acceleration are used to maintain control of the club during the carry. This throwing method is an example of exploiting dynamics in order to increase the capability of a manipulator.

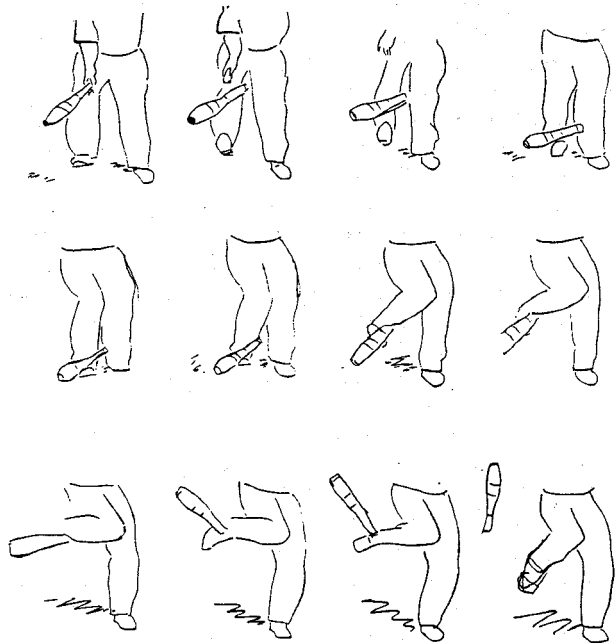


Figure 1: Throwing a club with a foot.

During the carry, the club tracks the motion of the foot by means of a dynamic grasp. The contacts between the foot and the club, along with Coulomb’s law of friction, define a friction cone in wrench (force/torque) space comprising the set of wrenches that can be applied to the club. Given the mass parameters of the club, this convex friction cone can be mapped to a convex cone of possible club accelerations in acceleration/angular acceleration space. This acceleration cone is fixed with respect to the club and in time, and it defines a constraint on the foot accelerations. Provided the acceleration of the foot (minus the gravitational acceleration) remains inside this cone, a dynamic grasp is possible.

We will derive the conditions for the successful execution of both the carry and the release for a simplified model of Kass’ method. The club is modeled as a thin cylinder of uniform mass. The club is thrown by the foot of a single degree-of-freedom leg, which rotates in the vertical plane with a constant angular acceleration α about a fixed axis. A coordinate frame is attached to the club at its center of mass, with the x -axis pointing in the direction of motion of the center of mass and the y -axis pointing toward the axis of rotation. The angle between the club and the x -axis is denoted ψ . The distance from the center of mass to the axis of rotation is r . See Figure 2.

The center of mass sweeps out a 90 degree arc during the carry, from directly below the rotation axis ($\theta = 0$ degrees) to directly to the right ($\theta = 90$ degrees). When θ reaches 90 degrees, the foot is decelerated and the club is thrown with a vertical linear velocity.

The club’s angular acceleration, velocity, and position may be written as functions of the time of the carry:

$$\ddot{\theta}(t) = \alpha \quad (1)$$

$$\dot{\theta}(t) = \alpha t \quad (2)$$

$$\theta(t) = \frac{1}{2} \alpha t^2. \quad (3)$$

Defining n to be the desired number of club rotations from the time of the launch to the time the club returns to the launch position, the trajectory of the club during the carry and flight is




Figure 2: Notation for club throwing.

characterized by the following equations:

$$\alpha = \frac{gn}{r} \quad (4)$$

$$T_c = \sqrt{\frac{\pi r}{gn}} \quad (5)$$

$$\dot{\theta}_r = \sqrt{\frac{gn\pi}{r}} \quad (6)$$

$$v_r = \sqrt{gn\pi r} \quad (7)$$

$$T_f = 2\sqrt{\frac{n\pi r}{g}} \quad (8)$$

$$h = \frac{n\pi r}{2} \quad (9)$$

where T_c is the time elapsed during the carry, $\dot{\theta}_r$ is the angular velocity at the release, v_r is the linear velocity at the release, T_f is the time of flight of the club, and h is the maximum height of the center of mass above the release point.

3.1 Carrying conditions

The forces (measured in the coordinate system fixed on the club) which must be applied to the club during the carry are given by the following equations:

$$\begin{aligned} f_z(\theta) &= m\alpha r + mg \sin \theta \\ &= mgn + mg \sin \theta \end{aligned} \quad (10)$$

$$\begin{aligned} f_y(\theta) &= m\dot{\theta}^2 r + mg \cos \theta \\ &= 2\theta mgn + mg \cos \theta \end{aligned} \quad (11)$$

$$\begin{aligned} \tau(\theta) &= m\rho^2 \alpha \\ &= \frac{m\rho^2 gn}{r} \end{aligned} \quad (12)$$

where m is the mass of the club and ρ is its radius of gyration. The angle ψ and the contacts between the foot and the club should be chosen to keep these forces inside the wrench friction cone.

3.2 Release conditions

Here we will consider the simplest type of release: simultaneous breaking of all contacts between the foot and the club. For this to be possible, each contact point on the foot must be able to accelerate away from the club during the release. Equivalently,




Figure 3: Forces applied to the club during the carry of a one rotation throw. \hat{n} is normal to the club.




Figure 4: Constraints on the contact locations: (a) finite club length; (b) release constraints; (c) the combined constraints and the normal forces and torques during the trajectory (normalized to the unit circle).

each contact normal (into the club) must have a positive angular sense about the axis of rotation. If this condition is satisfied, the club instantaneously breaks contact for sufficiently large decelerations of the foot. (Note that this is only a necessary local condition for simultaneous release. The complete paths of the foot and the club as it spins away must be considered to ensure a clean release.)

3.3 Choosing contacts

Given a club of length 0.5, a club angle $\psi = -10$ degrees, a leg such that $r = 1$, and a goal to throw the club so that it rotates once in the air ($n = 1$), how should the foot contact the club?

Figure 3 shows the evolution of the forces which must be applied to the club during the carry. The coefficient of friction between the foot and the club must be large enough to include these forces in the total wrench friction cone. In this case, the coefficient of friction must be at least 0.7 for a successful carry.

The choice of contact points on the club is constrained by the length of the club, the release condition, and the forces and torques which must be applied to the club during the carry. To visualize these constraints, we can plot them in a force-torque space, where the force axis is the force in the direction \hat{n} normal to the club.

The limited club length imposes a constraint on the relative values of the normal contact force and the torque about the center of mass of the club. These constraints are illustrated in Figure 4(a), where the shaded forces are unattainable. The release condition provides an additional constraint, shown in Figure 4(b). These constraints are combined in Figure 4(c). Each point on the unit circle which is not in a shaded region represents a permissible contact. Also plotted are the normal forces and torques (normalized to the unit circle) that must be applied to the




Figure 5: Qualitatively different contact solutions.

object during the carry. We should choose our contacts (points on the unit circle) such that the smallest cone that contains them also contains the required forces and torques.

This example yields two qualitatively different two-contact solutions sets, illustrated in Figure 5. The contacts may be anywhere within the intervals shown. The up-down configuration of Figure 5(a) corresponds roughly to using the shin to provide a contact opposed to the instep of the foot, whereas Figure 5(b) corresponds to using only the instep to throw the club. Under the condition of simultaneous release, the up-down solution exists only in the range $-14.5^\circ \leq \psi \leq 3.4^\circ$.

3.4 Choosing the club angle ψ

We would like to avoid relying on a high coefficient of friction for successful manipulation. In the previous example, a relatively high coefficient of friction was necessary to maintain a dynamic grasp. Instead, we can vary the club angle ψ to make club-throwing robust to low friction.

Figure 6 again shows the forces that must be applied to the club during the carry, but now the angle ψ is chosen to be -38.7 degrees in order to minimize the required coefficient of friction. For this choice of ψ , the carry is successful for friction coefficients as low as 0.11. Figure 6 also indicates that the carry is impossible at slow speeds or with zero gravity for this friction coefficient. The gravitational and centripetal forces sum to produce a force trajectory with a small angular deviation in force space, thus permitting a low coefficient of friction for a properly chosen ψ . Using the ability to choose ψ , Figure 7 shows a graph of the minimum required friction coefficient as a function of the desired number of rotations during the club's flight.

3.5 Implementation

We built a simple club-throwing device to test the strategy outlined above. A picture of the pendular leg is shown in Figure 8. The wooden cylindrical club rests on a foot made of foam-covered wood in between two boards which comprise the leg. The contact between the foot and the club is of the type illustrated in Figure 5(b). The leg is mounted on a pulley which is actuated by falling weights. The leg with the club is statically balanced so that the falling weights provide a constant angular acceleration about the rotation axis. When the leg reaches approximately a horizontal position, it is decelerated by stiff braking springs, and




Figure 6: Choosing ψ to minimize the friction coefficient required during the carry.




Figure 7: The minimum required friction coefficient as a function of the number of club rotations.

the club is thrown with a simultaneous release.

Figure 9 shows a sequence of four frames from a single throw. The club is initially supported by a holder, because the club's starting position on the foot is not statically stable. When the firing pin is pulled, the foot sweeps the club off of the holder and carries it with a dynamic grasp as the leg accelerates. When the leg reaches the release point, it is sharply decelerated by the springs and the club is launched. The throw shown here results in approximately one rotation of the club above the release point.

3.6 Comments

The condition of simultaneous release is unnecessarily restrictive, as it limits the release linear and angular velocities of the club to those achievable by the one degree-of-freedom leg. With an up-down contact configuration, the club could roll about the downward pointing contact as it releases, increasing its angular velocity.

This section provided simple examples of synthesizing manipulator trajectories and contact configurations based on goals stated in terms of a desired number of club rotations. We would




Figure 8: The Pendular Pedipulator.

like to be able to simultaneously design manipulator trajectories and contact configurations for more general cases, perhaps while minimizing functions of the motor torques or the required friction coefficient.

4 Discussion

What is manipulation? One characterization is that it is the problem of achieving a configuration of several objects, using just a few motors. Typically the problem is decomposed by time-multiplexing the motors: the objects are moved one at a time by attaching the motors to them one at a time.

This approach is conservative and inefficient, because it neglects resources, other than motors, that can produce a desired motion. We will try to develop this point a little further. We will propose that the goal is to control the motion freedoms of one or more objects. At any given time, we can list these freedoms, and for each freedom identify the controlling mechanism. The mechanisms we will consider are:

- Kinematic constraint. For example, all motion freedoms of an asymmetric object in an enveloping grasp are constrained kinematically.
- Static friction. For example, planar motion of an object at rest on a table is prevented by static friction.
- Quasi-static force balance. For example, the motion of a peg sliding into a hole at moderate to low velocities is governed by a quasi-static force balance.
- Dynamic: Newton's second law. For example, all motion freedoms of a satellite are governed by Newton's second law, subject only to gravitation forces.

Quite often we are faced with unilateral kinematic constraints, requiring some additional mechanism to prevent the object from moving away from the constraint. The mechanisms literature speaks of "form closure" and "force closure", but since these terms are used somewhat differently in the literature on robot grasping, we will adopt a different terminology.

- Kinematic closure. A second unilateral kinematic constraint is present, so that the two together form a bilateral constraint.
- Static closure. A static force, such as gravity, is applied to the object which can only be balanced by continuing contact with the kinematic constraint.
- Quasi-static closure. A quasi-static force is present, perhaps because the object is being pushed across a surface, which can only be balanced by continuing contact with the kinematic constraint.
- Dynamic closure. Acceleration of the kinematic constraint causes an object acceleration and consequent dynamic load on the object, which can only be balanced by continuing contact with the kinematic constraint.

The different freedoms of an object might be controlled in different ways. For example, an object at rest on a table typically has three degrees of freedom controlled by a kinematic constraint (the table surface against the object bottom) with static closure due to gravity, while the remaining three freedoms are controlled by static friction.

Manipulation consists of phases in which each motion variable is controlled in some particular way, punctuated by events at which a motion variable switches from one control method to another. For example, consider the grasp of a disk on a table by a parallel-jaw gripper. Initially the disk has three freedoms determined by gravity-closed kinematic constraint, and three determined by static friction. Then one finger makes contact at one point of the disk, and the static friction is broken. The disk slides on the table as the finger pushes it. We still have three freedoms determined by gravity-closed kinematic constraint. Assuming that the disk-finger contact does not slip, it gives us one kinematic constraint, closed by the forces of frictional sliding, one freedom fixed by static friction with the pushing finger, and the remaining freedom (rolling along the finger face) determined by a quasi-static force balance. If the second finger now makes contact, then we now have kinematic closure of one freedom (translation perpendicular to finger faces), and the remaining two (rotation; translation tangential to finger faces) determined by static friction.

With this perspective, the notions of picking, placing, grabbing, catching, and throwing can all be viewed similarly. In a picking operation, control of freedoms is progressively transferred from the table to the hand. In placing, the opposite occurs. When handing an object from one hand to another, the processes are the same, but whether to call it a pick or a place depends on which hand is doing the talking. Even a throw might be viewed as a kind of place operation. The hand progressively transfers control of motion freedoms to Newton.

The problems are in the details. How do we arrange an orderly transfer, so that the ultimate goals are served, consonant with restrictions on available information, accuracy of effector motion control, and so forth? Pursuant to that issue, this paper explored the problem of carrying a cylinder using a dynamic grasp, and throwing it. The motion strategy was hand-crafted and tuned using Newtonian rigid-body mechanics and Coulomb friction.

Naturally we would like to develop more general solutions to the problems of dynamic manipulation, rather than handcrafting solutions to each problem. This is a daunting prospect, because dynamic manipulation imposes some costs. For kinematic manipulation, the state of a rigid object requires six parameters. For dynamic manipulation, we must include rate information, giving twelve state parameters per rigid object. Dynamic actions are harder to model than kinematic actions. In general,




Figure 9: A club throw.

dynamic manipulation depends on inertial properties, coefficients of restitution, and other properties that may be difficult to estimate. In short, general dynamic manipulation planning is intractable.

But general kinematic manipulation planning is intractable, too. The solution is not to build a general manipulation system, but rather to build an efficient system that addresses those few special cases of interest, by applying the simplest method available for each problem. For some problems dynamic manipulation is the simplest method.

Acknowledgments

Ideas about parts-feeding go back to conversations with Peter Will, David Grossman, and Tomas Lozano-Perez, and have been sustained by a continuing collaboration with Mike Erdmann. The SONY Factory Automation group, including Akira Kimura and Yoshihiro Kuroki, has been very helpful. Some ideas, including dynamic closure, arose in an extended discussion with Bruce Donald, Mike Erdmann, and Tomas Lozano-Perez. Thanks to Marc Raibert, Dan Koditschek, and Chris Atkeson for valuable discussions. Thanks to Michael Kass for inspiration and for sharing insights. This work is supported by the National Science Foundation under grant IRI-9114208.

References

- [1] E. W. Aboaf, C. G. Atkeson, and D. J. Reinkensmeyer. 1987. Task-level robot learning: Ball throwing. A.I. Memo 1006, Massachusetts Institute of Technology.
- [2] E. W. Aboaf, S. M. Drucker, and C. G. Atkeson. 1989. Task-level robot learning: Juggling a tennis ball more accurately. *Proceedings 1989 IEEE Int Conf Robotics and Automation*, pages 1290-1295, Scottsdale AZ.
- [3] R. L. Andersson. 1989. Understanding and applying a robot ping-pong player's expert controller. *Proceedings 1989 IEEE Int Conf Robotics and Automation*, pages 1284-1289, Scottsdale AZ.
- [4] H. Arai and S. Tachi. 1991. Position control of a manipulator with passive joints using dynamic coupling. *IEEE Transactions on Robotics and Automation*, 7(4):528-534.
- [5] H. Asada and Y. Kakumoto. 1990. The dynamic analysis and design of a high-speed insertion hand using the generalized centroid and virtual mass. *ASME Journal of Dynamic Systems, Measurement, and Control*, 112:646-652.
- [6] M. Bühler and D. E. Koditschek. 1990. From stable to chaotic juggling: Theory, simulation, and experiments. *Proceedings 1990 IEEE Int Conf Robotics and Automation*, pages 1976-1981, Cincinnati OH.
- [7] A. Christiansen, T. M. Mitchell, and M. T. Mason. 1990. Learning reliable manipulation strategies without initial physical models. *Proceedings 1990 IEEE Int Conf Robotics and Automation*, pages 1976-1981, Cincinnati OH.
- [8] D. T. Higdon and R. H. Cannon, Jr. 1963. On the control of unstable multiple-output mechanical systems. *ASME Winter Annual Meeting*.
- [9] H. Hitakawa. 1988. Advanced parts orientation system has wide application. *Assembly Automation*, 8(3):147-150.
- [10] J. K. Hodgins and M. H. Raibert. 1990. Biped gymnastics. *Int J Robotics Research*, 9(2):115-132.
- [11] D. E. Koditschek. 1993. Dynamically dexterous robots. *Robot Control: Dynamics, Motion Planning, and Analysis* M. W. Spong, F. L. Lewis, C. T. Abdallah, editors. New York: IEEE Press.
- [12] T. McGeer. 1990. Passive dynamic walking. *Int J Robotics Research*, 9(2):62-82.
- [13] H. Miura. Private communication.
- [14] F. Miyazaki. 1992. "Motion Planning and Control for a Robot Performer" (videotape).
- [15] M. H. Raibert. 1986. *Legged Robots That Balance*. Cambridge: MIT Press.
- [16] A. A. Rizzi and D. E. Koditschek. 1992. Progress in spatial robot juggling. *Proceedings 1992 IEEE Int Conf Robotics and Automation*, pages 775-780, Nice, France.
- [17] T. Sakaguchi, Y. Masutani, and F. Miyazaki. 1991. A study on juggling task. *IEEE/RSJ International Workshop on Intelligent Robots and Systems IROS '91*, Osaka, Japan, pages 1418-1423.

- [18] S. Schaal, C. G. Atkeson, and S. Botros. 1992. What should be learned? *Proceedings Seventh Yale Workshop on Adaptive and Learning Systems*, pages 199-204.
- [19] M. Shirai and A. Saito. 1989. Parts supply in Sony's general-purpose assembly system "SMART". *Jpn J Adv Automation Tech* 1: 108-111.
- [20] J. Slotine. 1992. "Vision Guided Motion Control" (videotape).
- [21] S. Takashima. 1991. Control of gymnast on a high bar. *IEEE/RSJ International Workshop on Intelligent Robots and Systems IROS '91*, Osaka, Japan, pages 1424-1429.
- [22] Y. Wang and M. T. Mason. 1992. Two dimensional rigid body collisions with friction. *Trans ASME; J Applied Mechanics*. 59(3): 635-642.
- [23] A. Witkin and M. Kass. 1988. Spacetime constraints. *Computer Graphics*, 22(4):159-168.