

Extrinsic Dexterity: In-Hand Manipulation with External Forces

Nikhil Chavan Dafle¹, Alberto Rodriguez¹, Robert Paolini², Bowei Tang², Siddhartha S. Srinivasa², Michael Erdmann², Matthew T. Mason², Ivan Lundberg³, Harald Staab³ and Thomas Fuhlbrigge³

¹ Department of Mechanical Engineering, Massachusetts Institute of Technology

² The Robotics Institute, Carnegie Mellon University

³ ABB Corporate Research Center

(nikhilcd, albertor)@mit.edu, (rpaolini, boweit)@cmu.edu, (siddh, me, matt.mason)@cs.cmu.edu, ivan.lundberg@se.abb.com, (harald.staab, thomas.a.fuhlbrigge)@us.abb.com

Abstract—“In-hand manipulation” is the ability to reposition an object in the hand, for example when adjusting the grasp of a hammer before hammering a nail. The common approach to in-hand manipulation with robotic hands, known as dexterous manipulation [1], is to hold an object within the fingertips of the hand and wiggle the fingers, or walk them along the object’s surface. Dexterous manipulation, however, is just one of the many techniques available to the robot. The robot can also roll the object in the hand by using gravity, or adjust the object’s pose by pressing it against a surface, or if fast enough, it can even toss the object in the air and catch it in a different pose. All these techniques have one thing in common: they rely on resources extrinsic to the hand, either gravity, external contacts or dynamic arm motions. We refer to them as “extrinsic dexterity”.

In this paper we study extrinsic dexterity in the context of regrasp operations, for example when switching from a power to a precision grasp, and we demonstrate that even simple grippers are capable of ample in-hand manipulation. We develop twelve regrasp actions, all open-loop and hand-scripted, and evaluate their effectiveness with over 1200 trials of regrasp and sequences of regraps, for three different objects (see video [2]). The long-term goal of this work is to develop a general repertoire of these behaviors, and to understand how such a repertoire might eventually constitute a general-purpose in-hand manipulation capability.

I. INTRODUCTION

In the robotics research literature “dexterity” often refers to the manipulation of an object in the hand, with the hand. By “extrinsic dexterity” we mean the manipulation of an object *in* the hand, using resources *extrinsic* to the hand. For example, if you tilt your hand so that an object rolls from your fingertips to your palm, you are using gravity, not just your hand. If you press two chopsticks against the tabletop to even them up, you are using the surface of the table and arm motions, not just your hand. If you toss an object upwards from your palm and catch it in your fingertips (Fig. 1), you are using object inertia and arm acceleration, not just your hand. In each case, you can do more, better, or faster by exploiting resources extrinsic to the hand.

Manipulation with extrinsic dexterity stands in contrast to the “dexterous hands” approach, which manipulates the

This work was supported by National Science Foundation [NSF-IIS-0916557], Army Research Laboratory [W911NF-10-2-0016] and ABB. This work does not necessarily reflect the position or the policy of the U.S. Government or ARL. No official endorsement should be inferred.

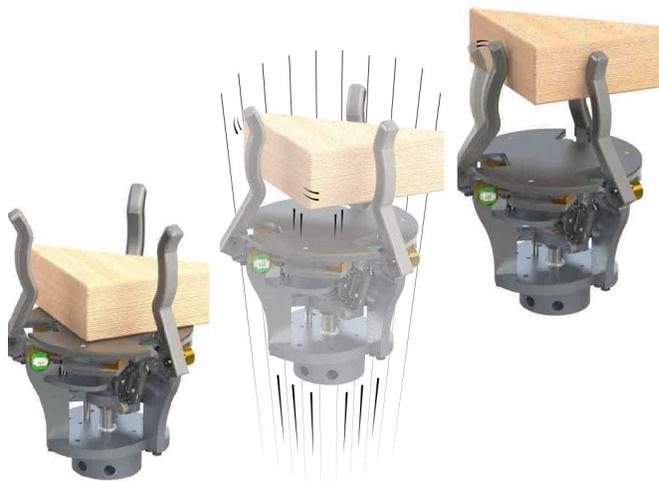


Fig. 1. Throw-to-fingertip regrasp. A high robot arm acceleration throws the object into the air for the fingers to catch.

object through finely controlled fingertip contacts. The dexterous hands approach started with Salisbury’s work [1], and has been thoroughly analyzed and extended since then. (See [3] for an overview of the work up to 2000). Given that it relies on resources intrinsic to the hand, we refer to it as “intrinsic dexterity”.

Extrinsic dexterity complements the intrinsic capabilities of the hand. In particular it is well suited to address tasks that require a significant adjustment of the object in the hand, such as the change from a power to a precision grasp [4] (Fig. 1), traditionally difficult to attain by the dexterous hands approach. In this paper we describe 12 such regraps, all open-loop and hand-coded, and all relying on external resources. The long term goal of this work is to develop a repertoire of regrasp actions sufficient to reach any grasp type from any other in a given grasp taxonomy—a repertoire to *navigate the grasp taxonomy*.

We represent the regrasp capability induced by such a repertoire with a *grasp graph* containing a directed arc for every regrasp (see example in Fig. 2). This paper only addresses a few specific grasps and objects for a given hand. Generalizing these actions so they can be applied across a broad range of grasps, objects, and hands, is an enormous

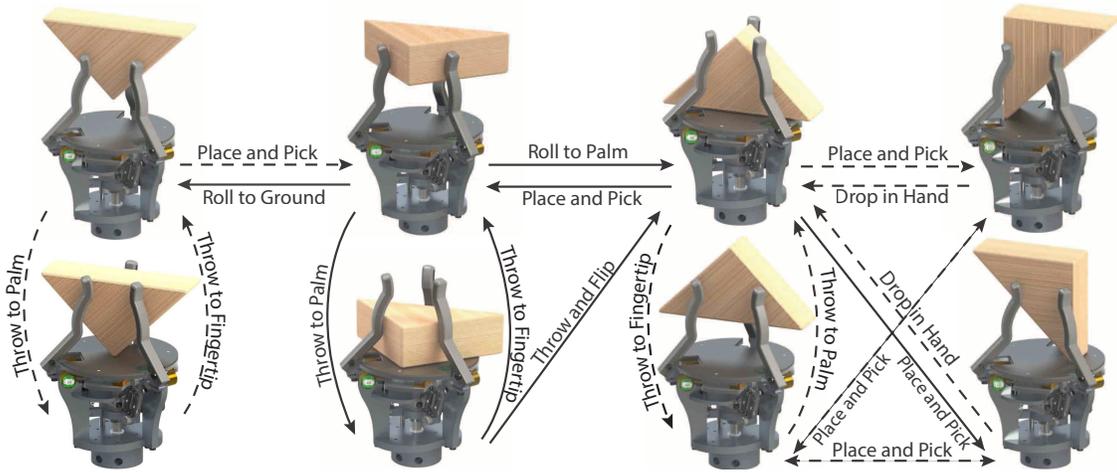


Fig. 2. Grasp graph for a triangular prism. Solid arrows represent implemented regrasps while dashed arrows indicate future ones.

challenge. Nonetheless extrinsic dexterity has advantages that justify this work. Extrinsic dexterity offers an incomparable ability to change grasps. In addition, by exploiting resources extrinsic to the hand, we shift part of the complexity of the manipulator to a possibly modular and adaptable environment. This allows the use of simpler hands which can be cheaper, lighter, smaller, and readily available.

II. RELATED WORK

This paper explores the application of extrinsic dexterity to regrasping objects. The dominant approach to regrasp is to avoid it. Manufacturing automation avoids regrasp through the design of specialized grippers, part feeders, work-holding fixtures, and tools [5]. Avoidance of regrasp is also a common consideration in the design of surgical instruments and procedures [6].

Nonetheless, robotics researchers have studied regrasp for a long time. The importance of regrasp was noted by Paul et al., and is readily apparent in the Instant Insanity demonstration [7, 8], and the Handey project [9, 10]. Both systems fit within the “pick-and-place” paradigm of autonomous manipulation, which rearranges the world by picking and placing one object at a time.

Dexterous manipulation, originally formulated by Salisbury [11, 1, 12] and contrary to the discrete nature of the pick-and-place approach, relies on the application of continuous controlled forces to the object through the fingertips. In Salisbury’s formulation the hand attains local controllability of an object by assuming three fingers in point contact with the object, and three actuators per finger. Such a hand is often referred to as a *dexterous hand*. Subsequent work relaxed those assumptions to include rolling [13, 14], sliding [15], finger gaiting [16, 17, 18, 19, 20, 21], and global controllability of object orientation [22, 23]. Despite numerous generalizations to Salisbury’s work, few attempt to alleviate the need for full hand actuation. A recent exception is by Odhner and Dollar [24] who demonstrate reduced local controllability of the object with an underactuated hand.

Most relevant to this paper is the work that exploits external resources for actuation, being in-hand manipulation one possible application. The more general use of external resources was noted by Lynch and Mason [25]:

“Other useful sources of control forces include gravity, the frictional kinematic constraints (floor, walls, obstacles) making up the robots environment, and dynamic forces. If the robot can reason about these forces, it can use a richer set of manipulation primitives, including pushing, throwing, and striking.”

We classify prior work by the resource employed: either external contacts, gravity, or dynamic motions. The following subsections explore those three categories.

A. Quasistatic actions (with external contacts)

Pushing against, or following, a flat surface is a common trick that humans use to align or relocate held objects. The mechanics of pushing an object under frictional contact with a flat surface was studied by Mason [26, 27], and later applied to manipulating polygonal objects in the plane [25], and to pushing objects into an enveloping grasp [28]. Nilsson [29] provides an earlier implementation of object manipulation by pushing and squeezing. Salisbury [30] and Brock [31] both propose to use external contacts as extra phalanges. The latter manipulates a held object by controlling the way it slips when pushed against an external contact. Paolini et al. [32] push a grasped cylinder against a flat surface to relocate it in the hand. In all cases, the environment helps to provide the wrench required to manipulate the object.

B. Passive Dynamic actions

Blessing or curse, gravity is a constant source of actuation that is always present. Erdmann and Mason [33] use it to manipulate a planar object lying on a tray by controlling its tilt. Erdmann [34] later uses it to manipulate a 3D object of unknown geometry lying against two palms. Both works plan transitions between a sequence of contact modes, relying on gravity to provide the wrench acting on the object.

C. Active Dynamic actions

Finally, we can manipulate an object held in the hand by relying on its inertia, and the effect of dynamic motions of the driving manipulator. We like to think that the powerful and precise arm actuators are repurposed to compensate for the possibly limited complexity of the hand.

Dynamic actions include throwing, catching, and juggling. Aboaf et al. [35, 36] and Burrige et al. [37] explored control strategies and learning techniques applied to catching and juggling. Starting from a more general formulation, Lynch and Mason [38] demonstrated controllability and planning of a variety of dynamic manipulation actions such as throwing, catching and rolling. They referred to this set of skills as “dynamic dexterity”. Another interesting mode of dynamic non-prehensile manipulation, explored by Bohringer et al. [39] and Vose et al. [40], is the use of friction force fields generated by custom vibration patterns to manipulate objects in the plane.

III. TERMINOLOGY

This paper uses the terms *grasp pose*, *grasp type*, *regrasp*, and *grasp graph*. Although we do not need precise definitions for them, we want to give an idea of what they refer to. Our usage doesn’t align perfectly with any previous work, but comes close to Zhang et al. [41]. Throughout, we assume there is a hand, with an associated hand reference frame, holding a single rigid object. We refer by *in-hand manipulation* to any way of moving the object with respect to the hand frame.

By *grasp pose* we mean the configuration of the system relative to the hand frame. This includes the pose of the object relative to the hand, and the configuration of the hand. In this paper, we use a simple hand [42] with three rigid fingers compliantly coupled to a single actuator. Each finger has a single joint, so the grasp pose is encoded as the object pose plus the three finger joint angles.

By *regrasp* we mean any action transferring the system from one grasp pose to another.

By *grasp type* we mean a subset of grasp pose space, like for example those commonly identified in a grasp taxonomy as “spherical enveloping” or “prismatic fingertip” [43].

We define the *grasp graph* as the directed graph with grasp types as vertices and regrasp actions as edges.

IV. REGRASP ACTIONS

This section catalogs the regrasp actions we implemented, working from the easy and familiar to the challenging and novel. All actions are implemented by a simple gripper with three fingers compliantly driven by a single actuator [42]. The hand has very limited “intrinsic dexterity”. In the absence of contact with the object or the environment, the hand is only capable of opening and closing, with all three fingers moving at the same rate.

The only intrinsic regrasp the hand can do is to squeeze an object, either resting on its palm or on a table. When the grip is tightened, contacts between hand and object may shift, leaving the object in a more stable and predictable pose.

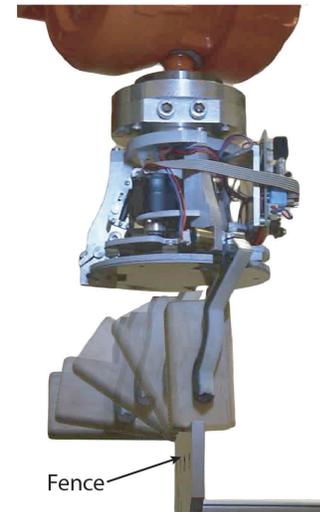


Fig. 3. Push-in-fingers regrasp: A block held between two fingertips rotates when the block makes contact with a fence.

Despite its simplicity, squeezing is a useful technique often exploited to reduce the uncertainty in the object’s pose [44, 45].

All other regrasp actions described in this section harness external resources rather than relying on the hand’s functionality. We divide them into three categories: quasistatic regrasps aided by contacts with the environment, passive-dynamic regrasps where the object is manipulated by the effect of gravity, and active-dynamic regrasps where fast arm motions use object inertia to reconfigure the object (see video [2]).

A. Quasistatic with External Contacts

In this section we describe five regrasps where the object is repositioned in the hand by exploiting contacts with the environment

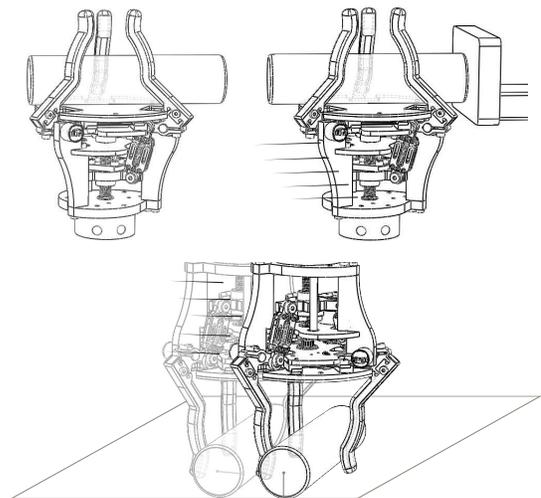


Fig. 4. Examples of quasistatic regrasps with external contacts: (top) push-in-enveloping regrasp; (bottom) roll-on-ground regrasp.

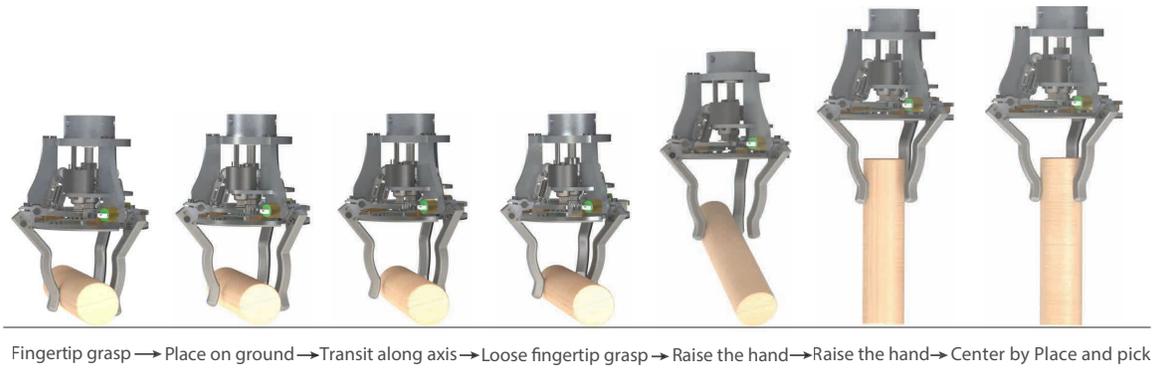


Fig. 5. Grasp pose changes involved in droop-in-fingers regrasp.

1) *Place-and-Pick*: The most widely used regrasp, both in research and industry, is to place the grasped object on a flat surface or fixture, and pick it up from a different direction. It fits within the pick-and-place manipulation paradigm and it is used, for example, in the Instant Insanity demo [7, 46] and the Handey system [9, 10, 47].

The place-and-pick regrasp is composed of three separate actions: *place*, *pick*, and *transit* (meaning to move the hand while the object rests on the ground) each of which can potentially be considered as an independent regrasp, if we think of the flat surface as another phalange, and the ungrasped object as a separate grasp type.

2) *Push-in-fingers*: (Fig. 3) The goal is to rotate an object about an axis determined by two fingertip contacts, similar to Brock’s work [31]. In our implementation, push-in-fingers rotates a rectangular block 90 degrees by moving the grasped block along a path obstructed by a fence.

3) *Push-in-enveloping*: (Fig. 4 top) The goal is to slide an object held in an enveloping grasp. The hand holds the object loosely against the palm and moves until the object makes contact with a fixture in the environment. A further motion against the fixture forces the object to slide in the hand. In previous work Paolini et al. [32] use this regrasp to relocate a grasped cylinder along its axis. By pressing both ends against a vertical post, the hand zeroes the uncertainty in the location of the cylinder along its axis.

4) *Roll-on-ground*: (Fig. 4 bottom) The goal is to rotate a cylinder about its main axis prior to picking it up. Starting with the cylinder lying on the ground and in a loose fingertip grasp, the hand moves parallel to the ground and perpendicular to the cylinder axis. If we assume slip at the fingertip contacts, and no slip at the ground contact line, the cylinder rolls through an angle determined by the hand’s linear motion.

5) *Droop-in-fingers*: (Fig. 5) The goal is to rotate a long thin object from an initial orientation parallel to the palm, to a final orientation perpendicular to the palm. For example, the hand might need to grasp a cylinder lying on the ground, and then insert it in a hole, requiring a change of grasp. The technique, illustrated in Fig. 5, involves grasping one end of the object loosely and lifting slowly. If the grasp is loose enough, the object’s weight causes the object to pivot in the

fingertips until it is vertical. Only after the pivoting motion is completed does the other end of the object lift off the ground, and the object aligns vertically.

B. Passive Dynamic

Here we describe regrasps governed by the dynamics of the object rolling on the finger phalanges under the effect of gravity. During these regrasp actions, the arm does not move at all, which avoids the need for synchronization between arm and hand motions. The outcome is consequently determined by the initial poses of hand and object.

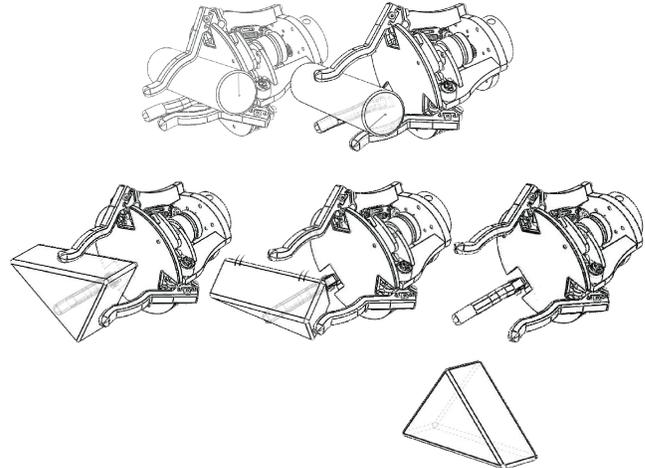


Fig. 6. Examples of passive dynamic regrasps with no arm motion: (top) roll-to-fingertip regrasp; (bottom) roll-to-ground regrasp

1) *Roll-to-fingertip*: (Fig. 6 top) The goal is to move the object from an enveloping grasp to a fingertip grasp by rolling on the fingers. The final pose can be adjusted by modifying the final opening of the hand. Fig. 6 (top) shows the application to a cylinder.

2) *Roll-to-ground*: (Fig. 6 bottom) The goal is to move the object from a fingertip grasp to rest on the ground, after some rotation. Fig. 6 (bottom) shows the application to a triangular prism with a rotation of 90 degrees. Roll-to-ground is more flexible than roll-to-fingertip, since the magnitude of the rotation is less limited by the length of the phalanges or the size of the hand.

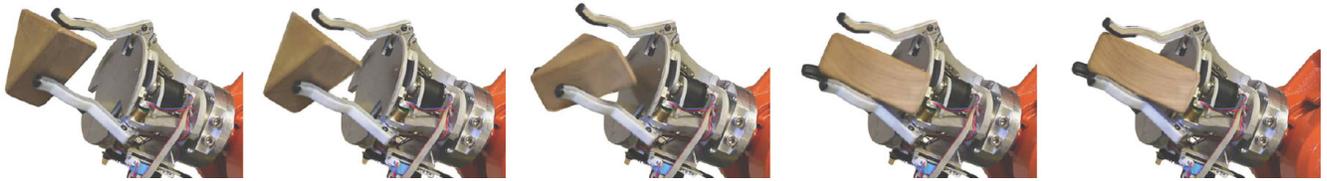


Fig. 7. Roll-to-palm regrasp: Transition from fingertip to enveloping with 90 degree rotation.

3) *Roll-to-palm*: (Fig. 7) Opposite to roll-to-fingertip, the object moves from a fingertip grasp to an enveloping grasp by rolling on the fingers. In Fig. 7 we show the application to a triangular prism. The rolling motion yields an offset rotation of 90 degrees.

C. Active Dynamic

The most challenging regrasp maneuvers involve arm motion coordinated with a dynamic motion of the object, as in [38, 48]. The most straightforward is a self-throw, where the object is thrown into the air and regrasped, which in theory could be used for almost any imaginable regrasp and for any object shape. The main challenges are sensitivity to initial conditions or throwing motion, and the careful timing required between arm, finger, and object.

Working in the earth's gravitational field means that an object in free fall has an acceleration of 1g, and the robot arm must be capable of comparable or higher accelerations to overcome it. It is unfortunate that many robot arms are not capable of such acceleration. By human standards it is not particularly high. By employing the whole body, baseball pitchers attain accelerations in excess of 100g [49]. More relevant, ordinary humans can easily attain hand accelerations exceeding 6g using elbow motion alone, as anyone can verify with the accelerometer in a smart phone.

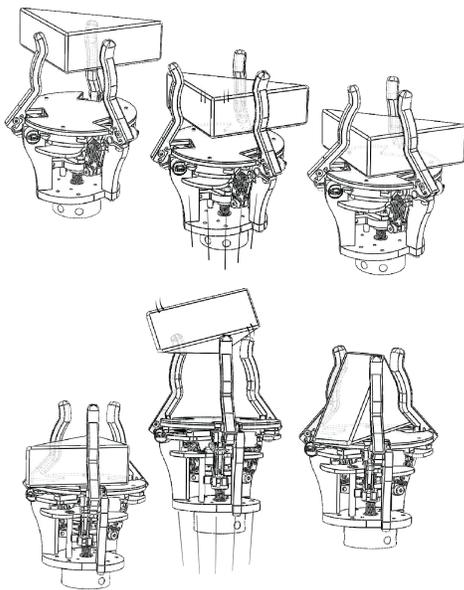


Fig. 8. Examples of active dynamic regrasps with arm motion: (top) throw-to-palm regrasp; (bottom) throw-and-flip regrasp.

Here we show a few variations of the self-throw idea:

1) *Throw-to-palm*: (Fig. 8 top) The goal is to transfer the object from a fingertip to an enveloping grasp, essential to many applications that require a firm grasp on an object. In our implementation, the robot holds the object in a fingertip grasp with the fingers pointed up and a horizontal palm. While the fingers open and the object drops, the arm is accelerated downward and then upward, in a profile that reduces the velocity mismatch at impact, helping the object settle into a stable pose. See the position and velocity profiles in Fig. 9. In practice, it is impossible to match velocities perfectly and some small bouncing remains. Recordings at 240 frames per second confirm that the arm's motion does substantially reduce bouncing.

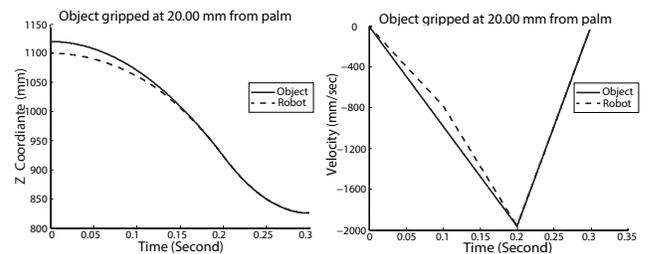


Fig. 9. Position (left) and velocity (right) of hand and object during throw-to-palm. The plan is for the object to touch the palm after 0.2sec, with matching velocities and positive acceleration of the palm.

2) *Throw-to-fingertip*: (Fig. 1) A more challenging task is to transfer the object from an enveloping grasp to a fingertip grasp by a self-throw. The hand starts with the fingers opened and pointed up, and the object resting on the palm. The hand is then accelerated upwards and then downwards at 2g throwing the object into the air. The fingers are then closed with the timing tuned to grip the object in the fingertips.

3) *Throw-and-flip*: (Fig. 8 bottom) Even more challenging is to change the grasp orientation while the object is in the air. In this maneuver, the object starts and ends in an enveloping grasp, but with an object rotation. Throw-and-flip was first observed empirically during the development of throw-to-fingertip for the triangular prism, and then tuned for performance and reliability. The prism is tossed in the air and during its fall it makes contact with the fingertips which causes it to rotate 90 degrees before falling into the palm.

4) *Vibration*: The goal is to attain a more predictable object pose by using vibration or dither to reduce the effect of friction.

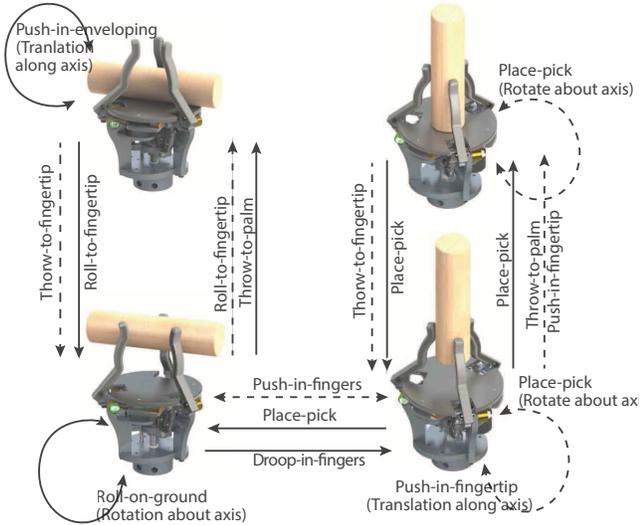


Fig. 10. Grasp graph for a cylinder. Solid arrows represent implemented regrasp, while dashed arrows show the regrasp that are not yet developed.

V. GRASP GRAPHS AND REGRASP SEQUENCES

For a given hand, a given object, and a give set of grasp types, we represent the repertoire of regrasp actions as a directed graph, called *grasp graph*, each arc corresponding to a regrasp. Figs. 2 and 10 show example graphs for a triangular prism and a cylinder respectively. A transition in the graph indicates the intended change in grasp type. In practice, failures occur, which brings stochasticity to the process. In a parallel line of work we are building stochastic data-driven models of these regrasp to quantify the expectations associated with grasp transitions.

The focus in this paper however, is to identify interesting regrasp actions that can be made to work robustly with only a manageable amount of tuning, and that can be sequenced to consolidate an in-hand manipulation capability. As seen in the accompanying video [2], this sometimes requires the addition of extra steps like squeeze or vibration designed to reduce uncertainty. In Section VI we evaluate the robustness of the regrasp and demonstrate examples of sequences.

A. Grasp Graph for a Triangular Prism (Fig. 2)

For the triangular prism in Fig. 2, we define grasp types for different choices of the face resting or opposite to the palm. In some cases, we distinguish between fingertip and enveloping grasps. We do not differentiate between symmetrical configurations of object or hand, although the graph is easily extended to do so. The grasp graph (Fig. 2) is strongly connected. If restricted to the actions implemented and tested for the triangular prism (solid arcs in the figure), we have a strongly connected sub-graph including three of the original 8 grasp types. Fig. 11 shows a sequence of regrasp exploiting that strong connectivity.

B. Grasp Graph for Cylinder (Fig. 10)

For the cylindrical object, we identified four interesting grasp types based on the orientation of the principal axis rel-

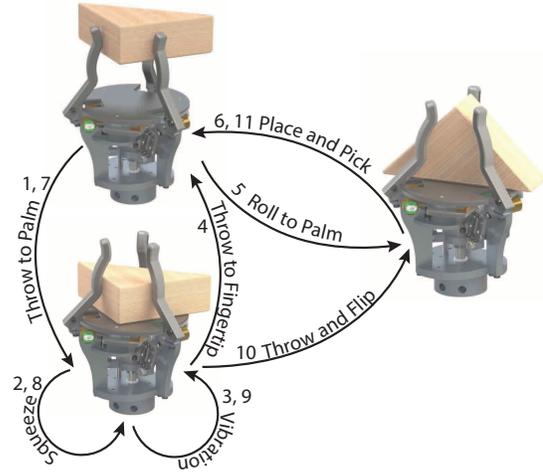


Fig. 11. Sequence of regrasp for a triangular prism. The path starts in the top-left grasp type and follows the arrows according to their numbers.

ative to the hand—parallel or perpendicular—and on whether it was held in a fingertip enveloping configuration. The full graph, as shown in Fig. 10, is strongly connected.

VI. EXPERIMENTAL RESULTS

This section describes experiments with individual regrasp actions and regrasp sequences. For the experiments we use a simple gripper with three rigid fingers compliantly connected to a single motor [42]. The hand is attached to a 6-DOF industrial manipulator (ABB IRB140) with a maximum speed of 2.5 m/sec and maximum acceleration of 20 m/sec² ($\sim 2g$). The workspace is composed of a table covered with rubber for friction and simple fixtures to support some of the regrasp. All actions are open loop and tuned by hand.

A. Individual Regrasps

To test the individual regrasp actions, we executed them 50 times in a row. The actions were calibrated and tuned by hand before the experiment, and then run 50 times at once without further adjustments. In some cases it was possible to remove the human from the loop, but typically, before each trial, a human placed the object as accurately as possible in a predefined initial pose. Success means that the object ended in the expected grasp type. Failure means the object ended in the wrong grasp type, or on the table. Table I shows the outcome of the experiments. The notes on limitations or typical failures are based on the authors' experience in designing and tuning the regrasp, not just on the formal 50-trial experiments.

We also explored the sensitivity to variations in the initial pose of the prism for the roll-to-palm action. Beginning with the prism grasped within the fingertips, as in Fig. 7, we introduce errors uniformly distributed over ± 20 mm in X and Y , ± 3 mm in Z , and ± 20 degrees rotation about Z , being X and Y parallel to the palm and Z perpendicular. Roll-to-palm was successful 580 out of 600 trials. The most sensitive directions were translations in Y and Z . Fig. 12 shows the success ratio plotted against those two dimensions.

Regrasp Title	Successes/50	Limitation/Typical Failure
Roll-to-palm	50	—
Roll-to-ground	34	Prism falls flat on ground
Roll-to-fingertip	50	Small error in net rotation
Throw-to-palm	50	Drop is not bounceless
Throw-to-fingertip	50	Object pose not always the same
Throw-and-flip	41	Prism caught in fingertips
Vibration	50	Change in pose not significant
Push-in-fingers	50	Slip between fingertips & prism
Roll-on-ground	50	Small error in net rotation
Droop-in-fingers	50	—
Push-in-enveloping	50	—
Squeeze	50	Fine reorientation not possible

TABLE I
OUTCOME OF REGRASP EXPERIMENTS

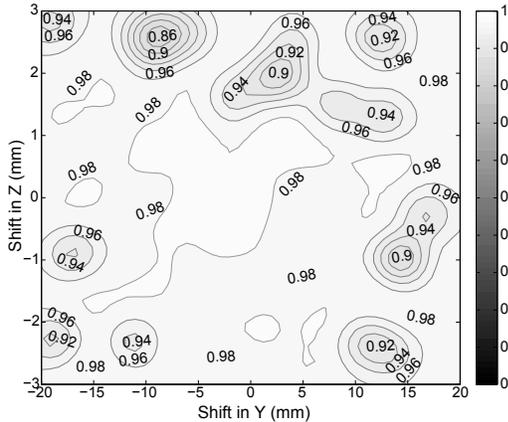


Fig. 12. Sensitivity of the probability of success of roll-to-palm to errors in the initial pose of the prism. The contours are generated from 600 experiments with random noise added to the initial pose of the prism in the three cartesian axis, as well as in its orientation in the axis perpendicular to the palm. The plot shows the two most significant axis.

B. Regrasp Sequences and Loops

In order to explore to what degree the regrasps can be sequenced without human intervention, we implemented one regrasp sequence and two regrasp loops.

Failures tend to occur because either actions are not as compatible to each other as we expected, or because errors accumulate from one action to the next. These limitations were not necessarily appreciable when developing the individual regrasp. The key observation is that for regrasp actions to be “usable”, it is important to have an understanding of the amount of uncertainty they can handle in the input, as well as the amount of uncertainty they produce in the output.

1) *Long regrasp sequence for a triangular prism:* We executed 30 times in a row the long sequence in Fig. 11, out of which 15 was successful. In most failures, throw-and-flip shifted the object to the correct grasp type but not centered well enough for the rest of the sequence to proceed. Another failure mode occurred when the object was not at the right height after the throw-to-fingertip regrasp. Recall that roll-to-palm is sensitive to the height.

2) *Back and forth between fingertip and enveloping grasps:* The throw-to-fingertip regrasp is sensitive to tim-

ing. To test it more thoroughly we programmed a loop in combination with the converse throw-to-palm regrasp, and ran it continuously for 50 trials. There were no failures, although variations of several millimeters were observed in the resulting fingertip grasps.

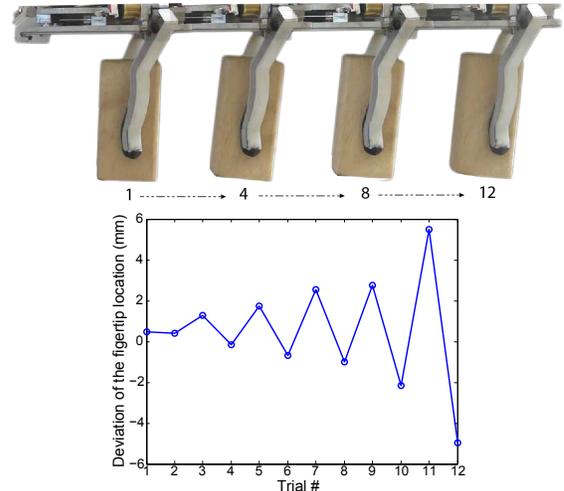


Fig. 13. Propagation of error for push-in-fingers regrasp

3) *Continuous rotation of square prism in fingers:* The push-in-fingers regrasp can be run in a loop all by itself, when programmed to rotate a square block by 90 degrees, as in Fig. 3. Typical errors in the object position are sub-millimeter, but the errors accumulate when the regrasp is executed in a loop, and the experiment typically fails after 8 or 9 trials, when the fingers reach an edge of the block, as illustrated in Fig. 13.

VII. DISCUSSION

Our results demonstrate that in-hand manipulation is possible with a hand dramatically simpler than typical dexterous hands. The key is to exploit the motions of the arm, object inertia, gravity, and external contacts: *extrinsic dexterity*.

This project is at an early stage, so the focus has been on inventing a repertoire of actions complete enough to navigate the grasp taxonomy, and to demonstrate their mechanical feasibility. The actions were hand-programmed and tuned for the objects and robot of choice. Most of the actions produce discrete changes of grasp pose, often involving changes in grasp type as well. Some of these operations seem to be well suited to the problem of changing between precision and power grasps, which is important in unstructured environments. It is important to notice, though, that many of the operations require a precision, speed, and acceleration that is well within the capabilities of traditional industrial manipulators, but not the arms more typically associated with research on unstructured anthropic environments.

The *grasp graph* represents a repertoire of regrasps for a particular object and hand, as well as their connectivity. It helps plan sequences of regrasps which can transfer the object from one pose to another, when the change is not possible in a single step. A key observation is that chaining

regrasps effectively requires an understanding of how much uncertainty it can handle in the input as well as how much it produces at the output.

Continuing work aims to study the basic principles underlying these actions, to combine analytical models with empirical data to characterize the operations, to automate the planning of the arm and hand motions, and to adapt them to different objects and tasks.

REFERENCES

- [1] J. K. Salisbury Jr., "Kinematic and Force Analysis of Articulated Hands," PhD Dissertation, Stanford University, 1982.
- [2] N. Chavan Daffe, A. Rodriguez, R. Paolini, B. Tang, S. Srinivasa, M. Erdmann, M. Mason, I. Lundberg, H. Staab, and T. Fuhlbrigge, "Regrasping objects using extrinsic dexterity," *IEEE Int Conf. on Robotics and Automation*, 2014. [Online]. Available: <https://www.youtube.com/watch?v=WAPgjZCI1gl>
- [3] A. Bicchi, "Hands for dexterous manipulation and robust grasping: a difficult road toward simplicity," *IEEE T Robotic Autom*, vol. 16, no. 6, pp. 652–662, 2000.
- [4] J. Napier and R. H. Tuttle, *Hands*. Princeton University Press, 1993.
- [5] G. Monkman, S. Hesse, and R. Steinmann, *Robot grippers*. John Wiley and Sons, 2007.
- [6] M. Kroh, K. El-Hayek, S. Rosenblatt, B. Chand, P. Escobar, J. Kaouk, and S. Chalikhonda, "First human surgery with a novel single-port robotic system: cholecystectomy using the da Vinci Single-Site platform," *Surgical Endoscopy*, vol. 25, no. 11, pp. 3566–3573, 2011.
- [7] R. Paul, K. Pingle, J. Feldman, and A. Kay, "Instant insanity," film, 1971.
- [8] J. Feldman, G. Feldman, G. Falk, G. Grape, J. Pearlman, I. Sobel, and J. Tenenbaum, "The Stanford hand-eye project," in *Int. Joint Conf. on Artificial Intelligence*, 1969, pp. 521–526.
- [9] P. Tournassoud, T. Lozano-Pérez, and E. Mazer, "Regrasping," in *IEEE Int. Conf. on Robotics and Automation*, vol. 4, 1987, pp. 1924–1928.
- [10] T. Lozano-Perez, J. Jones, E. Mazer, P. O'Donnell, W. Grimson, P. Tournassoud, and A. Lanusse, "Handey: A robot system that recognizes, plans, and manipulates," in *IEEE Int Conf. on Robotics and Automation*, vol. 4, 1987, pp. 843–849.
- [11] M. Mason and J. Salisbury Jr, *Robot Hands and the Mechanics of Manipulation*. Cambridge MA: MIT Press, 1985.
- [12] J. Salisbury and J. Craig, "Articulated hands: Force control and kinematic issues," *Int J Robot Res*, vol. 1, no. 1, pp. 4–17, 1982.
- [13] A. Bicchi and R. Sorrentino, "Dexterous manipulation through rolling," in *IEEE Int Conf. on Robotics and Automation*, 1995, pp. 452–457.
- [14] L. Han, Y. Guan, Z. Li, Q. Shi, and J. Trinkle, "Dextrous manipulation with rolling contacts," in *IEEE Int Conf. on Robotics and Automation*, 1997, pp. 992–997.
- [15] M. Cherif and K. Gupta, "Planning quasi-static fingertip manipulations for reconfiguring objects," in *IEEE T Robotic Autom*, vol. 15, 1999, pp. 837–848.
- [16] R. Fearing, "Simplified grasping and manipulation with dextrous robot hands," *IEEE Journal of Robotics and Automation*, vol. 2, no. 4, pp. 188–195, 1986.
- [17] J. Hong, G. Lafferriere, B. Mishra, and X. Tan, "Fine manipulation with multifinger hands," in *IEEE Int. Conf. on Robotics and Automation*, 1990, pp. 1568–1573.
- [18] T. Omata and K. Nagata, "Planning reorientation of an object with a multifingered hand," in *IEEE Int Conf. on Robotics and Automation*, 1994, pp. 3104–3110.
- [19] L. Han and J. Trinkle, "Dextrous manipulation by rolling and finger gaiting," in *IEEE Int Conf. on Robotics and Automation*, 1998, pp. 730–735.
- [20] D. Rus, "In-Hand Dexterous Manipulation of Piecewise-Smooth 3-D Objects," *Int J Robot Res*, vol. 18, no. 4, pp. 355–381, 1999.
- [21] T. Schlegl and M. Buss, "Hybrid closed-loop control of robotic hand regrasping," in *IEEE Int Conf. on Robotics and Automation*, 1998, pp. 3026–3031.
- [22] A. Marigo and A. Bicchi, "Rolling bodies with regular surface: controllability theory and applications," *IEEE Transactions on Automatic Control*, vol. 45, no. 9, pp. 1586–1599, 2000.
- [23] A. Bicchi and A. Marigo, "Dexterous grippers: Putting nonholonomy to work for fine manipulation," in *Int J Robot Res*, vol. 21, 2002, pp. 427–442.
- [24] L. Odhner and A. Dollar, "Dexterous manipulation with underactuated elastic hands," in *IEEE Int Conf. on Robotics and Automation*, 2011, pp. 5254–5260.
- [25] K. M. Lynch and M. T. Mason, "Stable pushing: Mechanics, controllability, and planning," *Int J Robot Res*, vol. 15, no. 6, pp. 533–556, 1996.
- [26] M. T. Mason, "Manipulator grasping and pushing operations," Ph.D. dissertation, Massachusetts Institute of Technology, 1982.
- [27] —, "Mechanics and planning of manipulator pushing operations," *Int J Robot Res*, vol. 5, no. 3, pp. 53–71, 1986.
- [28] M. Dogar and S. Srinivasa, "A framework for push-grasping in clutter," in *Robotics: Science and Systems VII*, 2011.
- [29] N. J. Nilsson, "Shakey the robot," SRI Int, Tech. Rep. 323, 1984.
- [30] K. Salisbury, "Whole arm manipulation," in *Robotics Research: The Fourth Int Symposium*, R. C. Bolles and B. Roth, Eds. Cambridge, MA, USA: MIT Press, 1988, pp. 183–189.
- [31] D. Brock, "Enhancing the dexterity of a robot hand using controlled slip," in *IEEE Int Conf. on Robotics and Automation*, vol. 1, 1988, pp. 249–251.
- [32] R. Paolini, A. Rodriguez, S. S. Srinivasa, and M. T. Mason, "A Data-Driven Statistical Framework for Post-Grasp Manipulation," *Int J Robot Res - To appear*, 2013.
- [33] M. A. Erdmann and M. T. Mason, "An exploration of sensorless manipulation," *IEEE Journal of Robotics and Automation*, vol. 4, no. 4, pp. 369–379, 1988.
- [34] M. A. Erdmann, "An exploration of nonprehensile two-palm manipulation," *Int J Robot Res*, vol. 17, no. 5, 1998.
- [35] E. Aboaf, S. Drucker, and C. Atkeson, "Task-level robot learning: Juggling a tennis ball more accurately," in *IEEE Int Conf. on Robotics and Automation*, 1989, pp. 1290–1295.
- [36] E. W. Aboaf, C. G. Atkeson, and D. J. Reinkensmeyer, "Task-level robot learning: Ball throwing," *Massachusetts Institute of Technology*, vol. 1006, 1987.
- [37] R. Burridge, A. Rizzi, and D. Koditschek, "Toward a dynamical pick and place," in *IEEE/RSJ Int Conf. on Intelligent Robots and Systems*, vol. 2, 1995, pp. 292–297.
- [38] K. M. Lynch and M. T. Mason, "Dynamic nonprehensile manipulation: Controllability, planning and experiments," *Int J Robot Res*, vol. 18, no. 1, pp. 64–92, 1999.
- [39] K.-F. Bohringer, V. Bhatt, B. R. Donald, and K. Y. Goldberg, "Algorithms for sensorless manipulation using a vibrating surface," *Algorithmica*, no. 3-4, pp. 389–429, 2000.
- [40] T. Vose, P. Umbanhowar, and K. Lynch, "Vibration-induced frictional force fields on a rigid plate," in *IEEE Int Conf. on Robotics and Automation*, 2007, pp. 660–667.
- [41] H. Zhang, K. Tanie, and H. Maekawa, "Dextrous manipulation planning by grasp transformation," in *IEEE Int Conf. on Robotics and Automation*, 1996, pp. 3055–3060.
- [42] M. T. Mason, A. Rodriguez, S. S. Srinivasa, and A. S. Vazquez, "Autonomous Manipulation with a General-Purpose Simple Hand," *Int J Robot Res*, vol. 31, no. 5, pp. 688–703, 2012.
- [43] M. Cutkosky, "On grasp choice, grasp models, and the design of hands for manufacturing tasks," *IEEE T Robotic Autom*, vol. 5, no. 3, pp. 269–279, 1989.
- [44] R. C. Brost, "Automatic Grasp Planning in the Presence of Uncertainty," in *IEEE Int Conf. on Robotics and Automation*, vol. 3, 1986, pp. 1575–1581.
- [45] K. Y. Goldberg, "Orienting Polygonal Parts without Sensors," *Algorithmica*, vol. 10, no. 204, pp. 201–225, 1993.
- [46] R. P. Paul, "Modeling, trajectory calculation and servoing of a computer controlled arm." Stanford Artificial Intelligence Laboratory, Technical Report AIM- 177, 1972.
- [47] T. Lozano-Pérez, J. Jones, P. O'Donnell, and E. Mazer, *Handey: a robot task planner*. MIT Press Cambridge, MA, USA, 1992.
- [48] T. Senoo, Y. Yamakawa, S. Mizusawa, A. Namiki, M. Ishikawa, and M. Shimajo, "Skillful manipulation based on high-speed sensory-motor fusion," in *IEEE Int Conf. on Robotics and Automation*, May 2009, pp. 1611–1612.
- [49] E. Berkson, R. Aylward, J. Zachazewski, J. Paradiso, and T. Gill, "IMU arrays: The biomechanics of baseball pitching," *Orthopaedic Journal at Harvard Medical School*, 2006.