123

Null Space Grasp Control: Theory and Experiments

Robert Platt Jr., Andrew H. Fagg, Roderic A. Grupen

***Abstract*—A key problem in robot grasping is that of positioning the manipulator contacts so that an object can be grasped. In unstructured environments, contact positions are typically planned based on range or visual measurements that are used to reconstruct object geometry. However, because it is difficult to measure the complete object geometry precisely in common grasp scenarios, it is useful to employ additional techniques to adjust or refine the grasp using only local information. In particular, grasp control techniques can be used to improve a grasp by adjusting the contact configuration after making initial contact with an object by using measurements of local object geometry at the contacts. This paper proposes three variations on null space grasp control, an approach that combines multiple grasp objectives to improve a grasp. Two of these variations are theoretically demonstrated to converge to force closure configurations for arbitrary convex objects when grasping with two contacts. All variations are found to converge in simulation. Robot grasping experiments are reported that show the approach to be useful in practice.**

I. Introduction

A key problem in robot grasping is positioning the contacts so that the necessary grasping forces can be applied. At each contact, the forces that can be applied depend on the local surface characteristics, including object surface normal and curvature. In unstructured environments, visual occlusions and sensor error make it difficult for a robot to measure the exact surface geometry of an object to be grasped before making contact. Therefore, the contacts must be placed on the object surface based on predictions that may be inaccurate. These predictions must ultimately be verified by force feedback when the robot actually makes contact.

When the predictions are wrong, it is advantageous to be able to adjust the manipulator configuration based on the sensed contact forces. Few approaches currently exist for accomplishing this step. After the contacts are placed on the object, how does the robot determine whether the grasp is good enough? When it is not, what mechanism can be used to displace the contacts toward better grasp locations? These two questions are the focus of the current paper. We describe key features of null space grasp control, a non-linear control strategy that synthesizes a grasp by using local measurements at the contacts to adjust contact configuration. The approach is predicated on a mechanism for measuring object surface normal in the neighborhood of each contact. Starting from an arbitrary configuration of the contacts on (or near) the object surface, measurements of the local object surface normals at the contacts are used to calculate a contact displacement on the object surface. Our experimental work uses six-axis load cells to measure the object surface normal while lightly touching the object. A contact displacement control system realizes the desired displacement by lightly sliding the contacts over the object surface.

We build upon force residual and moment residual controllers first proposed by Coelho and Grupen [1]. Coelho proved a convergence result for regular convex prismatic objects when the two controllers executed in a particular sequence and showed experimental results on a robot manipulator [2]. The current paper extends this work. First, we link the grasp controller to *unit frictionless equilibrium*, a special case of a force closure grasp. Second, we propose the null space approach to grasp control where force and moment residual controllers execute simultaneously. Three versions of the null space control law are proposed that trade off sensory requirements with speed of convergence:

- the exact null space grasp controller,
- the approximate null space grasp controller, and
- the switching grasp controller.

Convergence proofs are provided for the exact controller and the switching controller. All three variations are compared in simulation. Finally, robot experiments are presented that demonstrate the approach to be a practical mechanism for using local contact feedback to validate and improve robot grasps.

II. Related Work

A significant body of grasping research considers the problem of grasping in isolation from sensing considerations. This research typically begins with the assumption that the object geometry is known and that it is possible to sense object pose. One research direction identifies sufficient geometric conditions for a good grasp. For example, Nguyen proposed searching the space of two-contact configurations for those where a line connecting the two contacts lies inside friction cones associated with both contacts. This idea is the basis for algorithms that calculate two-contact force closure contact configurations for two- and three-dimensional polyhedral objects [3], [4] and curved objects [5], [6], [7]. This type of approach was extended to four-fingered grasps of polyhedral objects by Sudsang and Ponce who characterized four classes of four-contact grasp configurations [8]. Given the constraints associated with each grasp class, force closure grasps were found using optimization techniques. These ideas can be extended to in-hand manipulation by using the kinematics of rolling contact to move between different geometrically characterized grasp configurations [9], [10], [11].

Another approach to grasp planning finds grasps that optimize measures of grasp quality. As with the planning approaches above, these also generally ignore the sensing issue.

R. Platt is with the Massachusetts Institute of Technology, Cambridge, MA, 02139 (email: rplatt@csail.mit.edu).

A. H. Fagg is with the University of Oklahoma, Norman, OK, 73019 (email: fagg@cs.ou.edu).

R. A. Grupen is with the University of Massachusetts Amherst, Amherst, MA, 01003 (email: grupen@cs.umass.edu).

For example, Li and Sastry use linear optimization techniques to find contact configurations that optimize quality measures associated with the eigenvalues of the grasp map (relationship between contact loads and object loads) [12]. Kirkpatrick *et al.* and Ferrari and Canny propose optimizing a quality measure proportional to the radius of the largest sphere that can be inscribed in the convex hull of contact wrenches [13], [14]. Mirtich and Canny propose efficient planning algorithms for two and three contacts based on related quality measures [15].

In contrast to the above, a significant body of work combines sensing and planning in a two-step process where sensing occurs completely in advance of planning and control. For example, several researchers estimate the silhouette of an object to be grasped from different camera images and use the resulting spatial silhouette as input to a grasp planning algorithm. For the purposes of planning, the silhouette may be approximated by piecewise segments [16], [17], [18], smooth curves [19], or convex polyhedra [20]. In this context, some researchers also consider the feasibility of various grasp configurations in terms of manipulator kinematics [21], [22].

The application of tactile sensing to grasping in this paper is related to prior work that uses tactile information to estimate aspects of object shape and relative pose. Early work by Allen and Michelman modeled the surface of an unknown object as a superquadric using visual and tactile measurements [23]. Jia and Erdmann estimated contact position and object twist using an observer that was theoretically and empirically demonstrated to converge [24]. Haidacher and Hirzinger experimentally demonstrated an object localization method that matches tactile measurements to a best-fit object configuration [25]. Several researchers have solved a similar problem by applying statistical methods [26], [27], [28].

It is notable that all the work described above explicitly or implicitly divides grasp synthesis and manipulation into temporally separate perceptual and control processes. In contrast, the null space grasp control method characterized in this paper uses measurements continuously throughout the grasp synthesis process to adjust manipulator contact configuration. Our approach is more closely related to the work of Son, Howe, and Hager who combine visual and tactile "control primitives" to grasp a rod using a two-fingered gripper [29]. Using continuous tactile feedback, a gripper is re-oriented about a single axis so that it becomes better aligned with an object for grasping. Similarly, Yoshimi and Allen visually estimate the relative configuration of the object and manipulator and servo into a desired grasp configuration [30]. Another example of this type of approach are the provably-correct reactive grasping algorithms proposed by Teichmann and Mishra. These algorithms displace two or three manipulator contacts into a grasp configuration based on continually updated tactile feedback [31].

III. Grasp Objective Functions and Force Closure

The key idea of grasp control is to displace the contacts from an initial configuration on the object surface into a grasp configuration using measurements of local object geometry at the contacts. The grasp controller reaches grasp configurations by following the gradients of two objective functions: the unit frictionless force residual and the unit frictionless moment residual. These two objective functions lead the system into *unit frictionless equilibrium* configurations. This section introduces the notion of unit frictionless equilibrium as well as the two objective functions and relates them to force closure, a common quantitative measure of a grasp.

A. Grasp Objective Functions

For the purposes of the following development, it is useful to introduce the notion of *wrench*. A wrench, $\mathbf{w} = (\mathbf{f}^T, \mathbf{m}^T)^T$, is a screw that represents a combined force, $\mathbf{f}$, and moment, $\mathbf{m}$. Assume that all wrenches are expressed in a reference frame attached to the object located at the centroid of the contacts. A system of k contacts touching an object is in *equilibrium* when the sum of the wrenches applied to the object at each contact (the contact wrenches) is zero:

$$\sum_{i=1}^{k} \mathbf{w}_i = \mathbf{0}, \tag{1}$$

where $\mathbf{w}_i$ is the i^{th} contact wrench. We define *unit frictionless equilibrium* to be the special case of equilibrium were all contacts apply unit forces normal to the object surface:

Definition 1: A system of contacts is in unit frictionless equilibrium when it is in equilibrium and the contact wrenches, $\mathbf{w}_1 \dots \mathbf{w}_k$, satisfy:

$$\mathbf{w}_i = \begin{pmatrix} \hat{\mathbf{n}}_i \\ \mathbf{r}_i \times \hat{\mathbf{n}}_i \end{pmatrix},$$

where $\hat{\mathbf{n}}_i$ and $\mathbf{r}_i$ are the unit object surface normal and the position of the i^{th} contact, respectively.

When a two-contact system is in unit frictionless equilibrium, the contacts are in an antipodal configuration (parallel and intersecting contact normals). When a three-contact system is in unit frictionless equilibrium, the contact normals lie in a plane and intersect at a single point.

The proposed grasp control approach reaches unit frictionless equilibrium by descending the unit frictionless force residual and moment residual error functions. The squared unit frictionless force residual is defined to be:

$$\epsilon_f = \frac{1}{2}\mathbf{f}^T\mathbf{f}, \quad \mathbf{f} = \sum_{i=1}^{k} \hat{\mathbf{n}}_i. \tag{2}$$

When the unit frictionless force residual is zero, then all of the unit normals are balanced. Such a configuration will be known as unit frictionless force equilibrium. The squared unit frictionless moment residual is defined to be:

$$\epsilon_m = \frac{1}{2}\mathbf{m}^T\mathbf{m}, \quad \mathbf{m} = \sum_{i=1}^{k} \mathbf{r}_i \times \hat{\mathbf{n}}_i. \tag{3}$$

When the unit frictionless moment residual is zero, then the system is in unit frictionless moment equilibrium.

B. Relationship to Force Closure

Force closure is a way of quantifying the term "grasp." A force closure contact configuration can resist arbitrary loads applied to the object (from gravity or other sources) by applying appropriate combinations of contact wrenches [32]. For two or more contacts, unit frictionless equilibrium is a special case of force closure for any non-zero coefficient of coulomb friction. This was demonstrated for three or more contacts by Ponce [33] and can be extended to two contacts if it is assumed that the contacts are able to apply frictional torsional loads about the contact normals (this is typically known as the "soft contact" assumption [34]):

Lemma 1: When at least one contact can apply frictional torsional loads about the contact normal as well as tangential frictional forces, then a sufficient condition for three-dimensional 2-finger force closure is non-marginal equilibrium.

A contact configuration is in *non-marginal equilibrium* when it is in equilibrium and all contact wrenches are *strictly* within (not on the edge of) their respective friction cones. Since unit frictionless equilibrium grasps apply forces only along the surface normals (at the center of the associated friction cone), these grasps must therefore be force closure when the contacts are able to apply positive tangential frictional forces. Lemma 1 is proven in the Appendix.

IV. The Force and Moment Residual Controllers

Grasp control synthesizes grasps by displacing contacts over the object surface into grasp configurations using local contact feedback. This section describes the contact displacement mechanism. It also describes the force and moment residual control laws that are combined by null space grasp control.

A. Mechanism for contact displacement

Grasp control is predicated on a mechanism for measuring the object surface normal near the contacts while displacing them over the object surface. This can be accomplished in two ways: 1) by touching the object lightly so as to make the necessary measurements using force sensors without disturbing the object, and 2) by placing the contacts near enough to the object to be able to detect surface normal using non-contact sensors.

Our experimental work takes the first approach by using six-axis load cells mounted in the robot fingertips to measure object surface normal and contact forces while touching the object (see Section VII for hardware details). A simple torque control law is used to touch the object lightly:

$$\ddot{q}^* = K_p(\tau^* - \tau) - K_d\dot{q},$$

where $\ddot{q}^*$ is a commanded finger acceleration calculated by a proportional term on finger joint torque error, $\tau^* - \tau$, and a damping term on actual joint velocity, $\dot{q}$ [35]. Joint torque, τ, is calculated using force measurements from the fingertip load cell. While touching the object with the load cell, contact wrench measurements coupled with knowledge of the convex contact geometry are used to calculate the object surface normal [36].

A potential problem with contact displacement while touching is that the process causes small unintended object displacements. Although this was not a significant problem in our experimental work (see Section VII), null space grasp control can also be implemented without touching the object by using non-contact proximity sensors. For example, Teichmann and Mishra's implementation of reactive grasping using a parallel jaw gripper uses optical proximity sensors to measure local object geometry without touching in the context of a similar grasp displacement strategy [37]. Similarly, Walker and Salisbury's PMET manipulator uses optical proximity sensors to measure distance to the object surface without touching. Object surface normal is calculated by differentiating a series of distance measurements [38]. Instead of optical sensing, LIDAR might also be used to measure local object surface curvature when the scale of manipulation is large enough. Finally, in the future, new technologies such as electric field pretouch sensing may be used in ways similar to the above [39].

B. Force Residual Controller

Assume that the controller interacts with a second-order continuous spatial object with two or three contacts. The force residual controller follows the negative gradient of a unit-curvature approximation of the unit frictionless force residual (Equation 2). Let the surface of the object be parameterized by orthogonal parameter curves, u and v. Let $\mathbf{r}_i(u,v)$ describe the three-dimensional Cartesian position of the i^{th} contact as a function of the parameter curves. Let $\nabla_u\mathbf{r}_i$ and $\nabla_v\mathbf{r}_i$ denote $\frac{\partial\mathbf{r}_i}{\partial u_i}$ and $\frac{\partial\mathbf{r}_i}{\partial v_i}$, the tangents to the u and v parameter curves at contact i. Define the sense of the curves such that $(\nabla_u\mathbf{r}_i, \nabla_v\mathbf{r}_i, \hat{\mathbf{n}}_i)$ forms a right-hand orthonormal coordinate frame at each contact.

The gradient of the squared unit frictionless force residual (Equation 2) with respect to these surface coordinates is:

$$\frac{\partial\epsilon_f}{\partial\mathbf{u}} = \mathbf{f}^T J_f, \tag{4}$$

where $\mathbf{u} = (\mathbf{u}_1, \ldots, \mathbf{u}_k)$ is a vector describing the surface coordinates of k contacts, $\mathbf{f}$ is the unit frictionless force residual (Equation 2), and $J_f = \frac{\partial\mathbf{f}}{\partial\mathbf{u}}$ is the unit frictionless force residual Jacobian.

J_f may be decomposed into k partial derivatives:

$$J_f = \left(\frac{\partial\mathbf{f}}{\partial\mathbf{u}_1}, \ldots, \frac{\partial\mathbf{f}}{\partial\mathbf{u}_k}\right).$$

The i^{th} partial derivative can be expressed as follows:

$$\begin{aligned}\frac{\partial\mathbf{f}}{\partial\mathbf{u}_i} &= \frac{\partial\hat{\mathbf{n}}_\mathbf{i}}{\partial\mathbf{u}_i} \\ &= (\nabla_u\mathbf{r}_i, \nabla_v\mathbf{r}_i)\, K_i,\end{aligned}$$

where K_i is a 2×2 symmetric matrix of surface curvatures for contact i. Therefore, the unit frictionless force residual Jacobian is a matrix of surface tangents multiplied by a matrix of surface curvatures:

$$\begin{aligned}J_f &= (\nabla_u\mathbf{r}_1, \nabla_v\mathbf{r}_1, \ldots, \nabla_u\mathbf{r}_k, \nabla_v\mathbf{r}_k)\, K \\ &= \hat{J}_f K.\end{aligned}$$

where K is a $2k \times 2k$ symmetric block diagonal matrix comprised of K_i for each contact and $\hat{J}_f$ is a matrix whose columns are the object surface tangents at all contacts. Since for the unit sphere, K is identity and $\hat{J}_f = J_f$, we refer to $\hat{J}_f$ as the unit-curvature frictionless force residual Jacobian.

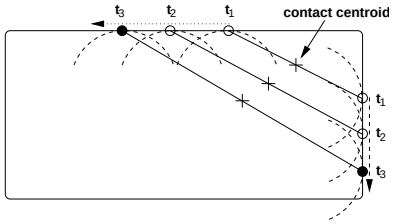

Fig. 1. The force residual controller calculates the force residual gradient by assuming that each contact normal will change as if the contact were moving on a sphere tangent to the object at the contact point. At each iteration of the controller, the gradient is recomputed using the spherical assumption.

The force residual controller follows the negative gradient of Equation 2 while assuming unit curvatures:

$$\dot{\mathbf{u}}_f = -\hat{J}_f^T \mathbf{f}. \tag{5}$$

To elucidate the effect of the force residual term of the grasp controller, consider the force residual controller executing for the planar rectangle illustrated in Figure 1. The force residual gradient assumes that the contacts are moving on surfaces with positive unit curvatures (*i.e.* spheres). The gradient with respect to the two contact positions is illustrated by the dashed arrows pointing tangent to the object surface. The controller sends this displacement to a control mechanism for displacing the contacts. On the next control cycle, the contacts will have moved in the direction of the dashed arrows and the gradient will be re-evaluated.

C. Convergence of the force residual controller

The force residual controller (Equation 5) can be shown to converge to unit frictionless force equilibrium configurations when grasping convex objects for two contacts. Consider the following Lyapunov function:

$$V = \frac{1}{2}\mathbf{f}^T\mathbf{f}. \tag{6}$$

The gradient of Equation 6 with respect to surface coordinates is:

$$\begin{aligned} \frac{\partial V}{\partial \mathbf{u}} &= \mathbf{f}^T \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \\ &= \mathbf{f}^T \hat{J}_f K. \end{aligned}$$

Therefore, the gradient of $\dot{V}_f$ along controller trajectories is:

$$\begin{aligned} \dot{V} &= \frac{\partial V}{\partial \mathbf{u}}\dot{\mathbf{u}}, \\ &= -\mathbf{f}^T \hat{J}_f K \hat{J}_f^T \mathbf{f}. \end{aligned}$$

Since K is always positive semi-definite for convex objects, it is clear that $\dot{V}$ is negative semi-definite.

Theorem 1: Let the object be convex, second-order continuous with finite maximum curvature. Then the two-contact force residual controller (Equation 5) converges to unit frictionless force equilibrium when execution does not begin with both contacts on the same face.

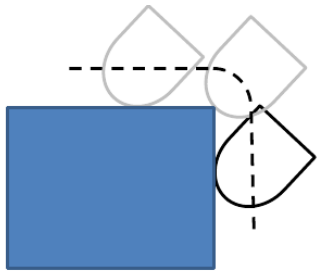

Fig. 2. Extruded object (dashed line) traced out by the center of a finger as it moves over the box.

Proof: Since $\dot{V}$ is negative semi-definite, the force residual controller (Equation 5) must be stable. It converges to configurations where $\dot{V}$ is zero: in unit frictionless force equilibrium, when both contacts are on the same face, or when the columns of K are orthogonal to $\hat{J}_f^T\mathbf{f}$. First, note that since $\dot{V}$ is negative semi-definite and V is at a maximum when both contacts are on the same face, the system never reaches a same-face configuration when a same-face initial configuration is prohibited. Second, consider the situation where the columns of K are orthogonal to $\hat{J}_f^T\mathbf{f}$. In this case, the object surface at each contact is flat in its direction of motion. Each contact continues to move along a flat surface until one contact reaches a region of positive curvature and the gradient of the Lyapunov function is again negative definite. While it is possible that the contact may reach another region where the object surface is flat in the direction of motion, V decreases every time a contact passes through a positive curvature region. Therefore, for objects with finite extent, V ultimately reaches zero in finite time and we conclude that the controller converges to unit frictionless force equilibrium. ■

The requirement by Theorem 1 for the object to be second-order continuous theoretically excludes polygonal objects. Nevertheless, these objects are not excluded in practice when a manipulator with rounded contacts is used. In this case, it is possible to define a corresponding extruded object that is traced out by a point on the interior of the rounded contact (see Figure 2). Configurations of the rounded contacts on the actual object map onto point contact configurations for the extruded object. See [15] for more detail on this argument. Theorem 1 can be applied to the extruded object and, since unit frictionless equilibrium configurations for the extruded object can be shown also to be unit frictionless equilibrium on the actual object, extended to the actual object.

D. Moment Residual Controller

The moment residual controller follows the gradient of the unit frictionless moment residual while making a specific curvature assumption. The gradient of Equation 3 is:

$$\frac{\partial \epsilon_m}{\partial \mathbf{u}} = \mathbf{m}^T J_m$$

where

$$J_m = \left(\frac{\partial \mathbf{m}}{\partial u_1}, \frac{\partial \mathbf{m}}{\partial v_1}, \ldots, \frac{\partial \mathbf{m}}{\partial u_k}, \frac{\partial \mathbf{m}}{\partial v_k} \right).$$

The partial derivative of the unit frictionless moment residual with respect to u_i is:

$$\frac{\partial \mathbf{m}}{\partial u_i} = \nabla_u \mathbf{r}_i \times \hat{\mathbf{n}}_\mathbf{i} + \mathbf{r}_i \times \frac{\partial \hat{\mathbf{n}}_\mathbf{i}}{\partial u_i}.$$

Rather than incorporating surface curvature information into the moment residual gradient, the moment residual controller sets the second term to zero, effectively assuming zero surface curvature at the contacts:

$$\begin{aligned} \dot{u}_i &= \mathbf{m}^T \left(\nabla_u \mathbf{r}_i \times \hat{\mathbf{n}}_i \right) \\ &= -\mathbf{m}^T \nabla_v \mathbf{r}_i. \end{aligned}$$

Coelho refers to this simplification as the "planar assumption" [1]. Extending this argument to the entire moment residual control law, we have:

$$\dot{\mathbf{u}}_m = -\hat{J}_m^T \mathbf{m}, \tag{7}$$

where

$$\hat{J}_m = \left(-\nabla_v \mathbf{r}_1, \nabla_u \mathbf{r}_1, \ldots, -\nabla_v \mathbf{r}_k, \nabla_u \mathbf{r}_k \right)$$

is the zero-curvature frictionless moment residual Jacobian.

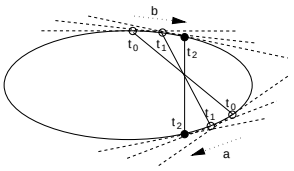

Fig. 3. The moment residual controller calculates the moment residual gradient by assuming that object geometry is a plane tangent to the object at each point of contact. At each iteration of the controller, the gradient is re-computed assuming a plane tangent to the current set of contact points.

To clarify the differences between the moment residual control gradient (Equation 7) and the exact gradient, consider the planar object in Figure 3. The approximation "thinks" that the contacts will move as if the local surface were flat as illustrated by the dotted lines. Following the gradient would cause contact a to move to the left and contact b to move to the right as illustrated by the dashed arrows.

V. Null space grasp control

Null space grasp control is an approach to combining the force and moment residual controllers in a way that realizes force closure grasps for arbitrary convex objects. This section proposes exact and approximate null space grasp control. The exact method projects the moment residual controller displacements into the null space of the gradient of the unit frictionless force residual (Equation 2) and is provably convergent for two contacts. Based on our simulations, this approach reaches unit frictionless equilibrium configurations faster than the other approaches studied in this paper. However, since it is difficult to measure object surface curvature, this method is difficult to implement. As a result, we also propose the approximate null space grasp controller in this section and the switching grasp controller in the next section.

The null space grasp controller assures that the moment residual controller does not cause the system to ascend the unit frictionless force residual by projecting moment residual control into the null space of the unit frictionless force residual gradient (Equation 4), $\frac{\partial \epsilon_f}{\partial \mathbf{u}}$:

$$\dot{\mathbf{u}}^* = -\hat{J}_f^T \mathbf{f} - \mathcal{N}\left(\mathbf{f}^T \hat{J}_f K \right) \hat{J}_m^T \mathbf{m}. \tag{8}$$

Since

$$\mathbf{f}^T \hat{J}_f K \mathcal{N}\left(\mathbf{f}^T \hat{J}_f K \right) \dot{\mathbf{y}} = 0, \tag{9}$$

for arbitrary contact displacements, $\dot{\mathbf{y}}$, $\dot{V}$ for this control law is still negative semi-definite and the result of Theorem 1 is unchanged.

A. Force and moment residual controllers for two contacts

This subsection introduces notation for two contacts that simplifies the subsequent development of the force and moment residual controllers. Let the v parameter curve of the object surface parameterization pass through both contacts at an identical tangent such that $\nabla_v \mathbf{r}_1 = \nabla_v \mathbf{r}_2$. Then, the force residual control gradient becomes:

$$\begin{aligned} \dot{\mathbf{u}}_f &= -\begin{pmatrix} \nabla_u \mathbf{r}_1^T \\ \nabla_v \mathbf{r}_1^T \\ \nabla_u \mathbf{r}_2^T \\ \nabla_v \mathbf{r}_2^T \end{pmatrix} \mathbf{f} \\ &= -\begin{pmatrix} \alpha \\ 0 \\ -\alpha \\ 0 \end{pmatrix}, \end{aligned} \tag{10}$$

where the substitution,

$$\alpha = \nabla_u \mathbf{r}_1^T \hat{\mathbf{n}}_2 = -\nabla_u \mathbf{r}_2^T \hat{\mathbf{n}}_1, \tag{11}$$

has been made (see Lemma 2 in the Appendix for the demonstration that $\nabla_u \mathbf{r}_1^T \hat{\mathbf{n}}_2 = -\nabla_u \mathbf{r}_2^T \hat{\mathbf{n}}_1$).

For two contacts, the moment residual gradient is:

$$\begin{aligned} \dot{\mathbf{u}}_m &= -\hat{J}_m^T \mathbf{m} \\ &= -\begin{pmatrix} -\nabla_v \mathbf{r}_1^T \\ \nabla_u \mathbf{r}_1^T \\ -\nabla_v \mathbf{r}_2^T \\ \nabla_u \mathbf{r}_2^T \end{pmatrix} \left(\mathbf{r}_1 \times \hat{\mathbf{n}}_1 + \mathbf{r}_2 \times \hat{\mathbf{n}}_2 \right). \end{aligned}$$

Since the origin of the reference frame is at the contact centroid, the two contact position vectors are opposite, $\mathbf{r}_1 = -\mathbf{r}_2$, and the gradient becomes:

$$\dot{\mathbf{u}}_m = -\begin{pmatrix} \mathbf{r}_1^T (\nabla_u \mathbf{r}_1 - \nabla_u \mathbf{r}_2) \\ \mathbf{r}_1^T (\nabla_v \mathbf{r}_1 - \hat{\mathbf{n}}_2 \times \nabla_u \mathbf{r}_1) \\ \mathbf{r}_1^T (\nabla_u \mathbf{r}_1 - \nabla_u \mathbf{r}_2) \\ -\mathbf{r}_1^T (\nabla_v \mathbf{r}_1 - \hat{\mathbf{n}}_1 \times \nabla_u \mathbf{r}_2) \end{pmatrix}.$$

The notation in this equation and others to follow is simplified with the following substitutions:

$$\begin{aligned} p &= \mathbf{r}_1^T \nabla_u \mathbf{r}_1, && (12)\\ q &= \mathbf{r}_1^T \nabla_u \mathbf{r}_2, && (13)\\ s &= \mathbf{r}_1^T \nabla_v \mathbf{r}_1, && (14)\\ a &= \mathbf{r}_1^T (\hat{\mathbf{n}}_2 \times \nabla_u \mathbf{r}_1), && (15)\\ b &= \mathbf{r}_1^T (\hat{\mathbf{n}}_1 \times \nabla_u \mathbf{r}_2). && (16) \end{aligned}$$

Then, the expression for $\dot{\mathbf{u}}_m$ is:

$$\dot{\mathbf{u}}_m = -\begin{pmatrix} p-q \\ s-a \\ p-q \\ b-s \end{pmatrix}. \quad (17)$$

B. Convergence of the null space grasp controller for two contacts

Since it is already established that Equation 8 converges to unit frictionless force equilibrium, all that remains is to show that it also converges to unit frictionless moment equilibrium. We establish this for two contacts. Consider the following second-order continuous positive definite function defined over two-contact configurations on convex objects where the system is in unit frictionless force equilibrium:

$$\begin{aligned} W &= \frac{1}{4}(\mathbf{r}_1 - \mathbf{r}_2)^T (\mathbf{r}_1 - \mathbf{r}_2) \\ &= \mathbf{r}_1^T \mathbf{r}_1. \end{aligned} \quad (18)$$

For two contacts, the derivative of W with respect to surface coordinates for two contacts is:

$$\frac{\partial W}{\partial \mathbf{u}} = 2\begin{pmatrix} \mathbf{r}_1^T \nabla_u \mathbf{r}_1 \\ \mathbf{r}_1^T \nabla_v \mathbf{r}_1 \\ -\mathbf{r}_1^T \nabla_u \mathbf{r}_2 \\ -\mathbf{r}_1^T \nabla_v \mathbf{r}_2 \end{pmatrix}^T = 2\begin{pmatrix} p \\ s \\ -q \\ -s \end{pmatrix}^T. \quad (19)$$

Theorem 2: Let the object be convex, second-order continuous with finite maximum curvature. For two contacts, the null space grasp controller converges to unit frictionless equilibrium when execution does not begin with both contacts on the same face.

Proof:

The gradient of W along the trajectories of the composite null space controller is:

$$\begin{aligned} \dot{W} &= \frac{\partial W}{\partial \mathbf{u}} \left(\dot{\mathbf{u}}_f + \mathcal{N}\left(\mathbf{f}^T \hat{J}_f K\right) \dot{\mathbf{u}}_m \right) \\ &= -2(p+q)\alpha - 2\begin{pmatrix} p \\ s \\ -q \\ -s \end{pmatrix}^T A \begin{pmatrix} p-q \\ s-a \\ p-q \\ b-s \end{pmatrix}, \end{aligned} \quad (20)$$

where $A = \mathcal{N}\left(\mathbf{f}^T \hat{J}_f K\right)$ is a positive definite projection matrix.

Note that A is never in the null space of $\dot{\mathbf{u}}_m$ or $\frac{\partial W}{\partial \mathbf{u}}$: A projects to zero the component of $\dot{\mathbf{u}}_m$ that is parallel with $K\dot{\mathbf{u}}_f$. Since K is composed of block diagonal positive definite matrices, we have that

$$K\dot{\mathbf{u}}_f = -\begin{pmatrix} \kappa_{11} \\ 0 \\ -\kappa_{33} \\ 0 \end{pmatrix} \alpha$$

where κ_{11} and κ_{33} are positive. In view of Equation 17, it is clear that $\dot{\mathbf{u}}_m$ is never parallel with $K\dot{\mathbf{u}}_f$. Also, notice that p is equal to q only when the two contacts are concurrent. Since this is prohibited by the assumption that execution does not begin with contacts on the same face, $\frac{\partial W}{\partial \mathbf{u}}$ is never parallel with $K\dot{\mathbf{u}}_f$.

We now show convergence to unit frictionless moment equilibrium. Given Theorem 1 and the consideration in Equation 9, we have that the null space grasp controller converges to unit frictionless force equilibrium. As $\mathbf{f}$ approaches zero, Lemma 3 requires that α approaches zero and therefore that the first term of Equation 20 approaches zero. In view of Lemma 4, $\frac{\partial W}{\partial \mathbf{u}}\dot{\mathbf{u}}_m$ is always negative semi-definite. Since A is positive definite, the second term of Equation 20 is also negative semi-definite.

As a result, Equation 20 is always negative semi-definite and the controller converges to a configuration where Equation 20 is zero. Since A is never in the null space of $\dot{\mathbf{u}}_m$ or $\frac{\partial W}{\partial \mathbf{u}}$, the second term of Equation 20 is zero only when $\frac{\partial W}{\partial \mathbf{u}}\dot{\mathbf{u}}_m$ is zero. Since the two contacts are assumed never to be concurrent, this only occurs when p, q, s, a, and b are zero. When this happens, note that $\mathbf{r}_1 = -\mathbf{r}_2$ is normal to the surface tangent at each contact and that $\mathbf{m}$ is therefore zero.

■

Theorem 2 can be combined with Lemma 1 to conclude that the exact null space grasp controller converges to force closure configurations.

C. Approximate null space grasp controller

The exact null space grasp controller (Equation 8) requires knowledge of the object surface curvature, K, in order to calculate the null space projection matrix, $\mathcal{N}(\mathbf{f}^T \hat{J}_f K)$. Since it may be difficult to measure local object surface curvature at the contacts, we consider alternatives to the exact formulation of the control law. One approach is the "approximate" null space controller that projects the moment residual controller into the null space of $\hat{J}_f$:

$$\dot{\mathbf{u}}^* = -\hat{J}_f^T \mathbf{f} - \mathcal{N}\left(\hat{J}_f\right) \hat{J}_m^T \mathbf{m}. \quad (21)$$

This controller has not been proven to converge. However, simulated results (see Section VII-A) suggest that it does converge slower than the exact null space controller but faster than the switching controller proposed in the next section. For insight into how the controller works, consider the null space projection matrix, $\mathcal{N}\left(\hat{J}_f\right)$. When a two-contact system is not in unit frictionless force equilibrium, the rank of $\hat{J}_f$ is three and the rank of $\mathcal{N}\left(\hat{J}_f\right)$ is therefore one. The rank of the null space projection term rises to two when $\mathbf{f} = 0$. This suggests that, similar to the switching controller, this controller allows the moment residual term to converge faster after the system reaches unit frictionless force equilibrium.

VI. Switching grasp control

Like the approximate null space grasp controller, the switching grasp controller does not require object surface curvature at the contacts to be measured. When the unit frictionless force residual is large it executes the force residual controller by itself. Once the unit frictionless force residual falls below a threshold, then the controller displaces the contacts according to the sum of the force residual and the moment residual control gradients. This controller is proven to converge to unit frictionless equilibrium.

A. Switching grasp controller

The switching grasp controller switches between executing the force residual controller when $\|\mathbf{f}\| > \beta$ and executing the moment residual controller when $\|\mathbf{f}\| \leq \beta$. This is accomplished using an indicator variable:

$$N_a = \begin{cases} 1 \text{ if } \|\mathbf{f}\| \leq \beta \\ 0 \text{ otherwise} \end{cases}$$

The resulting controller is:

$$\dot{\mathbf{u}}_a = \dot{\mathbf{u}}_f + N_a \dot{\mathbf{u}}_m. \tag{22}$$

B. Convergence of the switching grasp controller for two contacts

In order to establish convergence, we evaluate the gradients of V and W for the two cases, $\|\mathbf{f}\| \leq \beta$ and $\|\mathbf{f}\| > \beta$. The derivative of V with respect to surface coordinates for two contacts is:

$$\frac{\partial V}{\partial \mathbf{u}} = \begin{pmatrix} \alpha \\ 0 \\ -\alpha \\ 0 \end{pmatrix}^T K.$$

When $\|\mathbf{f}\| > \beta$, then $N_a = 0$ and $\dot{\mathbf{u}}_a = \dot{\mathbf{u}}_f$:

$$\begin{aligned} \dot{V}_{\|\mathbf{f}\|>\beta} &= \frac{\partial V}{\partial \mathbf{u}} \dot{\mathbf{u}}_f \\ &= -\begin{pmatrix} \alpha \\ 0 \\ -\alpha \\ 0 \end{pmatrix}^T K \begin{pmatrix} \alpha \\ 0 \\ -\alpha \\ 0 \end{pmatrix} \\ &= -(\kappa_{11} + \kappa_{33})\alpha^2, \end{aligned} \tag{23}$$

where κ_{11} and κ_{33} are positive diagonal elements of K. The gradient of W when $\|\mathbf{f}\| > \beta$ is:

$$\begin{aligned} \dot{W}_{\|\mathbf{f}\|>\beta} &= -2\begin{pmatrix} p \\ s \\ -q \\ -s \end{pmatrix}^T \begin{pmatrix} \alpha \\ 0 \\ -\alpha \\ 0 \end{pmatrix} \\ &= -2(p+q)\alpha. \end{aligned} \tag{24}$$

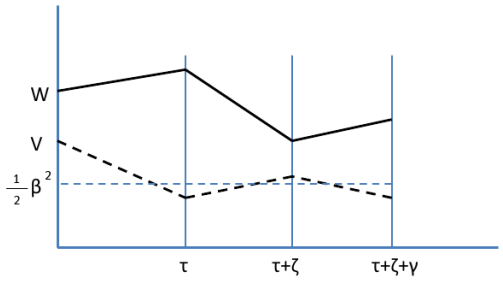

Fig. 4. Cartoon of V and W over time during switching. At time τ, the controller switches from $N_a = 0$ to $N_a = 1$, causing W to subsequently decrease. At time $\tau + \zeta$, the controller switches back to $N_a = 0$, causing a possible increase in W. At time $\tau + \zeta + \gamma$, when the controller switches to $N_a = 1$ again, W is lower than it was at time τ.

When $\|\mathbf{f}\| \leq \beta$, then $N_a = 1$ and $\dot{\mathbf{u}}_a = \dot{\mathbf{u}}_f + \dot{\mathbf{u}}_m$. In this situation, we have:

$$\begin{aligned} \dot{V}_{\|\mathbf{f}\|\leq\beta} &= \frac{\partial V}{\partial \mathbf{u}}(\dot{\mathbf{u}}_f + \dot{\mathbf{u}}_m) \\ &= -\begin{pmatrix} \alpha \\ 0 \\ -\alpha \\ 0 \end{pmatrix}^T K \begin{pmatrix} p-q+\alpha \\ s-a \\ p-q-\alpha \\ b-s \end{pmatrix} \\ &= -(\kappa_{11} + \kappa_{33})\alpha^2 - (\kappa_{11} - \kappa_{33})(p-q)\alpha \\ &\quad -\kappa_{12}(s-a)\alpha - \kappa_{34}(s-b)\alpha, \end{aligned}$$

where κ_{11} and κ_{33} are positive diagonal elements of K, and κ_{12} and κ_{34} are off-diagonal elements of K. The gradient of W when $\|\mathbf{f}\| \leq \beta$ is:

$$\begin{aligned} \dot{W}_{\|\mathbf{f}\|\leq\beta} &= \frac{\partial W}{\partial \mathbf{u}}(\dot{\mathbf{u}}_f + \dot{\mathbf{u}}_m) \\ &= -2\begin{pmatrix} p \\ s \\ -q \\ -s \end{pmatrix}^T \begin{pmatrix} p-q+\alpha \\ s-a \\ p-q-\alpha \\ b-s \end{pmatrix} \\ &= -2(p-q)^2 - 2s(s-a) - 2s(s-b) \\ &\quad -2(p+q)\alpha. \end{aligned} \tag{25}$$

The following theorem establishes convergence for the switching controller.

Theorem 3: Let the object be convex, second-order continuous, with finite maximum curvature. Then the switching grasp controller (Equation 22) converges to a threshold around unit frictionless moment equilibrium. The size of the threshold can be made arbitrarily small by decreasing β.

Proof:

To show convergence of $\mathbf{m}$, we show that each time the controller switches from $N_a = 0$ to $N_a = 1$, the value of W decreases until a threshold proportional to β is reached. Let τ be an arbitrary iteration of the controller when N_a has just switched from 0 to 1 such that $\|\mathbf{f}\| \leq \beta$ or when the controller has started execution in a configuration where $\|\mathbf{f}\| \leq \beta$. At time $\tau + \zeta$, $\dot{\mathbf{u}}_f + \dot{\mathbf{u}}_m$ has executed for ζ steps such that the

last step caused $\|\mathbf{f}\|$ to cross the switching threshold such that:

$$V_{\tau+\zeta} \leq \frac{1}{2}\beta^2 + \dot{V}^{\tau+\zeta-1}_{\|\mathbf{f}\|\leq\beta},$$

where $\dot{V}^{\tau+\zeta-1}_{\|\mathbf{f}\|\leq\beta}$ is the value of $\dot{V}_{\|\mathbf{f}\|\leq\beta}$ at time $\tau+\zeta-1$. At this point, the controller switches to $N_a = 0$ and $\dot{\mathbf{u}}_f$ executes for another γ iterations until $\|\mathbf{f}\| \leq \beta$ again. Suppose that each of the γ iterations causes V to change by at least $\dot{V}^{min}_{\|\mathbf{f}\|>\beta} = -(\kappa_{11}+\kappa_{33})\beta^2$. Then, the maximum integer number of iterations of $\dot{\mathbf{u}}_f$ required to bring V below $\frac{1}{2}\beta^2$ is:

$$\begin{aligned}\gamma &\leq \max\left\{0, \left\lceil \frac{V_{\tau+\zeta} - \frac{1}{2}\beta^2}{-\dot{V}^{min}_{\|\mathbf{f}\|>\beta}} \right\rceil\right\} \\ &\leq \left\lceil \frac{|\dot{V}^{\tau+\zeta-1}_{\|\mathbf{f}\|\leq\beta}|}{-\dot{V}^{min}_{\|\mathbf{f}\|>\beta}} \right\rceil \\ &\leq 2 + \frac{\Lambda_{\tau+\zeta-1}}{(\kappa_{11}+\kappa_{33})\beta},\end{aligned}$$

where

$$\Lambda_{\tau+\zeta-1} = (\kappa_{11}-\kappa_{33})(p-q) + \kappa_{12}(s-a) + \kappa_{34}(s-b)$$

evaluated at time $\tau+\zeta-1$.

Consider how W changes during the ζ iterations between times τ and $\tau+\zeta$ when $N_a = 1$:

$$\begin{aligned}W_{\tau+\zeta} - W_\tau &= \dot{W}^{\tau}_{\|\mathbf{f}\|\leq\beta} + \ldots + \dot{W}^{\tau+\zeta-1}_{\|\mathbf{f}\|\leq\beta} \\ &\leq \zeta \dot{W}^{\tau+\zeta-1}_{\|\mathbf{f}\|\leq\beta},\end{aligned}$$

where $\dot{W}^{t}_{\|\mathbf{f}\|\leq\beta}$ is the change in W caused by the t^{th} iteration of the controller. The last inequality above uses the fact that for small α, the magnitude of $\dot{W}_{\|\mathbf{f}\|\leq\beta}$ is minimized for small values of W (at time $\tau+\zeta-1$). Substituting Equation 25 into the above and using Lemma 5, we have:

$$\begin{aligned}W_{\tau+\zeta} - W_\tau &\leq 2\zeta\left[-H_{\tau+\zeta-1} + |p+q|\beta\right] \\ &\leq 2\zeta\left[-H_{\tau+\zeta-1} + \|\mathbf{r}_1\|\beta^2\right],\end{aligned}$$

where $H_{\tau+\zeta-1} = (p-q)^2 + s(s-a) + s(s-b)$ is evaluated at time $\tau+\zeta-1$ (when it is largest) and we have used the fact that the magnitude of $\mathbf{f}$ never exceeds β between time τ and $\tau+\zeta-1$.

Now, consider how W changes during the γ iterations between times $\tau+\zeta$ and $\tau+\zeta+\gamma$ when $N_a = 0$:

$$\begin{aligned}W_{\tau+\zeta+\gamma} - W_{\tau+\zeta} &= \dot{W}^{\tau+\zeta}_{\|\mathbf{f}\|>\beta} + \ldots + \dot{W}^{\tau+\zeta+\gamma-1}_{\|\mathbf{f}\|>\beta} \\ &\leq \gamma \dot{W}^{max}_{\|\mathbf{f}\|>\beta},\end{aligned}$$

where, as above, $\dot{W}^{t}_{\|\mathbf{f}\|\leq\beta}$ is the change in W caused by the t^{th} iteration of the controller and $\dot{W}^{max}_{\|\mathbf{f}\|\leq\beta} = \max_t\left\{\dot{W}^{t}_{\|\mathbf{f}\|\leq\beta}\right\}$. Using Lemma 5, we have:

$$W_{\tau+\zeta+\gamma} - W_{\tau+\zeta} \leq 2\gamma r_{max}\mathbf{f}^2_{\tau+\zeta},$$

where $\mathbf{f}_{\tau+\zeta}$ is the unit frictionless force residual at the beginning of the $\tau+\zeta$ controller iteration and r_{max} is the maximum value of $\|\mathbf{r}_1\|$ between time $\tau+\zeta$ and $\tau+\zeta+\gamma$.

We can use $V_{\tau+\zeta}$ to bound $\mathbf{f}^2_{\tau+\zeta}$:

$$\begin{aligned}\mathbf{f}^2_{\tau+\zeta} &= 2V_{\tau+\zeta} \\ &\leq \beta^2 + 2\dot{V}^{\tau+\zeta-1}_{\|\mathbf{f}\|\leq\beta} \\ &= \beta^2 - 2(\kappa_{11}+\kappa_{33})\beta^2 - 2\Lambda_{\tau+\zeta-1}\beta.\end{aligned}$$

Dropping the $\tau+\zeta-1$ subscript from Λ, we have:

$$\begin{aligned}W_{\tau+\zeta+\gamma} - W_{\tau+\zeta} &\leq 2\gamma r_{max} \\ &\quad \left[\beta^2 - 2(\kappa_{11}+\kappa_{33})\beta^2 - 2\Lambda\beta\right] \\ &\leq \left(4 + 2\frac{\Lambda}{(\kappa_{11}+\kappa_{33})\beta}\right) r_{max} \\ &\quad \left[\beta^2 - 2(\kappa_{11}+\kappa_{33})\beta^2 - 2\Lambda\beta\right].\end{aligned}$$

Combining $W_{\tau+\zeta} - W_\tau$ and $W_{\tau+\zeta+\gamma} - W_{\tau+\zeta}$, we have:

$$\begin{aligned}W_{\tau+\zeta+\gamma} - W_\tau &\leq 2r_{max}\beta^2\left[\zeta - 4(\kappa_{11}+\kappa_{33}) + 2\right] \\ &\quad +2\Lambda\beta\left[\frac{r_{max}}{\kappa_{11}+\kappa_{33}} - 4r_{max} - 2\right] \\ &\quad -2\zeta H_{\tau+\zeta-1} - \frac{4\Lambda^2}{\kappa_{11}+\kappa_{33}}. \qquad (26)\end{aligned}$$

The key thing to note about Equation 26 is that the first two terms are factors of β^2 and β, respectively, and, by applying Lemma 4, the last two terms are negative semi-definite. Therefore, each time the controller switches from $N_a = 0$ to $N_a = 1$, W decreases until $H_{\tau+\zeta-1}$ reaches a threshold around the origin no larger than the first two terms in Equation 26. Recall that p does not equal q and s does not equal a or b unless the two contacts are concurrent. However, since a concurrent configuration is prohibited by the assumption that execution does not begin with the contacts on the same face, H approaches zero only when p, q, s, a, and b approach zero. Since the first and second terms of Equation 26 can be made arbitrarily small by decreasing β, we can force the switching controller to converge to a configuration with p, q, s, a, and b arbitrarily close to zero. Since p, q, and s are zero only when $\mathbf{r}_1$ is orthogonal to the surface tangents at the contacts, we conclude that the system converges to a threshold around unit frictionless moment equilibrium that can be lowered by reducing β.

■

VII. Experiments

The three controllers proposed in this paper were compared with each other in simulation. The approximate null space grasp controller was also tested in practice using Dexter, a bimanual dexterous humanoid robot at the University of Massachusetts Amherst.

A. Experiment 1: Simulation

The simulations explored grasping a spatial ellipsoid with principle axis lengths 1, 2, and 3 using two contacts. In order to focus the experiment on the relative performance of the controllers in the absence of the effects of manipulator kinematics or control, the two contacts were modeled as free-floating points constrained to the surface of the ellipsoid. The switching controller executed with a force threshold parameter

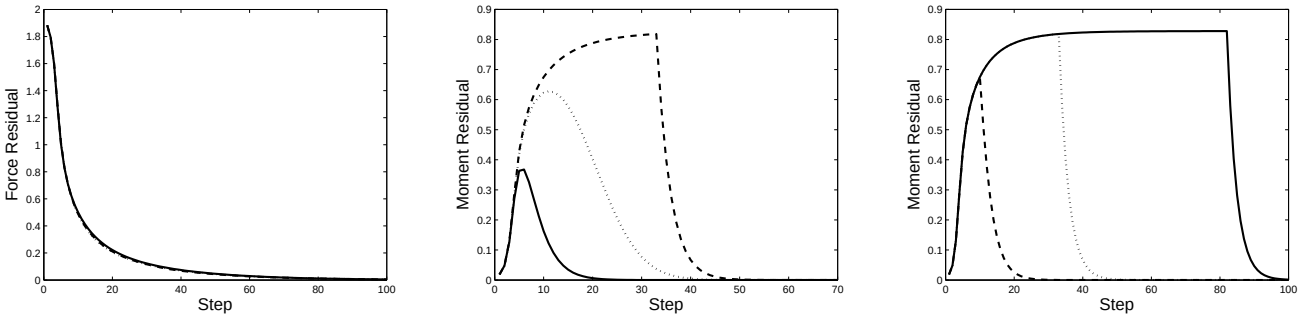

(a) (b) (c)

Fig. 5. Simulation results. The first two panels show unit frictionless force residual (a) and moment residual (b) as a function of controller iteration for a representative simulation of the three controllers. The solid line is the exact null space controller, the dotted line is the approximate null space controller, and the dashed line is the switching controller. The last panel (c) illustrates the performance of the switching controller for different values of β. The dashed line is $\beta = 0.5$, the dotted line is $\beta = 0.1$, and the solid line is $\beta = 0.01$.

of $\beta = 0.1$. The approximate null space controller evaluated the null space projection matrix, $\mathcal{N}(\hat{J}_f) = I - \hat{J}_f^+ \hat{J}_f$ (a^+ denotes the damped-least-squares inverse [40] of a), with a damping parameter of 0.01. In the simulations, there is no direct correspondence between step size and time. In general, the control law can be executed as fast as the mechanism for contact displacement and contact sensing allows.

The simulation was executed 100 times with the contacts initialized in randomly selected locations on the ellipsoid. All three controllers converged to a neighborhood around unit frictionless equilibrium in all cases. Figure 5(a) and (b) illustrate representative force and moment residual trajectories for the three controllers. Figure 5(c) compares the performance of the switching controller for three different values of β starting from the same initial contact configuration.

The results are consistent with what might intuitively be expected. All three controllers have essentially the same performance with respect to the unit frictionless force residual. This reflects the fact that before converging to a neighborhood around unit frictionless force equilibrium, all controllers follow essentially the same force residual control gradient. The three controllers differ in their unit frictionless moment residual performance. The exact controller converges the fastest, the approximate controller converges next fastest, and the switching controller converges slowest. We found that it was possible to change the relative performance of the approximate controller and the switching controller by adjusting the damping parameter and the β parameter, respectively. Although Figure 5(c) indicates that the switching controller works for the ellipsoid with high values for β, this is likely not to be true for arbitrary objects. Also, note that Figure 5(c) indicates apparently equal convergence of the controller to zero for all values of β. This suggests that the convergence bound derived at the end of the proof of Theorem 3 describes the worst-case behavior of the controller. We hypothesize that the switching controller will out-perform this bound for many objects.

(a) (b) (c)

Fig. 6. Three objects for which the grasp controller was tested.

B. Experiment 2: Dexter grasping a towel roll

This experiment was performed using Dexter, a bimanual dexterous humanoid robot at the University of Massachusetts Amherst. Dexter consists of two whole arm manipulators (WAMs), two Barrett hands equipped with six-axis load cells at the fingertips, and a Bisight stereo camera system. Contact displacements were realized by a hybrid force-position controller that applied a small inward force at each contact while displacing the contacts tangent to the surface. The contacts tracked the velocities specified by the grasp controller as closely as the manipulator kinematic constraints allowed.

The approximate null space grasp controller synthesized 58 two-contact grasps of the vertical towel roll (10cm diameter and 20cm tall) shown in Figure 6(a). On each trial, the grasp controller began execution in a randomly selected configuration relative to the object and continued until controller convergence or until the human operator detected that the manipulator had collided with the environment. Two of the three fingers on the Barrett hand were grouped together as a single contact (a virtual finger) [41], [42]. In this experiment and in experiment 3, computational time was negligible relative to the speed of arm motion.

Figures 7(a) and 7(b) show the density of hand orientations before and after executing the grasp controller. Hand orientation is measured by the angle between the line that connects the two virtual contacts and the towel roll major axis. The Figures show that for the vertical towel roll, the two-contact grasp controller aligned the hand orthogonal to the major axis of the cylinder on most of the grasp trials.

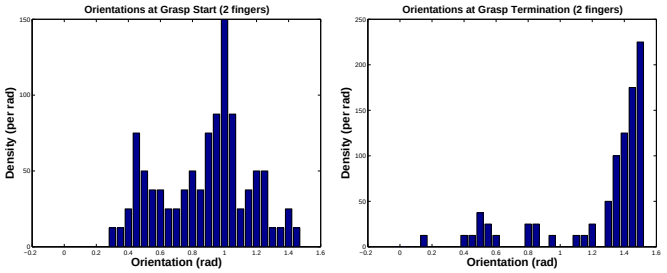

(a) (b)

Fig. 7. Experiment 2 (towel roll, two contacts): the distribution of contact orientations before, (a), and after, (b), the grasp controller has executed. Orientation is the angle between a line that passes between the two grasp contacts and the major axis of the object (see text).

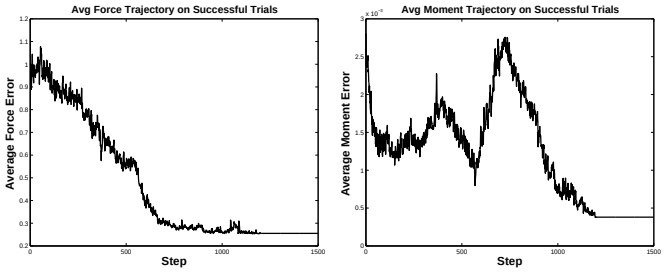

(a) (b)

Fig. 8. Experiment 2 (towel roll, two contacts): average squared force residual, (a), and average squared moment residual, (b), for the grasp trials that terminated near the peak at $\pi/2$ in Figure 7(b). The horizontal axis is in milliseconds.

However, on a few trials, the controller terminated near the small peak at 0.45 radians in Figure 7(b). These trials were terminated by the human operator because the Barrett hand palm collided with the object. These collisions highlight the fact that, without any provision for obstacle avoidance or configuration optimization, limitations on contact mobility may interfere with grasp controller performance. On these grasp trials, one of the grasp contacts was on the top of the cylinder while the other was on the side. As the grasp controller displaced the contacts around the object, it did not take the limited aperture of the Barrett hand into account and caused a collision.

Figure 8 illustrates the average force and moment residual error trajectories for the grasp trials that comprise the peak near $\pi/2$ in Figure 7(b). Notice that the moment residual error begins to converge only after convergence of the force residual controller is complete. This is consistent with the proofs of Theorems 2 and 3 that suggest that moment residual convergence depends on force residual convergence. Figure 8(a) shows the average force error (squared force residual) while Figure 8(b) shows the average moment error. The horizontal axis in both figures is grasp controller step. The graphs illustrate that, on average, both force and moment errors converge to configurations with small wrench residuals in approximately 1000 steps (20 seconds, not including the time taken to tare the fingertip load cells.)

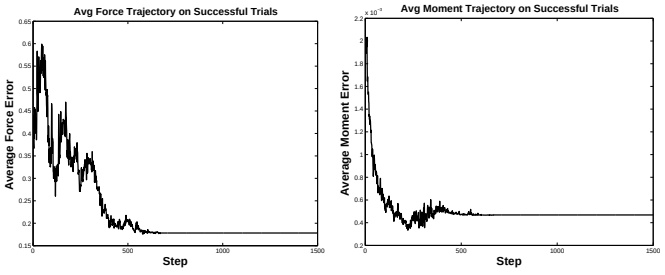

(a) (b)

Fig. 9. Experiment 3 (squirt bottle, two contacts): average force residual, (a), and moment residual, (b). The horizontal axis is in milliseconds.

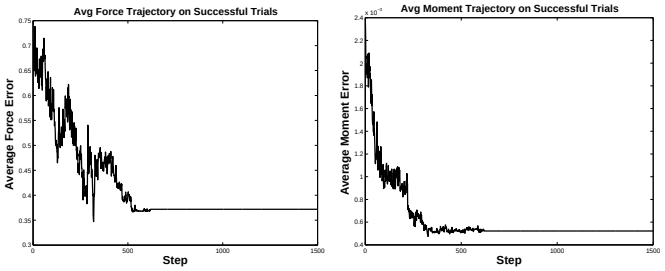

(a) (b)

Fig. 10. Experiment 3 (detergent bottle, two contacts): average force residual, (a), and moment residual, (b). The horizontal axis is in milliseconds.

C. Experiment 3: Grasping a Squirt Bottle and a Detergent Bottle

In the third experiment, the approximate null space grasp controller was executed for the squirt bottle and detergent bottle shown in Figures 6(b) and 6(c). The experimental procedure was the same as that used in experiments 2. On each trial, the grasp controller started from a randomly selected configuration. 28 grasp synthesis trials were executed for the squirt bottle and 31 grasps for the detergent bottle. Whereas the grasp controller had problems with kinematic limitations of the manipulator when grasping the cylinder, there were no such problems with the squirt and detergent bottles because, for these objects, the grasp controller tended away from grasp configurations that caused the manipulator to collide with the table. Figures 9 and 10 show that the grasp controller found low-error grasps for these objects. These results demonstrate that although the controllers are theoretically correct only for convex objects, they perform well for arbitrary objects in practice.

VIII. Conclusion

Rather than planning contact positions based on global geometric information, grasp control uses local contact measurements to synthesize grasps. An analogy can be drawn between the use of manipulator compliance in insertion tasks and grasp control in grasping tasks. In both cases, a control law using local force feedback is used to adjust what might initially be only an approximate solution. Both methods can

make the fine adjustments in manipulator configuration that are extremely difficult to achieve in other ways.

This paper focuses on a theoretical understanding of null space grasp control. Three different variants on the controller are proposed: exact null space grasp control, approximate null space grasp control, and switching grasp control. Exact control and switched control are theoretically demonstrated to converge to unit frictionless equilibrium contact configurations. Nevertheless, all three controllers are found to converge in simulation from arbitrary initial contact configurations. The approximate null space controller has been tested extensively using Dexter, a bimanual dexterous humanoid robot at the University of Massachusetts Amherst, and found to work well.

From a theoretical perspective, an important remaining question is whether convergence can be established for the approximate null space controller of Equation 21. Our experimental results suggest that this controller works well. However, the controller has not yet been shown to converge for all convex objects. From a broader perspective, there are many ways that force information might be used to assist robot grasping. Intuition suggests that humans rely on a sense of touch to grasp without looking at the object and to make grasping more robust. We expect that this will continue to be an important research question in the future.

APPENDIX

The following Lemma was used in Section III-B and is proven below.

Lemma 1: When the contacts can apply frictional torsional loads about the contact normal as well as tangential frictional forces, then a sufficient condition for three-dimensional two-finger force closure is non-marginal equilibrium.

Proof: Let $\mathbf{f}_1$ and $\mathbf{f}_2$ be equilibrium forces on the object. Let a_1 and a_2 be equilibrium contact moments (induced by the soft contacts) about the surface normals. Let $\mathbf{r}_1$ and $\mathbf{r}_2$ be the contact positions in a coordinate frame centered outside the object. Let $\mathbf{f}$ and $\mathbf{m}$ be the components of an arbitrary wrench applied to the object. Let β be the component of $\mathbf{m}$ orthogonal to $\mathbf{r}_1 - \mathbf{r}_2$. Let α be the other component.

Since the system is in equilibrium, we have that $\mathbf{f}_1 + \mathbf{f}_2 = 0$ and $\mathbf{r}_1 \times \mathbf{f}_1 + \mathbf{r}_2 \times \mathbf{f}_2 + a_1 + a_2 = 0$. Let $\mathbf{f}_1' = \mathbf{f}_1 - \mathbf{f} + \mathbf{v}$, $\mathbf{f}_2' = \mathbf{f}_2 - \mathbf{v}$, $a_1' = a_1 - \alpha$, and $a_2' = a_2$ where

$$\mathbf{v} = (\mathbf{x}_1 \times \mathbf{f} - \beta) \times (\mathbf{x}_1 - \mathbf{x}_2).$$

Then, we have that $\mathbf{f}_1' + \mathbf{f}_2' = -\mathbf{f}$ and $a_1' + a_2' + \mathbf{r}_1 \times \mathbf{f}_1' + \mathbf{r}_2 \times \mathbf{f}_2' = -\mathbf{m}$. Therefore, it is possible to resist an arbitrary wrench, $\mathbf{f}$ and $\mathbf{m}$, as long as $\mathbf{f}_1'$, $\mathbf{f}_2'$, a_1', and a_2' are within their friction cones. Following the argument in [33], for any force difference $\mathbf{c}$, it is possible to apply the net force, $\mathbf{f}_1' = \gamma \mathbf{f}_1 + \mathbf{c}$ by increasing γ sufficiently. Similarly, arbitrary moments about the contact normal can be applied. ■

The following three lemmas are used in Section VI.

Lemma 2: Let $(\mathbf{u}_1, \mathbf{v}_1, \mathbf{n}_1)$ and $(\mathbf{u}_2, \mathbf{v}_2, \mathbf{n}_2)$ be two orthonormal right-handed coordinate frames such that $\mathbf{v}_1 = \mathbf{v}_2$. Then $\mathbf{n}_1^T \mathbf{u}_2 = -\mathbf{n}_2^T \mathbf{u}_1$.

Proof:

Let R be a rotation matrix that describes the relationship between the two coordinate frames:

$$(\mathbf{u}_2, \mathbf{v}_2, \mathbf{n}_2) = R(\mathbf{u}_1, \mathbf{v}_1, \mathbf{n}_1).$$

Let Φ describe a 90 degree rotation about $\mathbf{v}_1 = \mathbf{v}_2$ such that:

$$\mathbf{n}_1 = \Phi \mathbf{u}_1,$$

and

$$\mathbf{n}_2 = \Phi \mathbf{u}_2.$$

Then:

$$\begin{aligned} \mathbf{n}_1^T \mathbf{u}_2 &= \mathbf{n}_1^T R \mathbf{u}_1 \\ &= (\Phi \mathbf{u}_1)^T R \Phi^T \mathbf{n}_1 \\ &= \mathbf{u}_1^T \Phi^T R \Phi^T \mathbf{n}_1. \end{aligned}$$

Since both Φ and R rotate about $\mathbf{v}_1 = \mathbf{v}_2$, these rotation matrices commute:

$$\Phi^T R \Phi^T = R \Phi^T \Phi^T.$$

However, notice that since Φ^T rotates through 90 degrees, $\Phi^T \Phi^T$ rotates through 180 degrees, or:

$$\begin{aligned} \mathbf{n}_1^T \mathbf{u}_2 &= \mathbf{u}_1^T \Phi^T R \Phi^T \mathbf{n}_1 \\ &= \mathbf{u}_1^T R \Phi^T \Phi^T \hat{\mathbf{n}}_1 \\ &= -\mathbf{u}_1^T R \mathbf{n}_1 \\ &= -\mathbf{u}_1^T \mathbf{n}_2. \end{aligned}$$

■

Lemma 3: Let $(\mathbf{u}_1, \mathbf{v}_1, \mathbf{n}_1)$ and $(\mathbf{u}_2, \mathbf{v}_2, \mathbf{n}_2)$ be two orthonormal right-handed coordinate frames such that $\mathbf{v}_1 = \mathbf{v}_2$ and $\mathbf{n}_1^T \mathbf{n}_2 \leq 0$. Then

$$\|\mathbf{n}_2^T \mathbf{u}_1\| \leq \|\mathbf{f}\| \leq \sqrt{2} \|\mathbf{n}_2^T \mathbf{u}_1\|,$$

where $\mathbf{f} = \mathbf{n}_1 + \mathbf{n}_2$.

Proof: Let β be the magnitude of the angle between $\mathbf{n}_1$ and $\mathbf{n}_2$. Since $\mathbf{n}_1^T \mathbf{n}_2 \leq 0$, then β must be bounded by: $90 \leq \beta \leq 180$.

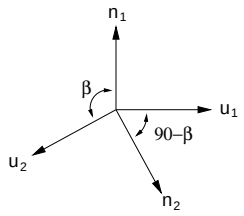

Fig. 11. Geometry of β.

Since $\mathbf{v}_1 = \mathbf{v}_2$, then $\mathbf{n}_1$, $\mathbf{n}_2$, $\mathbf{u}_1$, and $\mathbf{u}_2$ lie in a plane. By the geometry of the situation, the magnitude of the angle between $\mathbf{n}_2$ and $\mathbf{u}_1$ is $\beta - 90$ (see Figure 11). Therefore:

$$\begin{aligned} \|\mathbf{n}_2^T \mathbf{u}_1\| &= \cos(\beta - 90) \\ &= \|\sin \beta\| \\ &= 2 \|\sin \frac{\beta}{2} \cos \frac{\beta}{2}\|. \end{aligned}$$

Let $\mathbf{h}$ be the unit vector such that:

$$\mathbf{n}_1 + \mathbf{n}_2 = \|\mathbf{f}\|\mathbf{h},$$

and $\|\mathbf{f}\| = \|\mathbf{n}_1 + \mathbf{n}_2\|$. Then

$$\|\mathbf{f}\| = \|\mathbf{h}^T\mathbf{n}_1\| + \|\mathbf{h}^T\mathbf{n}_2\|.$$

Let γ be the angle between $\mathbf{h}$ and $\mathbf{n}_1$ such that

$$\begin{aligned}\cos\frac{\beta}{2} &= \cos\gamma \\ &= \|\mathbf{h}^T\mathbf{n}_1\| \\ &= \|\mathbf{h}^T\mathbf{n}_2\| \\ &= \frac{\|\mathbf{f}\|}{2}.\end{aligned}$$

Since $\sin\frac{\beta}{2} \leq 1$, we have that:

$$\begin{aligned}\|\mathbf{n}_2^T\mathbf{u}_1\| &= \|2\sin\frac{\beta}{2}\cos\frac{\beta}{2}\| \\ &\leq 2\|\cos\frac{\beta}{2}\| \\ &\leq \|\mathbf{f}\|.\end{aligned}$$

Also, since $\beta \geq 90$ by assumption, then $\sin\frac{\beta}{2} \geq \frac{1}{\sqrt{2}}$ and we have:

$$\begin{aligned}\|\hat{\mathbf{n}}_2^T\mathbf{u}_1\| &= \|2\sin\frac{\beta}{2}\cos\frac{\beta}{2}\| \\ &\geq \sqrt{2}\|cos\frac{\beta}{2}\| \\ &\geq \frac{1}{\sqrt{2}}\|\mathbf{f}\|.\end{aligned}$$

Combining the above bounds on $\|\mathbf{f}\|$, we have:

$$\|\hat{\mathbf{n}}_2^T\mathbf{u}_1\| \leq \|\mathbf{f}\| \leq \sqrt{2}\|\hat{\mathbf{n}}_2^T\mathbf{u}_1\|.$$

■

Lemma 4: Let $(\mathbf{u}_1, \mathbf{v}_1, \mathbf{n}_1)$ and $(\mathbf{u}_2, \mathbf{v}_2, \mathbf{n}_2)$ be two coordinate frames such that $\mathbf{v}_1 = \mathbf{v}_2$ and $\mathbf{n}_1^T\mathbf{n}_2 \leq 0$. Let $s = \mathbf{r}^T\mathbf{v}_1$, $a = \mathbf{r}^T(\mathbf{n}_2 \times \mathbf{u}_1)$, and $b = \mathbf{r}^T(\mathbf{n}_1 \times \mathbf{u}_2)$ for an arbitrary vector, $\mathbf{r}$. Then $s = 0$ implies that $a = 0$ and $b = 0$. Also, $sa \leq 0$ and $sb \leq 0$.

Proof: Since $\mathbf{v}_1 = \mathbf{v}_2$, we have that $\mathbf{u}_1$,$\mathbf{n}_1$, $\mathbf{u}_2$, and $\mathbf{n}_2$ are orthogonal to $\mathbf{v}_1$. Therefore, $\mathbf{n}_2 \times \mathbf{u}_1 = \gamma\mathbf{v}_1$ and $\mathbf{n}_1 \times \mathbf{u}_2 = \eta\mathbf{v}_1$ where $\gamma = (\mathbf{n}_2 \times \mathbf{u}_1)^T\mathbf{v}_1$ and $\eta = (\mathbf{n}_1 \times \mathbf{u}_2)^T\mathbf{v}_1$. a and b can be rewritten: $a = \gamma s$ and $a = \eta s$. Therefore, we have that $s = 0$ implies that $a = 0$ and $b = 0$.

Note that γ and η must be negative:

$$\begin{aligned}\gamma &= (\mathbf{n}_2 \times \mathbf{u}_1)^T\mathbf{v}_1 \\ &= \mathbf{n}_2^T(\mathbf{u}_1 \times \mathbf{v}_1) \\ &= \mathbf{n}_2^T\mathbf{n}_1 \\ &\leq 0,\end{aligned}$$

and

$$\begin{aligned}\eta &= (\mathbf{n}_1 \times \mathbf{u}_2)^T\mathbf{v}_1 \\ &= \mathbf{n}_1^T(\mathbf{u}_2 \times \mathbf{v}_2) \\ &= \mathbf{n}_1^T\mathbf{n}_2 \\ &\leq 0.\end{aligned}$$

We can conclude that sa and sb are negative because:

$$\begin{aligned}sa &= \mathbf{r}^T\mathbf{v}_1\mathbf{r}^T(\mathbf{n}_2 \times \mathbf{u}_1) \\ &= \gamma\mathbf{r}^T\mathbf{v}_1\mathbf{v}_1^T\mathbf{r} \\ &\leq 0,\end{aligned}$$

and

$$\begin{aligned}sb &= \mathbf{r}^T\mathbf{v}_1\mathbf{r}^T(\mathbf{n}_1 \times \mathbf{u}_2) \\ &= \eta\mathbf{r}^T\mathbf{v}_1\mathbf{v}_1^T\mathbf{r} \\ &\leq 0.\end{aligned}$$

■

Lemma 5: Let $(\mathbf{u}_1, \mathbf{v}_1, \mathbf{n}_1)$ and $(\mathbf{u}_2, \mathbf{v}_2, \mathbf{n}_2)$ be two orthogonal coordinate frames such that $\mathbf{v}_1 = \mathbf{v}_2$. Then $|\mathbf{r}^T\mathbf{u}_1 + \mathbf{r}^T\mathbf{u}_2| \leq \|\mathbf{r}\|\|\mathbf{n}_1 + \mathbf{n}_2\|$.

Proof: Notice that $\mathbf{n}_1$ and $\mathbf{u}_1$ are related by the same rotation matrix that relates $\mathbf{n}_2$ and $\mathbf{u}_2$: $\mathbf{n}_1 = R\mathbf{u}_1$ and $\mathbf{n}_2 = R\mathbf{u}_2$. Therefore:

$$\begin{aligned}\|\mathbf{u}_1 + \mathbf{u}_2\| &= \|R^T\mathbf{n}_1 + R^T\mathbf{n}\| \\ &= \|\mathbf{n}_1 + \mathbf{n}_2\|\end{aligned}$$

and

$$\|\mathbf{r}^T\mathbf{u}_1 + \mathbf{r}^T\mathbf{u}_2\| \leq \|\mathbf{r}_1\|\|\mathbf{n}_1 + \mathbf{n}_2\| \tag{27}$$

where the last inequality used Lemma 3. ■

REFERENCES

[1] J. Coelho and R. Grupen, "A control basis for learning multifingered grasps," *Journal of Robotic Systems*, vol. 14, no. 7, pp. 545–557, 1997.

[2] J. Coelho, "Multifingered grasping: Grasp reflexes and control context," Ph.D. dissertation, University of Massachusetts, 2001.

[3] V. Nguyen, "Constructing force-closure grasps," in *IEEE Int'l Conf. Robotics Automation*, vol. 3, April 1986, pp. 1368–1373.

[4] ——, "Constructing stable grasps in 3d," in *IEEE Int'l Conf. Robotics Automation*, vol. 4, March 1987, pp. 234–239.

[5] J. Ponce and B. Faverjon, "On computing three-finger force-closure grasps of polygonal objects," *IEEE Transactions on Robotics and Automation*, vol. 11, no. 6, pp. 868–881, 1995.

[6] I.-M. Chen and J. Burdick, "Finding antipodal point grasps on irregularly shaped objects," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 4, pp. 507–512, 1993.

[7] Y.-B. Jia, "Computation on paramtric curves with an application in grasping," *International Journal of Robotics Research*, vol. 23, no. 7-8, pp. 825–855, 2004.

[8] A. Sudsang and J. Ponce, "New techniques for computing four-finger force-closure grasps of polyhedral objects," in *IEEE Int'l Conf. Robotics Automation*, vol. 2, May 1995, pp. 1355–1360.

[9] L. Han, Y. Guan, Q. Li, Z. Shi, and J. Trinkle, "Dextrous manipulation with rolling contacts," in *IEEE Int'l Conf. Robotics Automation*, 1997.

[10] L. Han and J. Trinkle, "Dextrous manipulation by rolling and finger gaiting," in *IEEE Int'l Conf. Robotics Automation*, 1998, pp. 992–997.

[11] M. Cherif and K. Gupta, "Planning for in-hand dextrous manipulation," in *Third Workshop on the Algorithmic Foundations of Robotics*, 1998, pp. 103–117.

[12] Z. Li and S. Sastry, "Task-oriented optimal grasping by multifingered robot hands," in *IEEE Int'l Conf. Robotics Automation*, vol. 4, March 1987, pp. 389–394.

[13] C. Ferrari and J. Canny, "Planning optimal grasps," in *IEEE Int'l Conf. Robotics Automation*, May 1992, pp. 2290–2295.

[14] D. Kirkpatrick, B. Mishra, and C. Yap, "Quantitative steinitz's theorems with applications to multifingered grasping," in *20th ACM Symp. on Theory of Computing*, May 1990, pp. 341–351.

[15] B. Mirtich and J. Canny, "Easily computable optimum grasps in 2-d and 3-d," in *IEEE Int'l Conf. Robotics Automation*, 1994, pp. 739–747.

[16] A. Hauck, J. Ruttinger, M. Sorg, and G. Farber, "Visual determination of 3d grasping points on unknown objects witha binocular camera system," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 1999, pp. 272–278.

[17] A. Morales, G. Recatala, P. Sanz, and A. Pobil, "Heuristic vision-based computation of planar antipodal grasps on unknown objects," in *IEEE Int'l Conf. Robotics Automation*, 2001, pp. 583–588.

[18] K. Stanley, J. Wu, A. Jerbi, and W. Gruver, "A fast two dimensional image based grasp planner," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 1999, pp. 266–271.

[19] D. Perrin, O. Masoud, C. Smith, and N. Papanikolopoulos, "Unknown object grasping using statistical pressure models," in *IEEE Int'l Conf. Robotics Automation*, 2000, pp. 1054–1059.

[20] T. Yoshikawa, M. Koeda, and H. Fujimoto, "Shape recognition and grasping by robotic hands with soft fingers and omnidirectional camera," in *IEEE Int'l Conf. Robotics Automation*, 2008, pp. 299–304.

[21] C. Borst, M. Fischer, and G. Hirzinger, "Calculating hand configurations for precision and pinch grasps," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 2002.

[22] A. Morales, P. Sanz, A. Pobil, and A. Fagg, "Vision-based three-finger grasp synthesis constrained by hand geometry," *Robotics and Autonomous Systems*, vol. 54, no. 6, pp. 496–512, 2006.

[23] P. Allen and P. Michelman, "Acquisition and interpretation of 3-d sensor data from touch," *IEEE Transactions on Robotics and Automation*, vol. 6, no. 4, pp. 397–404, 1990.

[24] Y. Jia and M. Erdmann, "Pose and motion from contact," *International Journal of Robotics Research*, vol. 18, no. 5, pp. 466–490, 1999.

[25] S. Haidacher and G. Hirzinger, "Estimating finger contact location and object pose from contact measurements in 3-d grasping," in *IEEE Int'l Conf. on Robotics and Automation*, 2003, pp. 1805– 1810.

[26] S. Chhatpar and M. Branicky, "Particle filtering for localization in robotic assemblies with position uncertainty," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 2005, pp. 3610– 3617.

[27] A. Petrovskaya, O. Khatib, S. Thrun, and A. Ng, "Bayesian estimation for autonomous object manipulation based on tactile sensors," in *IEEE Int'l Conf. on Robotics and Automation*, 2006, pp. 707–714.

[28] K. Gadeyne and H. Bruyninckx, "Markov techniques for object localization with force-controlled robots," in *10th Int'l Conf. on Advanced Robotics*, 2001.

[29] J. Son, R. Howe, J. Wang, and G. Hager, "Preliminary results on grasping with vision and touch," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, vol. 3, November 1996, pp. 1068–1075.

[30] B. Yoshimi and P. Allen, "Integrating real-time vision and manipulation," in *Proc. of 13th Hawaii Int'l Conf. on System Sciences*, vol. 5, January 1997, pp. 178–187.

[31] M. Teichmann and B. Mishra, "Reactive algorithms for 2 and 3 finger grasping," in *International Symposium on Intelligent Robotic Systems*, July 1994, pp. 1931–1936.

[32] B. Mishra, J. Schwartz, and M. Sharir, "On the existence and synthesis of multifinger positive grips," *Algorithmica*, vol. 2, no. 1-3, pp. 541–558, 1987.

[33] J. Ponce, S. Sullivan, A. Sudsang, J. Boissonnat, and J. Merlet, "On computing four-finger equilibrium and force-closure grasps of polyhedral objects," *Int. J. Rob. Res.*, pp. 11–35, 1996.

[34] R. Murray, Z. Li, and S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. CRC Press, 1994.

[35] L. Sciavicco and B. Siciliano, *Modelling and Control of Robot Manipulators*, 2nd ed. Springer, 2000.

[36] A. Bicchi, J. Salisbury, and D. Brock, "Contact sensing from force measurements," *International Journal of Robotics Research*, vol. 12, no. 3, 1993.

[37] M. Teichmann and B. Mishra, "Reactive algorithms for grasping using a modified parallel jaw gripper," in *IEEE Int'l Conf. Robotics Automation*, vol. 3, May 1994, pp. 1931–1936.

[38] S. Walker and K. Salisbury, "Pushing using learned manipulation maps," in *IEEE Int'l Conf. on Robotics and Automation*, 2008, pp. 3808–3813.

[39] R. Wistort and J. Smith, "Electric field servoing for robotic manipulation," in *IEEE Int'l Conf. on Intelligent Robots and Systems*, 2008, pp. 494–499.

[40] C. Wampler, "Manipulator inverse kinematic solutions based on vector formulations and damped least squares methods," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 16, pp. 93–101, 1986.

[41] R. Platt, "Learning and generalizing control-based grasping and manipulation skills," Ph.D. dissertation, University of Massachusetts, September 2006.

[42] R. Platt, A. H. Fagg, and R. Grupen, "Extending fingertip grasping to whole body grasping," in *IEEE Int'l Conference on Robotics and Automation*, 2003, pp. 2677– 2682.

Robert Platt Jr. Robert Platt is a Research Scientist in the Computer Science and Aritificial Intelligence Laboratory at the Massachusetts Institute of Technology. Between 2005 and 2009, he was a researcher at NASA Johnson Space Center working on the Robonaut project. He received a Ph.D. in Computer Science in 2006 from the University of Massachusetts, Amherst. He is interested in robot manipulation and assembly with an emphasis on perception and control in noisy and imperfectly observed environments.

Andrew H. Fagg Andrew H. Fagg is an Associate Professor of Computer Science at the University of Oklahoma. He holds a BS in Applied Mathematics/Computer Science from Carnegie-Mellon University, and a MS and a PhD in Computer Science from the University of Southern California. His research focuses on the computational issues surrounding the symbiotic relationships between humans and machines. In particular, he is interested in primate and robot learning of motor skills and task-oriented representations; reaching, grasping, and manipulation; brain-machine interfaces; and interactive art.

Roderic A. Grupen Rod Grupen is a professor of Computer Science at the University of Massachusetts Amherst where he directs the Laboratory for Perceptual Robotics. Grupen received a B.A. in Physics from Franklin and Marshall College, a B.S. in Mechanical Engineering from Washington University, a M.S. in Mechanical Engineering from Penn State University, and a Ph.D. in Computer Science from the University of Utah in 1988. His research concerns integrated mobile manipulation in systems that learn autonomously and acquire cumulative knowledge during extended interactions with the environment applied to personal robotics and healthcare.