

# Accurate Energetic Constraints for Passive Grasp Stability Analysis

Maximilian Haas-Heger\*, Matei Ciocarlie†

\*email: m.haas@columbia.edu †email: matei.ciocarlie@columbia.edu

Department of Mechanical Engineering; Columbia University, New York, NY 10027

**Abstract**—Passive reaction effects in grasp stability analysis occur when the contact forces and joint torques applied by a grasp change in response to external disturbances applied to the grasped object. For example, nonbackdrivable actuators (e.g. highly geared servos) will passively resist external disturbances without an actively applied command; for numerous robot hands using such motors, these effects can be highly beneficial as they increase grasp resistance without requiring active control. We introduce a grasp stability analysis method that can model these effects, and, for a given grasp, distinguish between disturbances that will be passively resisted and those that will not. We find that, in order to achieve this, the grasp model must include accurate energetic constraints. One way to achieve this is to consider the Maximum Dissipation Principle (MDP), a part of the Coulomb friction model that is rarely used in grasp stability analysis. However, the MDP constraints are non-convex, and difficult to solve efficiently. We thus introduce a convex relaxation method, along with an algorithm that successively refines this relaxation locally in order to obtain solutions to arbitrary accuracy efficiently. Our resulting algorithm can determine if a grasp is passively stable, solve for equilibrium contact forces and compute optimal actuator commands for stability. Its implementation is publicly available as part of the open-source GraspIt! simulator.

## I. INTRODUCTION

The analysis of the stability of a grasp is a foundational aspect of multi-fingered robotic manipulation. Determining the ability of a grasp to resist given disturbances, formulated as external wrenches applied to the grasped object, is equivalent to computing the stability of a multi-body system with frictional contacts under applied loads. Problems of this kind are thus pervasive in grasp analysis and may be encountered in many other scenarios that require simulation of general rigid bodies with frictional contacts.

A major complication in the simulation of multi-body systems is the accurate modeling of friction. Of particular interest to us is the Maximum Dissipation Principle (MDP), which is a part of the Coulomb friction model. This principle, which informally states that friction attempts to dissipate as much energy as possible in the presence of motion, imposes constraints that are both non-smooth and non-convex [44]. Existing exact formulations of the MDP are NP-complete [37], and therefore do not allow for efficient computation of solutions. To avoid this, current grasp models do not include the MDP at all, and therefore do not suffer from the same




Fig. 1: A grasping scenario where a hand establishes multiple frictional contacts (numbered 1-4) with a target object. A disturbance pushing the object “up” (such as  $w_1$ ) can be resisted (by contact forces  $c_2$  and  $c_4$ ), but only if a preload has been actively applied by the hand. In contrast, a disturbance pushing “down” ( $w_2$ ) will be resisted (by  $c_1$  and  $c_3$ ) passively, without the need for any preload, as long as the joint are not backdrivable. By correctly modeling energy conservation constraints, our proposed grasp model can capture such effects.

complexity. In consequence, however, these grasp models cannot capture what we call *passive stability*. As we argue in this paper, in order to determine the passive stability of a grasp, inclusion of the MDP in the friction model is imperative in order to maintain energy conservation laws.

In robotic manipulation, passive resistance arises due to the highly geared actuation of most commonly used hands. Large gearing ratios between the actuators and the joints mean the fingers are nonbackdrivable thus providing potentially rigid constraints for the grasped object. To illustrate this concept, consider the grasp in Fig. 1. Does the grasp remain stable if we apply either disturbance  $w_1$  or  $w_2$  to the grasped object? In order to resist those disturbances, contact forces must arise that balance them. In either case, there exist contact forces that satisfy a simple friction law (illustrated by red friction cones) and balance the disturbance. These contact forces in turn must be balanced by torques at the finger joints. However, it is clear that contact forces  $c_2$  and  $c_4$  can only arise if we have previously loaded the joints such that there is sufficient normal force at contacts 2 and 4 to sustain the friction

forces required. An appropriate *preload* is required, where the actuators ‘squeeze’ the object prior to the application of  $w_1$ . In contrast, assuming the joints are nonbackdrivable, contact forces  $c_1$  and  $c_3$  will arise in response to  $w_2$  irrespective of preload, and with no action required on the part of the actuators. We refer to this phenomenon as *passive resistance*.

Passive resistance allows practitioners to greatly simplify robotic grasping. We can often apply actuator torques that close the fingers around the object without necessarily worrying if these will balance out once contact is made: the fingers jam as the hand squeezes the object, and the gearboxes between joints and actuators provide additional structural torques. If the grasp geometry is adequately chosen, the equilibrium joint torques arise passively when the fingers squeeze the object between them and a stable grasp arises. For example, in the grasp in Fig. 1 it is sufficient to actively load the joints of one finger. The nonbackdrivability of the other finger means the object will be stably grasped and equilibrium joint torques arise passively in the non-actuated finger.

The same phenomenon can allow a grasp to withstand a range of disturbances applied to the object without a change in the actuator commands. If chosen wisely, the initially applied preload is sufficient to balance the object throughout the task and various corresponding different disturbances. This is the true power of passive resistance. However, while this approach works well in practice, the field currently lacks the tools to formally analyze this effect: we need a method that can determine which range of disturbances will be passively resisted given a specific preload.

A grasp model that accounts for these effects must capture the interplay of contact forces, joint torques and external wrenches. It must be able to accurately predict how joint torques and external wrenches are transmitted through the object and distributed across the contacts, which is complicated by the fact that in general robotic grasps are statically indeterminate (or hyperstatic) [42]. Nonlinearities due to the unilateral nature of contacts (a contact can only push, not pull) as well as the nonbackdrivability of highly geared joints complicate the analysis of such problems even in the frictionless case. As we have mentioned earlier, the introduction of friction poses even greater difficulty, as the Coulomb friction model is both nonsmooth and nonlinear, and inclusion of the MDP makes it non-convex.

Despite significant advances, no model proposed to date meets all these requirements. As we discuss in detail in Sec. II, some existing models make assumptions that are not met in practical manipulation tasks. Others approximate one or more of the constraints above, are applicable only in two dimensions, or do not provide convergence guarantees. Thus, the state-of-the-art grasp analysis tools cannot capture passive stability phenomena.

In this paper, we introduce a quasi-static model that addresses all the constraints above for general, three-dimensional grasps. We show that inclusion of the MDP is vital to the physical accuracy and usefulness of such a model and present a convex relaxation that allows for efficient solution of problems involving the MDP. While our model is based on a linear approximation of friction cones, we introduce a

computationally efficient method that can successively tighten this approximation up to the desired accuracy. A very useful property of tightening approaches is that they allow for strong guarantees: if, at any stage of refinement, our model fails to find a solution, we can guarantee that no solution exists to the exact problem. This allows for early exit from computation in cases where equilibrium cannot exist. It is, to the best of our knowledge, the first time that a model has been proposed that can handle three-dimensional frictional constraints that include the MDP, up to arbitrary accuracy (and thus approaching the solution to the exact problem) and in a computationally efficient fashion.

Depending on the choice of variables and optimization objective, our model can be used for a wide range of queries. In this paper, we illustrate its applicability to quasi-static grasp stability analysis by answering multiple queries on a number of example grasps. The queries we show here include: Given applied joint torques, will the grasp be stable in the presence of a specified external disturbance, assuming passive resistance effects? Alternatively, given applied joint torques, what is the largest disturbance that can be passively resisted in a given direction? Finally, given a disturbance, what are the optimal joint torques that a grasp can apply for stability? We believe these are all useful tools in the context of grasp analysis, and plan to expand the use of this model to other types of problems in the future.

## II. RELATED WORK

### A. Grasp Force Distribution

The most basic approaches to grasp stability analysis are concerned with the existence of stabilizing contact forces. A grasp has *force closure* if the hand can apply arbitrary wrenches to the object through the contacts. Analysis of force closure grasps dates back to Reuleaux (1876) [40]. Salisbury [41] proposed an analytical method to test for force closure. However, perhaps the most commonly used method to test for force closure was introduced by Ferrari and Canny [16] who developed an efficient geometric method for computing the space of possible resultant wrenches a hand can impart on a grasped object with contact forces obeying a linearized friction constraint. Thus, their algorithm can answer what we will call the *existence problem*: Given a disturbance applied to the grasped object, are there contact forces that satisfy a set of simplified, linearized friction constraints, and can balance the disturbance? In case of a positive answer their method can furthermore determine the magnitude of the equilibrium contact forces relative to the magnitude of a worst-case disturbance - a useful property, which has been utilized as a grasp quality metric in numerous planning algorithms proposed since. Kirkpatrick et al. [25] had previously described a similar metric for frictionless contacts.

However, the specific choice of grasp force, or as Bicchi [9] calls it the *force distribution problem* is crucial for the stability of a grasp. Salisbury et al. [42] investigated the conditions for a grasp to become overconstrained for a variety of different contact types and were the first to express contact wrenches as the sum of a particular and a homogeneous solution. From this

it followed that the static indeterminacy could be alleviated if contact wrenches can be actively controlled [24]. This insight provided the foundation for the large body of work on grasp force optimization, which is concerned with finding optimal grasp forces in the space of contact forces possible under a friction law. Given a positive answer to the *existence problem* we can compute a set of contact forces, which will balance a given disturbance and are optimal with respect to some objective (e.g. minimizing the magnitude of the forces.)

The first works posing the synthesis of grasp forces as a convex optimization problem were by Kerr and Roth [24] as well as Nakamura et al. [33]. While these works required a linearization of the friction cone, later contributions [13, 20] have proposed formulating the friction cone constraints as Linear Matrix Inequalities (LMI) such that no linearization is necessary. More recently Boyd et al. [11] formulated the problems as second-order cone problems (SOCPs). They proposed a custom interior-point algorithm that exploits the structure of force optimization problems and is thus very efficient.

These methods allow us to compute optimal contact forces and the joint torques necessary to balance them. Thus, for any specific wrench on the object encountered throughout a task we can compute the specific optimal joint torques for stability. In order to use this in practice, however, we have to make a string of assumptions:

- First, we assume perfect knowledge of the disturbance to the object that must be balanced at all times;
- Second, we assume that we can actively control the contact forces at every contact;
- Third, we assume that we can actively control the joint torques required for equilibrium;
- Fourth and finally, we assume that we can accurately control the torque output of the hand actuators.

In the majority of robotic manipulation tasks these assumptions do not hold. First, the exact disturbance wrench acting on an object is difficult to compute - it requires knowledge of the mass and inertial properties as well as the exact trajectory of the object. Any additional disturbance can not be accounted for unless the fingers are equipped with tactile sensors. Second, many robotic hands are kinematically deficient and contain links with limited mobility. This means, for example, that we can not directly control the contact force at a contact on the palm of the hand for instance. Forces at such a contact can only arise passively by transmission of the disturbance on the grasped object or forces at other contacts through the object.

Third, the kinematics of the hand may not permit explicit control of the torques at every individual joint. This is the case for the class of underactuated hands, where joint torques by definition may not be independently controlled but are determined by the kinematic composition of the hand. Finally, most robotic hands use highly geared motors, which makes accurate sensing and control of the torques at the hand joints all but impossible.

Under these assumptions a positive answer to the *existence problem* is a necessary but by no means a sufficient condition for the ability of a grasp to resist a given disturbance. It just indicates the existence of contact forces that satisfy the friction laws and can balance the applied disturbance but does

not guarantee they will arise, unless we assume that we are actively controlling the contact forces to that end.

A method to alleviate some of the limitations due to the above assumptions is to take into account the flexibility of the object and the elements of the hand. Salisbury [41] used this approach to derive the stiffness matrix of a grasp and developed a framework to test for the stability of a grasp. Nguyen [34] modeled each contact as a virtual spring (first introduced by Hanafusa et al. [21]) and showed that any force closure grasp can be made stable by applying appropriate forces at each contact. Cutkosky et al. [14] extended this work to take into account the compliance of structural elements of the hand and object as well as effects due to changing geometry such as contact location changes due to rolling. However, these compliance models cannot accurately model nonlinearities due to breaking contacts or friction.

Bicchi [8, 9, 10] was the first to point out the limitations of an assumption that is central to the works of early authors such as Salisbury [42, 41] and Kerr [24]. Bicchi showed that the force generation capabilities of kinematically deficient hands are limited (no active control over forces at palm contacts for instance.) He proposed a decomposition of the space of possible contact forces into those actively controllable and those that may only arise passively taking into account the kinematics of the hand. This decomposition can be used to synthesize optimal grasp forces or to derive a quantitative grasp quality metric. However, the compliance model used [14] uses a linear friction model that does not satisfy any of the constraints associated with Coulomb friction. Prattichizzo et al. [39] derived grasp quality metrics using the same compliance model but somewhat alleviated the issues of a linear friction model by distinguishing between sticking, breaking and sliding contacts. Under their model sliding contacts may not exhibit any frictional forces at all and hence the results are overly conservative. Furthermore, the run-time of their algorithm grows exponentially with the number of contacts.

These 'traditional' grasp stability metrics have become a common tool in many grasp planning algorithms (to train the Dex-Net grasp planner [27] for instance.) Seldom, however, are actuator commands actively controlled in response to measured disturbance as described above. Instead, the actuators are commanded to close the fingers around an object and often a stable grasp arises. The reasons for this lie in the design of most robotic hands in use today: Nonbackdrivable joints provide passive resistance, which allows for equilibrium contact forces to arise by themselves. This phenomenon is not at all well understood.

## B. Rigid Body Kinematics

All analyses introduced thus far assume that we have some degree of control over the contact wrenches. Palmer [35] investigated the stability of arrangements of rigid polygonal bodies without this simplifying assumption and showed that the problem of determining stability is co-NP complete. Matickalli et al. [29] pointed out that the existence of a solution to the equilibrium equations is in fact a necessary but not

sufficient condition for stability. Pang et al. [36] pointed out that the equations of equilibrium are insufficient for the determination of stability of rigid body contact problems due to the possibility of false positives. Here lies the difficulty of the problem introduced in Section I and Fig. 1.

Trinkle et al. [50] developed theory in order to predict the motion of a rigid body in the plane under quasistatic assumptions (i.e. inertial effects are considered negligible) and Coulomb friction - specifically including the *Maximum Dissipation Principle* (MDP) [44]. This principle states that given relative motion at the contact the frictional forces must do maximum work and is of foundational importance for the modeling of rigid body mechanics. In further work [37] Pang et al. cast the contact constraints as an uncoupled complementarity problem (UCP) and showed that problems of this type are NP-complete. In our own work [19] we showed that exact solutions of rigid body equilibrium with breaking and sliding contacts can be obtained for planar grasps in polynomial time. The algorithm presented makes use of the piece-wise convex behavior of friction forces in two dimensions and the contact motion constraints imposed by rigid body assumptions (the latter observation having been made by Mason [28].)

For three dimensional grasps Baraff [5, 6] suggested iterative schemes to attempt to approach exact satisfaction of the friction law. They however also noted that their algorithm may not converge to find the correct solution and thus cannot provide any guarantees. Trinkle et al. [51] showed that the Coulomb friction law with maximum dissipation can be cast as a mixed nonlinear complementarity problem (mixed NCP), which is difficult solve. Therefore, they propose linearizing the friction cone by approximating it as a pyramid, which allows for a formulation of the friction constraint as a linear complementarity problem (LCP). Problems of this type can be solved with Lemke’s algorithm. A downside of this approximation is that the friction force direction inside the linearized cone that maximizes energy dissipation is generally not the same one that opposes motion. Furthermore, the linearization violates the assumption of isotropic friction. An immensely influential time stepping scheme for the solution of multi-rigid-body dynamics with Coulomb friction that makes use of this framework became known as the Stewart-Trinkle formulation [45].

Another very influential time-stepping scheme - the Anitescu-Potra formulation [3] - can be obtained by omission of the constraint stabilization term from the Stewart-Trinkle formulation. Anitescu and Tasora [2, 46] showed that using the LCP formulation of the pyramidal friction cone approximation in general leads to nonconvex solution sets. LCP problems with nonconvex solution sets contain reformulated instances of the Knapsack problem and are therefore NP-hard. Thus, problems of this type are difficult to solve. The authors develop an iterative algorithm to solve them that converges to the solution of the original problem. They achieve this through successive convex relaxation effectively solving subproblems that have the form of strictly convex quadratic programs. They note that their algorithm is only guaranteed to converge for ‘sufficiently small’ friction coefficients but provide a lower bound for convergence. In further work Anitescu et al. [4]

propose a cone-complementarity approach in order to alleviate the shortcomings of the LCP approach with linearized cones. They develop an iterative method that converges under fairly general conditions but may allow bodies to behave as if they were in contact although they have drifted apart [22]. Kaufman et al. [23] showed that solving for contact forces when relative velocities are known can be achieved by solving a quadratic program (QP). Similarly, relative velocities can be solved for when contact forces are known. However, problems where both are to be solved for simultaneously are non-convex. The authors present an algorithm which iterates between the two convex QPs until a solution of the required accuracy is found. Although it works well in practice convergence is not guaranteed with their approach.

Todorov [48] takes a completely new approach by deriving nonlinear equations for the dynamic contact problem that implicitly satisfy the complementarity conditions. In order to solve the resulting nonlinear equations. Todorov proposes a Gauss-Newton approach with specific adaptations such as a novel linesearch procedure. Todorov [49] and Drumwright et al. [15] independently developed formulations that are relaxing the complementarity conditions such that a convex optimization problem is recovered. Song et al. [43] propose a compliant model of the contacts between nominally rigid bodies. Both normal and tangential contact forces are determined by viscoelastic constitutive relations coupling them to local deformations. Pang et al. [38] showed that the LCP formulation for discretized friction cones first introduced by Trinkle [51] can be cast as a Mixed Integer Program (MIP). While this method still suffers from the same inaccuracies as the LCP formulation they demonstrated its sufficiency for the control of a robotic gripper in simulation.

In previous work [18] we modeled grasps allowing for breaking contacts using MIPs. In that work we did not treat the MDP explicitly but instead introduced an iterative algorithm in order to mitigate inaccuracies due to approximation of the friction law. It also contains a direct comparison of our method with those developed by Ferrari and Canny [16] as well as Prattichizzo et al. [39].

The work presented in this paper builds on our previous work. We demonstrate that the MDP can be leveraged for accurate analysis of passive grasp stability. We propose a framework for efficient solution of the resulting problems by approximating the friction law including the MDP and successively refining the approximation to obtain solutions arbitrarily close to the exact solution.

### III. GRASP MODEL

Consider a robotic hand that makes  $m$  contacts with a grasped object. The systems is initially at rest and we would like to determine if the system will remain at rest when a disturbance is applied to the object. Each contact is defined by a location on the surface of the grasped object and a normal direction determined by the local geometry of the bodies in contact. We chose the *Point Contact with Friction* model to describe the possible contact wrenches that may arise at the interfaces between the hand and the object. Therefore we

only consider contact forces and do not allow for frictional torques. This is a reasonable assumption for contacts between smooth and relatively stiff bodies. For any contact specific vector (such as the contact force), we will use subscripts  $n$  and  $t$  respectively to denote the components lying in the contact normal and contact tangent directions. We use the vector  $\mathbf{c} \in \mathbb{R}^{3m}$  to denote contact forces, where  $\mathbf{c}_i \in \mathbb{R}^3$  is the force at the  $i$ -th contact. Using the notation above,  $c_{i,n} \in \mathbb{R}$  is the normal component of this force, and  $\mathbf{c}_{i,t} \in \mathbb{R}^2$  is its tangential (i.e. frictional) component.

**Equilibrium:** The grasp map matrix  $\mathbf{G} \in \mathbb{R}^{6 \times 3m}$  maps contact wrenches into a frame fixed to the grasped object. We can now write the equilibrium equations for the grasped object where we collect all disturbances externally applied to the object (such as gravitational forces for instance) as  $\mathbf{w} \in \mathbb{R}^6$ . For a hand with  $l$  joints the transpose of the grasp Jacobian  $\mathbf{J} \in \mathbb{R}^{3m \times l}$  relates contact forces to the torques  $\boldsymbol{\tau} \in \mathbb{R}^l$  in the hand joints required for hand equilibrium.

$$\mathbf{G}\mathbf{c} + \mathbf{w} = 0 \quad (1)$$

$$\mathbf{J}^T\mathbf{c} + \boldsymbol{\tau} = 0 \quad (2)$$

These equilibrium equations can predict the resultant wrench on the object and the joint torques required to balance given contact forces, but, since neither  $\mathbf{G}$  nor  $\mathbf{J}^T$  are typically invertible, they have no predictive capabilities in the opposite directions: we can not use them to predict the contact response to known external wrenches, or joint torques. They thus fail to capture effects where contact forces are transmitted (either amongst the joints or between the joints and the external environment) through the object itself, and cannot account for many of the phenomena illustrated in Section I.

**Virtual motion:** To resolve the structural indeterminacy of the grasp and determine which contact forces will arise in response to a disturbance, we must introduce additional constitutive relations. Following the grasp compliance model introduced by Cutkosky and Kao [14] and later used by Bicchi [8, 9, 10], we model contact stiffness by introducing virtual springs at contacts. Normal forces must arise as a result of these springs being loaded through *virtual motion* of the grasped object. However, we significantly extend these models, by including constraints for unilateral contacts, and a more accurate friction model.

The constitutive relations we use relate contact forces to the relative *virtual motion between the hand and the object at the contacts*. This relative motion, expressed in the contact frames, is denoted by  $\mathbf{d} \in \mathbb{R}^{3m}$  (only the translational components are of interest to us), and can be formulated in terms of the overall virtual object motion  $\mathbf{r} \in \mathbb{R}^6$  and the virtual joint motion  $\mathbf{q} \in \mathbb{R}^l$ :

$$\mathbf{G}^T\mathbf{r} - \mathbf{J}\mathbf{q} = \mathbf{d} \quad (3)$$

The introduction of motion in our framework might initially seem to contradict our focus on the static equilibrium of a grasp. However, this allows us to resolve the structural indeterminacy of the grasp through the use of constitutive relations. It furthermore allows us to enforce nonsmooth constraints such as the unilaterality of the contacts and the non-backdrivability

of the joints. One can think of our approach as solving the first step of the dynamics of the hand-object systems in order to determine if it is in equilibrium. Thus, we consider all motion to be 'virtual', used as a tool to enforce results that are consistent with rigid body behaviors: *all contact forces must be consistent with some virtual motion of the grasped object and the joints, which become additional variables in our framework.*

**Normal forces:** The constitutive relation for normal forces assumes virtual springs of stiffness  $k$  along the contact normals. Thus, the normal force at a contact  $i$  is determined by the relative motion between the object and the robot hand at that contact in the direction of the contact normal. However, we also model unilateral contacts, which may only push on an object but can never pull (we do not concern ourselves with adhesion or similar effects). Thus, the normal force at a contact must be strictly non-negative. Furthermore, if the contact detaches the contact force must be zero. For simplicity and without loss of generality we can assume  $k = 1$ .

$$\begin{cases} c_{i,n} = -d_{i,n} & \text{if } d_{i,n} \leq 0 \\ c_{i,n} = 0 & \text{if } d_{i,n} > 0 \end{cases} \quad (4)$$

**Friction forces:** To model friction, we chose the Coulomb model. The first part of the Coulomb model provides an upper bound to the magnitude of the friction force given the normal force and the friction coefficient  $\mu_i$  (as all motion is virtual we consider all friction coefficients to be those of static friction.) This defines a cone  $\mathcal{F}_i$  at each contact (see Fig. 2a.)

$$\mathcal{F}_i(\mu_i, c_{i,n}) = \{\mathbf{c}_{i,t} : \|\mathbf{c}_{i,t}\| \leq \mu_i c_{i,n}\} \quad (5)$$

The second part of the Coulomb model concerns the Maximum Dissipation Principle (MDP) [44]: Given a relative contact motion, the friction force at that contact must maximize energy dissipation, while bounded by eq. (5).

$$\mathbf{c}_{i,t} \in \arg \min_{\mathbf{c}_{i,t} \in \mathcal{F}_i} \mathbf{c}_{i,t}^T \cdot \mathbf{d}_{i,t} \quad (6)$$

In the case of isotropic friction, the dissipation is maximized if the friction force is anti-parallel to the relative sliding motion and lies on the boundary of the cone  $\mathcal{F}_i$ . Thus, we can also directly express the friction force in terms of the normal force and the relative sliding motion. We must distinguish between two cases:

- At a contact that does not exhibit relative motion in a tangential direction (sliding) the friction force is constrained such that the contact force lies within the cone  $\mathcal{F}_i$
- If a contact does exhibit sliding, the friction force must oppose the direction of motion (or incipient acceleration), and the total contact force must lie on the friction cone edge.

The complete model can be formulated as follows:

$$\begin{cases} \|\mathbf{c}_{i,t}\| \leq \mu_i c_{i,n}, & \text{if } \|\mathbf{d}_{i,t}\| = 0 \\ \mathbf{c}_{i,t} = -\mu_i c_{i,n} \frac{\mathbf{d}_{i,t}}{\|\mathbf{d}_{i,t}\|}, & \text{otherwise} \end{cases} \quad (7)$$

Note that the formulation in eq. (7) requires the distinction between sliding contacts and those remaining at rest while the

original formulation in eqs. (5)&(6) holds in both cases. For isotropic friction they are equivalent, and we have thus arrived at an accurate constitutive relation describing the friction forces in terms of the relative contact motion.

**Joint torques:** Finally, we model the joints as non-backdrivable, in order to capture the behavior of the majority robotic hands driven by highly geared motors. This means that a joint  $j$  may only exhibit virtual motion in the direction that its commanded torque  $\tau_{j,c}$  is moving it in. The joint torque may exceed the commanded level, but only if this arises passively. This means that a joint that is being passively loaded beyond the commanded torque levels must be locked in place and may not move. A moving joint must apply the torque it was commanded to. Thus, we must also distinguish between two states for each joint.

$$q_j \geq 0, \quad \begin{cases} \tau_j \geq \tau_{j,c}, & \text{if } q_j = 0 \\ \tau_j = \tau_{j,c}, & \text{if } q_j > 0 \end{cases} \quad (8)$$

A joint with zero commanded torque may not move, as any torque arising from external factors will be absorbed by the gearing.

**Complete problem:** The system comprising eqs. (1)-(8) defines the static equilibrium formulation for a grasp. It is very general in nature, and can be considered as part of existence problems (e.g. given  $\tau$ , determine if  $r$  and  $c$  exist that balance a given  $w$ ), or optimization problems, with the addition of an objective (e.g. determine the optimal  $\tau$  that satisfies the existence problem above). Remaining agnostic to the exact query that is being solved we will refer to the *exact problem* as the following query: given a subset of  $\tau$ ,  $r$ ,  $c$  or  $w$ , determine the rest of these variables such that (1)-(8) are exactly satisfied.

The main difficulty of directly solving this exact problem lies in the non-convexity of the friction law when including the MDP - the constraints in (7) contain a nonlinear equality. Many of the works we have reviewed earlier seek to find approximations that lend themselves to efficient solvers. The compliance grasp model [14] makes a linear approximation, while Prattichizzo et al. [39] also make a distinction between sticking and sliding contacts. However, due to the non-convexity of the second part of (7) they must make the overly conservative approximation that a sliding contact may not apply any frictional forces at all.

Of course one could cast the MDP in its original minimization formulation instead: If we remove (7) from the exact problem and introduce instead a minimization objective (6) then an optimal solution to this new problem would be valid if it also satisfies (7). However, this formulation is of course still non-convex as the bilinear form in (6) is not positive definite [26]. However, this reformulation of the exact problem would allow for a global optimization approach based on a convex relaxation of the bilinear forms using McCormick underestimators [30], which relies on upper and lower bounds on the variables involved in the bilinear forms. We could then use a Spatial Branch-and-Bound (sBB) algorithm [52] to solve for the global optimum (see for example the  $\alpha$ BB algorithm [1] or the BARON [47] software for global optimization.)

In this paper we make use of the particular structure of the grasping problem to derive a relaxation that does not require variable bounds, since in the grasping problem bounds on contact forces and particularly virtual motions are not implicit to the problem and thus difficult to predetermine. Furthermore, we relax the constraints of the optimization problem instead of the objective. This allows us to use the objective to solve for interesting grasp characteristics such as the optimum actuator commands.

The friction models used in the grasp force optimization literature [42, 24, 33, 13, 20, 11] assume full control over the contact forces and therefore only have to concern themselves with stationary contacts. This removes the non-convexity that arises from the MDP. As all motion in our problems is virtual, one may indeed be tempted to ignore the sliding effects and use one of the friction models in the grasp force optimization literature. We can achieve this by solving the problem described by equations (1)-(5), (8) and neglecting equation (6). This approach, however, fails to capture the passive effects in the grasps we want to analyze. As we will show in Section VII solving (1)-(5), (8) without the MDP results in unphysical solutions that violate the laws of energy conservation. In contrast, as we will also show in Section VII, including the MDP in our formulation allows us to obtain physically meaningful results.

The approaches from the rigid body dynamics literature are perhaps more applicable to these problems, as they concern themselves with moving bodies and therefore must include some treatment of the MDP. However, due to the difficulty in solving the resulting problems [35, 2, 23] they are either computationally infeasible [50, 37] or make approximating assumptions [51, 45, 3, 4, 48, 49, 15, 43, 38] in order to allow for real-time simulation. Others make no convergence guarantees [5, 6, 18] or are only guaranteed to converge for small friction coefficients [46].

*To the best of our knowledge, no efficient solution has been proposed to date for a three-dimensional grasp model that includes an exact formulation of Coulomb friction such as eq. (7), or an equivalent reformulation. Our main contribution is a method to efficiently find solutions to progressively tighter approximations of this model, up to arbitrary accuracy.*

#### IV. FORMULATION AS MIXED INTEGER PROGRAM

Let us take a closer look at the grasp model constraints introduced so far. Constraints (1)-(3) are simple equality constraints. We note that the constraints in (4) and (8) exhibit a combinatorial nature - a type of constraint that can be found in mixed integer programming. In fact both the normal force and the joint model relationships can be cast as pairs of convex constraints with binary decision variables in an MIP that can then be solved using algorithms such as branch and bound. The friction law is more complex: the constraint for sliding contacts in (7) defines a non-convex set. Therefore, we must formulate a relaxation of this constraint such that we can solve the system as an MIP.

We start from the common linearized friction model which replaces the circular friction cone at contact  $i$  with its discretization as a polygonal cone (see Fig. 2.) Matrix  $D_i \in \mathbb{R}^{3 \times k}$

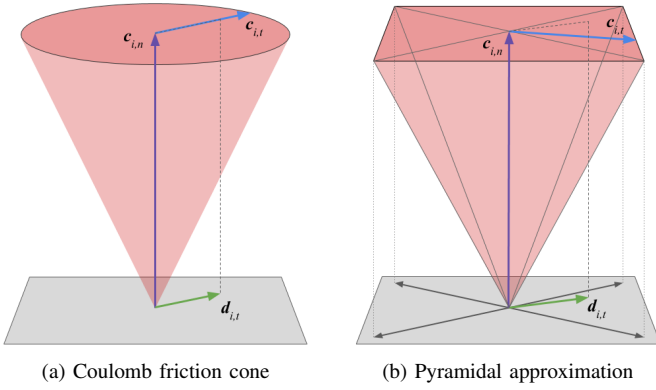


Fig. 2: Illustration of the exact Coulomb friction cone and a pyramidal approximation as we use it. The black arrows positively span the contact tangential plane and make up the matrix of basis vectors  $D_i$ .

contains as its columns a set of  $k$  vectors that positively span the contact tangential plane and thus the space of possible friction forces. Frictional forces can now be expressed as a positive linear combinations of these so called *friction edges* with weights  $\beta_i \in \mathbb{R}^k$ . Inequality constraints on a vector are to be understood in a piecewise fashion.

$$D_i \beta_i = c_{i,t}, \beta_i \geq 0 \quad (9)$$

We can express relative tangential contact motion as a weighted combination of the same friction edges, with weights  $\alpha_i$ . For reasons that will soon become apparent, we choose to express the negative of the tangential motion instead:

$$D_i \alpha_i = -d_{i,t}, \alpha_i \geq 0 \quad (10)$$

If the friction edges in  $D_i$  are arranged in an ordered fashion such that neighboring friction edges in the tangential plane are also neighbors in  $D_i$  we can constrain friction to (approximately) oppose motion by requiring that *the friction force lies in the same sector of the linearized friction cone as the negative of the tangential contact motion*. Without loss of the above properties, we require that at most two components of  $\beta_i$  can be non-zero and that non-zero components are either consecutive or lie at the first and last positions of vector  $\beta_i$ . This can be achieved by constraining  $\beta_i$  with a special ordered set  $z_i \in \mathbb{R}_{\geq 0}^{k+1}$  of type 2 (SOS2) [7], which has one more component than  $\beta_i$  itself.

$$\beta_{i,1} \leq z_{i,1} + z_{i,k+1}, \beta_{i,2} \leq z_{i,2}, \dots, \beta_{i,k} \leq z_{i,k}, z \in \text{SOS2} \quad (11)$$

A special ordered set of type 2 is a set of ordered non-negative numbers of which at most two can be non-zero. If two numbers in the SOS2 are non-negative they must be consecutive in their ordering. This type of constraint can be solved by MIP solvers and is hence admissible to our model. We now similarly constrain the weights  $\alpha_i$  that determine relative motion with the same SOS2 as in (11).

$$\alpha_{i,1} \leq z_{i,1} + z_{i,k+1}, \alpha_{i,2} \leq z_{i,2}, \dots, \alpha_{i,k} \leq z_{i,k} \quad (12)$$

Note that these constraints hold for both sliding as well as stationary contacts. If a contact is stationary then all components of  $\alpha_i$  must be zero. Thus, any two consecutive components of  $z_i$  may be non-zero and since  $\beta_i$  are the weights of the basis vectors in  $D_i$  the friction force may point in any direction in the tangent plane. If the contact slides then some components of  $\alpha_i$  must be nonzero. As both  $\alpha_i$  and  $\beta_i$  are constrained by  $z_i$  and only two consecutive components of  $z_i$  may be nonzero, this means the friction force must lie in the same sector of the friction pyramid as the negative of the tangential relative contact motion, but may not be collinear (see Fig. 2b.)

Finally, we can constrain the magnitude of the friction force in addition to its direction. For sliding contacts, friction must be maximized, while for stationary contacts it only has an upper bound. The friction edges in  $D_i$  are chosen to be unit vectors such that these constraints can be expressed as follows:

$$\begin{cases} e^T \beta_i \leq \mu_i c_{i,n}, & \text{if } \|d_{i,t}\| = 0 \\ e^T \beta_i = \mu^i c_{i,n}, & \text{otherwise} \end{cases} \quad (13)$$

where  $e = (1, 1, \dots, 1) \in \mathbb{R}^k$ . Constraint (13) can also be included in an MIP using a binary decision variable.

We now have a complete model of friction. For a finite value of  $k$ , this model is approximate; in the limit, as  $k \rightarrow \infty$  eq. (13) in combination with constraints (11 - 12) behave like the Coulomb friction model in (7).

## V. SUCCESSIVE HIERARCHICAL REFINEMENT

We can solve the complete system described by constraints (1)-(4), (8)-(13) as an MIP with algorithms such as branch and bound. In order to improve our approximation, we could choose a high number of edges for the discretized friction cones. In practice, however, that approach is not feasible as the time taken to solve an MIP is sensitive to the number of integer variables in the problem. As SOS2 constraints are implemented using binary variables a highly refined friction cone approximation quickly becomes computationally intractable.

### A. Intuition

Our approach is based on the key insight that one can obtain an equally accurate solution by solving a problem with a coarse friction cone approximation, and *successively refining the linearized friction constraints only in the region where friction forces arise*. Our approach thus proceeds as follows:

- We solve our problem using a coarse approximation of the friction cone (few friction edges). From the solution, we identify the sector of the linearized cone (the area between two edges) where both the friction force and the negative of the relative motion lie.
- To obtain a tighter bound, we add new friction edges that refine *only the sector identified above*. We then repeat the procedure with the new, selectively refined version of the friction cone.

An important characteristic of our method is that we can choose our friction edges so that, at any level of refinement, *the solution set to the approximate problem contains the solution*

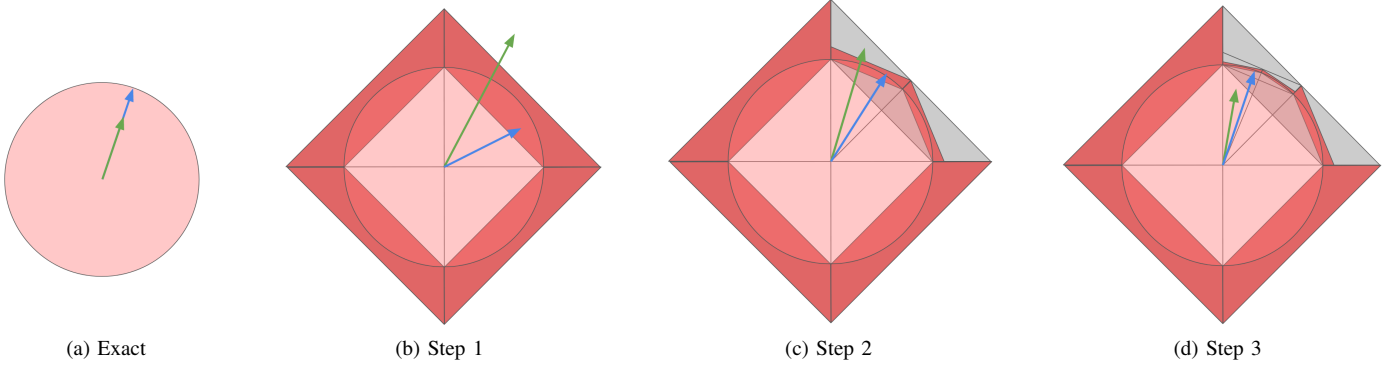


Fig. 3: First three steps of our algorithm refining the friction law relaxation successively and locally. The dark red regions are the space of feasible friction forces at a sliding contact. If there is no tangential motion the light red regions are added to the feasible friction force space. The green and blue arrows are  $\mathbf{d}_{i,t} = \mathbf{D}_i \boldsymbol{\alpha}_i$  and  $\mathbf{c}_{i,t} = \mathbf{D}_i \boldsymbol{\beta}_i$  respectively, which drive refinement of their local friction cone sector.

to the exact problem, assuming one exists. This concept is a necessary condition for a tightening approach and is visually illustrated in Fig. 3. For the exact problem, the feasible space of friction forces for sticking contacts is the inside of the circle, while the feasible space of friction forces for sliding contacts is the circle itself - a non-convex set. Consider now a rough approximation with four friction edges (Step 1). If we allow sticking friction to reside inside the areas shaded in either shade of red, while sliding friction must lie within the dark red border, the space of allowable solutions to the exact problem is contained inside our linear and piecewise convex approximation. Assume that, at this level of refinement, there is a solution to an equilibrium problem, with friction force lying inside the upper right sector. We refine this sector, again taking care that the space of allowable solutions to the exact problem is contained inside our refinement (Step 2). We continue this procedure (Step 3, etc.) until one of two things happen: we either reach a level of refinement where no solution exists, or we refine down to the point where the active sector is as small as we want it to be, bringing us arbitrarily close to the solution to the exact problem.

This refinement scheme provides two important advantages: If, at any point during the refinement, no solution exists that satisfies equilibrium, we can guarantee that no solution can exist to the exact version of the problem either. This guarantee immediately follows from the properties that the solution set at any refinement level includes the solution set at the next level, and that, in the limit, our discretization approaches the exact constraints. (Note that alternative discretizations of the friction constraints, such as the LCP formulations discussed previously, do not exhibit this property.) In practice, this means that, when no solution exists to the exact equilibrium problem, our algorithm can determine that very quickly, only solving relatively coarse refinement levels.

The second advantage our scheme provides is that when a solution does exist, we can typically refine it to high accuracy (a very close approximation to the solution of the exact problem) using relatively few friction edges. This is not theoretically guaranteed: in the worst case, our approach

could require all sectors to be fully refined before finding an adequate solution as the desired resolution, and may hence perform worse than using a fully refined friction discretization to begin with. However, we have never found that to be the case. Typically, only a small region of the discretization must be refined as the contact forces are also constrained by equilibrium relations (1) & (2) and will generally point in similar directions at all levels of refinement, leading to a very localized and targeted tightening of the relaxation. Thus, this algorithm is efficient enough to analyze complex grasps on a consumer PC to levels of refinement that are otherwise unachievable.

However, in order to achieve these abilities, we have to implement a friction cone refinement method meeting the requirement discussed above. We present the details of our chosen refinement scheme next.

### B. Implementation

Let us pick the initial basis vectors in  $\mathbf{D}$  such that the angle  $\gamma$  between all pairs of successive vectors is equal. We pick an initial angle  $\gamma = \pi/2$ . We refine our polyhedral friction cones by bisecting sectors defined by the non-zero components of  $\mathbf{z}$  and define the angle at which to stop refinement as  $\gamma/2^q$ . We now find the required length  $l_1$  of these initial friction edges such that the initial solution set contains the solution sets at all refinement levels.

$$l_1 = \prod_{r=1}^{q+1} \sec\left(\frac{\gamma}{2^r}\right) \quad (14)$$

Thus, we modify the friction edges in  $\mathbf{D}$  accordingly. We define vector  $\mathbf{f}_i$  to contain the lengths of the friction edges making up the friction cone approximation at a contact in an order corresponding to the order of the weights  $\boldsymbol{\beta}_i$ . Constraints (13) now become

$$\begin{cases} \mathbf{e}^T \boldsymbol{\beta}_i \leq \mu_i c_{i,n}, & \text{if } \|\mathbf{d}_{i,t}\| = 0 \\ \mathbf{e}^T \boldsymbol{\beta}_i \leq \mu_i c_{i,n}, \mathbf{f}_i^T \boldsymbol{\beta}_i \geq \mu_i c_{i,n}, & \text{otherwise} \end{cases} \quad (15)$$

We are now ready to solve the initial coarse relaxation problem defined by (1)-(4), (8)-(12) and (15). We find the two

---

**Algorithm 1** Grasp analysis through successive relaxation
 

---

**procedure** RELAXATION REFINEMENT

**Input:**
 $O$  - objective function

 $C$  - additional constraints

 $\gamma$  - initial refinement level (angle btw. friction edges)

 $q$  - max refinement level desired

 Initialize  $D$  with basis vectors of length  $l_1$  as in (14)

**do**

 Optimize  $O$  subject to (1)-(4), (8)-(12), (15) and  $C$ .

**if** no solution exists **then**

   **return** no feasible solution

**end if**

 refinement\_needed  $\leftarrow$  False

**for** each contact  $i$  **do**

   Find active edges in  $D_i$ 

    $\delta_i \leftarrow$  angle between active edges

   **if**  $\delta_i > \gamma/2^q$  **then**

      $p \leftarrow \log_2(\gamma/\delta_i) + 2$ 

     Add new edges to  $D_i$  of length  $l_p$  as in (16)

     Remove redundant edges from  $D_i$ 

     refinement\_needed  $\leftarrow$  True

   **end if**
**end for**
**while** refinement\_needed

**return** solution

**end procedure**


---

active friction edges  $d_1$  and  $d_2$  and create three new edges that point in the direction of  $d_1$ ,  $d_1 + d_2$  and  $d_2$  and have magnitude  $l_2$ . For this and all following refinements we have

$$l_p = \prod_{r=p}^{q+1} \sec\left(\frac{\gamma}{2^r}\right) \quad (16)$$

where  $p$  is the level of refinement of the sectors to be created. We insert the new friction edges between  $d_1$  and  $d_2$  in matrix  $D_i$  and remove any redundant friction edges (edges that are identical or edges that lie between any such edges). Solving the problem thus obtained we can continue refining the friction discretization until the angle of the active sectors at all contacts reach an angle of  $\gamma/2^q$ . We do not further refine any sector that has already reached this threshold. The overall method is shown in Algorithm 1.

Note that this convex relaxation does not rely on explicit bounds on the problem variables as would be the case were we using McCormick envelopes [30]. Instead, we use the bounds implicit in the friction cone constraint (7) in our relaxation. This is also a further advantage of our approach as the relaxation has an easily understood physical interpretation.

## VI. IMPROVING ROBUSTNESS TO GEOMETRICAL UNCERTAINTIES

In the above we outlined a grasp model that allows us to analyze the grasp stability given perfect information about the geometry of the grasp. We assume we know exactly the contact position and orientation. In practice however we

often encounter uncertainties, which can greatly affect the stability of a grasp. Even when using tactile sensors in order to locate contacts made between the hand and the object the contact normals (and hence orientation) are often difficult to obtain accurately. Therefore we would like to make our framework robust to discrepancies up to a certain magnitude. We introduce the method we use for this here, and illustrate its importance in the following section.

Let us suppose we have an upper bound on our uncertainty in the contact normal  $\eta$ . Thus, the actual contact normal lies in a space of possible contact normals that deviate by at most angle  $\eta$  from the nominal contact normal. In order for a grasp to be robust to deviations defined by this space we would like it to be robust in the worst-case. The worst-case contact normal is one such that its projection into the tangent plane opposes the relative tangential contact motion. In our space of contact normals such a normal would be the most effective at unloading the contact and hence destabilizing the grasp. The relative contact motion in direction of the worst-case normal would then be given by

$$\hat{d}_{i,n} = d_{i,n} \cos(\eta) - \|d_{i,t}\| \sin(\eta) \quad (17)$$

However, we need to find a linear approximation for  $\|d_{i,t}^T\|$  as including it exactly would introduce a nonconvex quadratic equality constraint. Fortunately we can use the amplitudes of the friction edges  $\alpha$ . Using (10) & (16) the summation of the product of all contact motion amplitudes with the length of the corresponding friction edges gives us an estimate of the magnitude of the relative tangential contact motion.

$$\|d_{i,t}\| \approx \sum_{s=1}^k l_p \alpha_{i,s} \quad (18)$$

where  $p$  is the level of refinement of the active sector (i.e. the sector corresponding to the nonzero components of  $\alpha_i$ .) This formulation is equivalent to one where all the friction edge vectors in  $D$  are of unit length where we could then omit  $l_p$ .

The problem with this formulation is that, at low resolutions of the friction cone, it overestimates the relative tangential contact motion. It is exact for contact motion parallel to to any friction edge (tangential motion of magnitude 1 will result in exactly one component of  $\alpha$  equal to 1). However, for any unit tangential motion that lies between two edges, the sum of the two active components of  $\alpha$  must be greater than 1 due to the triangle inequality. This effect diminishes at finer resolutions as friction edges become closer to parallel, but the destabilizing effect is potentially larger at coarser resolution.

However, recall that our refinement method requires that the solution set at coarser levels includes the solution set at more detailed levels. We thus require the destabilizing effect to be weaker at coarse resolutions and become stronger approaching its exact value as  $k \rightarrow \infty$ . Therefore we modify (18) such that it underestimates tangential motion, except at the midpoint between two edges where it is exact.

$$\|d_{i,t}\| \approx \frac{l_p}{l_{p+1}} \sum_{s=1}^k l_p \alpha_{i,s} \quad (19)$$




Fig. 4: Illustration of the shortcomings of neglecting the MDP when solving passive grasp stability problems. A force is applied to the object in the positive  $y$ -direction. The solver rotates the object around the  $y$ -axis (violet) generating normal forces (blue) due to the slight offset of the distal contacts. These normal forces are so large the frictional force component is barely noticeable. In this fashion a formulation without the MDP allows the solver to resist arbitrary disturbances through unphysical object motions.

As  $k \rightarrow \infty$  this estimation becomes exact. We now replace the normal relative contact motion in (4) with  $\hat{d}_{i,n}$  in order to obtain solutions that are robust to uncertainties in contact normal up to an angular discrepancy of  $\eta$ .

## VII. RESULTS

### A. Determining stability and computing contact forces

The simplest query we may make is to determine the stability of a grasp with respect to a given disturbance. We will use this application of our method to illustrate the importance of the friction model we introduce compared to previous literature, and extend to more complex queries in the following subsections.

Let us solve for the stability of the grasp shown in Fig. 1 (and again in Fig. 4) when we apply a force of 1N to the object in the  $y$ -direction *without preloading the joints at all*. To answer this query, we use Algorithm 1 as follows:

$$\text{objective:} \quad \text{none} \quad (20)$$

$$\text{additional constraints:} \quad \mathbf{w} = [0, 1, 0, 0, 0, 0]^T \quad (21)$$

$$\boldsymbol{\tau}_c = 0 \quad (22)$$

Note that, in the absence of an optimization objective, we are simply asking if a solution exists that satisfies all the constraints of the problem, equivalent to determining values for all the unknowns (contact forces, virtual motions, etc.) such that the grasp is stabilized. Using Gurobi [17] as a solver for the constituent MIPs, Algorithm 1 finds no feasible solution to this problem (i.e. it predicts the grasp is unstable in the presence of the given disturbance.) This is the expected result: the grasp may not resist a force of 1N in the  $y$ -direction without any preloading, as it is intuitively clear that the object will slide out.

What happens if, instead of our formulation incorporating the MDP, we use the friction models commonly used in the grasp force optimization literature [42, 24, 33, 13, 20, 11]? Together with equations (1)-(4) and (8) (which model unilateral contacts and nonbackdriveable joints) the complete query can then be formulated as a single MIP. Again using Gurobi as a solver, we find that the resulting MIP accepts a solution *regardless of the magnitude of the applied disturbance*. This is equivalent to saying that the grasp is stable in such cases even in the absence of any preload, which is clearly incorrect.

Why does a simpler friction model accept solutions that we know are physically unrealistic? To obtain further insight, we can study the virtual object motion returned by the solver in such a case: this motion corresponds to a rotation around the  $y$ -axis (along which also the disturbance is applied), “wedging” the object between the rigid, passively loaded fingers (Fig. 4a). The solver exploits a slight offset between contacts (Fig. 4b) to generate very large normal forces, thus permitting friction forces to balance the applied disturbance, regardless of its magnitude. However, the equilibrium contact forces computed in this fashion *do not oppose the virtual motion, and thus break the laws of energy conservation*. In order to obtain physically meaningful results it is necessary to include the MDP in the friction constraints.

In contrast, our formulation correctly captures the interplay between applied preload, and the ability to resist disturbances. Let us now consider the case where we apply a preload torque of 0.1 Nm at the proximal joints. Due to the simplicity of the grasp we can analytically determine the expected maximum force in the  $y$ -direction the grasp can withstand: a  $\sim 90$ mm contact moment arm and a friction coefficient of 1.0 results in a maximum total friction force applied to the object of  $\sim 2.2$  N. Using a similar formulation as before (no objective,  $\boldsymbol{\tau}_c = 0.1$ Nm at the proximal joints), we indeed find that Algorithm 1 accepts a solution (predicts stability, see Fig. 5) for a disturbance of 2.2N in the  $y$ -direction ( $\mathbf{w} = [0, 2.2, 0, 0, 0, 0]^T$ ), but finds no solution for a disturbance of 2.5N in the same direction ( $\mathbf{w} = [0, 2.5, 0, 0, 0, 0]^T$ ).

### B. Maximum resistance to disturbances in a given direction

So far we have used Algorithm 1 simply to check feasibility of a problem, without an optimization objective, and used it to “spot check” resistance to specific disturbances. Adding an objective allows us to formulate more powerful queries. For example, we can directly determine the exact *maximum disturbance* applied to the grasped object in a given direction that a grasp may resist purely passively.






Fig. 5: Equilibrium contact forces that are predicted to arise by our framework when a force of 2.2N is applied to the object in the  $y$ -direction. The proximal joints have both been preloaded with 0.1Nm. Note the refinement of the pyramidal friction cone approximation. Only a minority of sectors have been refined. Sectors that are not of interest in this specific grasp problem remain in a less refined state and thus only contribute little to the complexity of the overall problem. The sectors containing the equilibrium contact forces are  $< 1^\circ$  small.



(a) No contact normal uncertainty



(b) Robust to contact normal uncertainties of  $2.5^\circ$

Fig. 6: Resistible forces in the grasp plane of the grasp in Fig. 1 as predicted by our model and algorithm. In blue are forces that can be resisted even without the application of preloading torques at the joints. When loading the two proximal joints with 0.1Nm the maroon area is added to the resistible forces.

To achieve this, we prescribe a preload  $\tau_m$  for the actuators to be kept constant, and a direction  $\mathbf{d}$  along which to apply a disturbance to the object. To compute the largest magnitude

disturbance the grasp can withstand in that direction, we use Algorithm 1 as follows:

$$\text{objective: maximize } s \quad (23)$$

$$\text{additional constraints: } \mathbf{w} = s\mathbf{d} \quad (24)$$

$$\tau_c = \tau_m \quad (25)$$

To illustrate this, we will again use as an example the grasp in Fig. 1. The simplicity of this grasp allows us to verify the accuracy of predictions made by our algorithm by comparison to the intuitions formulated in Section I. We discretize the space of possible disturbance directions lying in the  $xy$ -plane with a  $1^\circ$  resolution. We then compute the maximum magnitude wrench that can be resisted along each of these discrete directions. We visualize the results in Fig. 6, where Fig. 6a plots results without considering robustness to contact normal uncertainty, while Fig. 6b assumes an uncertainty of  $2.5^\circ$ .

The results match our intuition that any downward force can be reacted without any loading of the fingers. Furthermore the model captures the need for finger loading in order to resist upward forces. It also shows an effect of passive finger loading for forces with nonzero  $x$  component: pushing sideways increases the amount of resistance to upwards forces.

The reason for the asymmetry of Fig. 6a however is not immediately obvious, as the grasp itself appears symmetric. In fact however, the two distal contacts are ever so slightly offset, causing the object to wedge itself stuck if enough leftward force is applied. The grasp in Fig. 7 makes it clearer why this behavior occurs - here the contacts are visibly offset. Note, that this wedging behavior is very different from what we observed when solving without the MDP (see Fig. 4). There, the wedging occurred no matter the applied wrench such that arbitrary wrenches could be resisted. Furthermore, the resulting contact forces did not satisfy energy conservation.

In contrast, in our framework only specific wrenches allow wedging to occur. These wrenches depend on the geometry of the grasp and are consistent with the rigid body statics of the grasp problem. The equilibrium contact forces predicted by our framework satisfy the MDP and hence energy conservation. However, as our method allows us to solve the rigid body problem very accurately, only a small offset is required for our model to predict wedging of the object; an offset that is easily within the accuracy of a typical triangular mesh. Note, that this behavior is fully consistent with the rigid body assumption and the predictions made by our framework are correct, albeit highly sensitive to the grasp geometry.

Of course in practice it is not advisable to rely on such volatile geometric effects. Therefore taking into account geometric uncertainties is of paramount importance for practical applications. Fig. 6b shows which forces can be robustly resisted when we consider the uncertainty in normal angle to be no larger than  $2.5^\circ$  (using the approach outlined in Section VI.) The resulting plot of resistible forces is approximately symmetric corresponding to the near-symmetry of the grasp. The indicated spaces of resistible forces both with and without a preload are consistent with our intuition. We note that the range of passively resistible forces in the  $xy$ -plane




Fig. 7: A grasp designed to highlight the 'wedging' effect. The two contacts on the distal link are offset with respect to each other allowing wedging to occur.




Fig. 8: Forces in the  $xy$ -plane predicted resistible by the grasp in Fig. 7 using our model and algorithm. In blue are forces that can be resisted even without the application of preloading torques at the joints. When loading the two proximal joints with  $0.1\text{Nm}$  the maroon area is added to the resistible forces. We also take into account a contact normal uncertainty of  $2.5$  degrees to make sure the grasp is robust to such discrepancies.

for this grasp was also experimentally validated in previous work [18], and the results shown in Fig. 6b match this existing data.

We showed how we can make the stability predictions less sensitive to the small scale geometric characteristics of the grasp and thus robust to uncertainties. At a larger scale, however, wedging effects can be robustly leveraged. If the contacts are offset such as in Fig. 7 the forces this grasp may resist robustly are shown in Fig. 8. Thus, if we know the range of disturbances likely to be encountered during a manipulation task our framework can be a valuable tool in picking an appropriate grasp.

One feature of Fig. 8 perhaps requires further elaboration: The sectors in the second and fourth quadrants of Fig. 8 where forces may only be resisted when a preload is applied but not otherwise. These sectors stand out because they appear thin and are surrounded by large areas where applied forces are resistible even without a preload. We thus investigated what effects cause these wedges in order to verify if these predictions are physically accurate:

The grasps shown in Fig. 1 and 7 were created such that all contacts are as close as possible to lying in a mutual plane. This is because three dimensional grasps can be very complex




Fig. 9: 3 dimensional grasp that highlights the necessity of a grasp to be able to withstand a range of forces applied to the object. During a pouring task gravity (pink) moves in the  $xy$ -plane.




Fig. 10: Forces in the  $xy$ -plane predicted resistible by the grasp in Fig. 9. In blue are forces that can be resisted even without preloading joint torques. When loading the proximal joints with  $0.1\text{Nm}$ , the maroon area is added to the resistible forces. We use a contact normal uncertainty of  $2.5$  degrees.

while two dimensional grasps allow us to use our intuition to validate the predictions made by our framework. A limitation of using GraspIt! in our grasp analysis is that it is difficult to create grasps that are truly two dimensional in nature. Due to the meshing of the finger and object geometries as well as the intricacies of collision checking in GraspIt! the contacts always lie somewhat offset from the central plane.

In the specific case of the grasp in Fig. 7 the contacts do not quite lie within the  $xz$ -plane, in which the forces applied to the object lie. This means that two contacts are generally not sufficient in order to balance an applied force. Let us investigate the wedge in the second quadrant of Fig. 8: When a force is applied along the  $(-1, 0)$  direction the distal contact of the left finger acts as a fulcrum and the object rotates clockwise loading the contact on the proximal contact of the right finger. When a force is applied in the  $(-1, 1)$  direction the left finger contact again acts like a fulcrum, however now the object rotates counter-clockwise loading the distal contact on the right finger.

Somewhere in between those two cases the applied force points almost directly at the fulcrum contact and instead of rotating the object is mostly pressed against the left finger

breaking both contacts on the right finger. Thus, only two contacts remain and the grasp becomes unstable. In this particular case the object would rotate out of the grasp around the  $y$ -axis as the two remaining contacts on the left finger do not lie in the same plane as the applied force.

Let us now consider the grasp in Fig. 9. Note that this grasp comprises four contacts (one on each distal link plus one on a proximal link) which do not lie on the same plane, and thus has to be analyzed in a three-dimensional framework. We consider here an apparent task the robot grasping the flask may need to execute. In order to pour a liquid contained in the flask it is necessary to tip it. If we choose to use the robot wrist for this tipping motion then the force of gravity acting on the flask and its contents lies in the  $xy$ -plane. The grasp must thus be able to resist such forces in order to complete its task successfully. Furthermore, we have a choice of direction in which to turn the flask in order to pour its content. Creating a visualization (shown in Fig. 10) as before we can deduce the need for a preload, and that it is more robust to turn the flask counter-clockwise. Thus, once a grasp has been established our framework can help in making decisions as to how a task is to be executed.

### C. Computing optimal actuator commands

The passive stability of a grasp is not only determined by its geometry: the actuator commands are equally important. Consider for example the grasp in Fig. 11. The three contacts the hand makes with the object all lie approximately in the  $xz$ -plane. Contacts 1 and 2 lie approximately on the  $x$ -axis and oppose each other. Let us assume we create a grasp by commanding the proximal joints of fingers 1 and 2 to each apply 0.1Nm. Let us vary the torque commanded at the proximal joint of finger 3 and observe the difference in passive stability. Specifically, we will use our framework to investigate the maximum disturbance on the object the grasp can resist in two directions. Fig. 12 shows the resulting predictions from our algorithm.

First we are interested in forces along the positive  $z$ -axis. As expected, the resistance is largest if no motor torque is applied by finger 3. Any load by this finger only adds to the disturbance and does not help in resisting it. When the torque applied by finger 3 reaches 0.09 Nm, it has completely removed any resistance to  $z$ -direction forces. This can be easily verified: The coefficient of friction chosen for this example is 0.45 and the moment arms from joint to contact are identical for all three fingers. As there cannot be any out-of-plane forces, the normal forces at all contacts will be proportional to the applied joint torque. Thus, applying 0.9Nm at finger 3 claims all possible contact friction at both contacts 1 and 2: no further forces in that direction can be resisted.

Let us now consider passive resistance to torques applied to the object around the  $x$ -axis. If we do not load finger 3 the object is only held by contacts 1 and 2. As both these contacts lie on the  $x$ -axis they cannot apply any torque to the object in that direction. Thus, the grasp cannot resist any torques around the  $x$ -axis unless we also load finger 3. The third finger provides the contact necessary for resisting the




Fig. 11: Example grasp to illustrate the importance of appropriately choosing grasp preloads. We assume a preload of 0.1Nm at the proximal joints of fingers 1 and 2 (axes marked as green dashed lines). We apply a range of preload torques at the proximal joint of finger 3 (axis marked as yellow dashed line) and evaluate passive stability.




Fig. 12: Force in the  $z$ -direction (blue) and moment in the  $x$ -direction (red) the grasp in Fig. 11 can resist for a range of preloads at finger 3.

torque on the object. The more we load finger 3, the larger the force at contact 3 and the larger the resistible torque. At some point however, as discussed above, the forces at finger 3 begin to overwhelm fingers 1 and 2 and the object slides out along the  $z$  axis even without any external disturbances.

This example shows the importance of preload and passive stability: for this grasp, loading finger 3 helps resistance against some disturbances, but hurts against others. The right amount of preload must thus be chosen based on the task. To the best of our knowledge, no existing grasp stability analysis method can make such predictions.

Using our model, we can also find actuator commands that are optimal with respect to any specific objective we chose. For example we may want to minimize the maximum torque

Equivalent number of friction edges	Without robustness scheme			With robustness scheme		
	Maximum wrench (N)	Time at full resolution (s)	Time with hierarchical refinement (s)	Maximum wrench (N)	Time at full resolution (s)	Time with hierarchical refinement (s)
4	>100	0.088	0.088	>100	0.11	0.11
8	7.07435	0.22	0.43	7.07435	0.24	0.46
16	6.34905	0.43	0.65	3.73377	0.45	1.02
32	3.71036	0.72	1.92	3.70539	0.97	1.33
64	3.70095	1.58	2.33	2.31479	1.83	3.09
128	3.69775	5.80	3.50	1.21805	5.12	10.9
256	3.69751	>600	4.87	1.21756	20.3	13.0
512	3.69681	-	4.34	1.21743	102	17.8
1024	3.07268	-	13.3	1.21740	394	22.9
2048	3.07267	-	15.0	1.21739	-	32.2

TABLE I: Performance of the analysis of the grasp in Fig. 9. We compute the magnitude of the largest force that can be applied in the negative  $y$ -direction without destabilizing the grasp. We do this for varying levels of refinement expressed as the number of sectors the friction cone approximation contains. We record compute times for both an approach where all sectors are already of the desired size and our hierarchical refinement approach.

a single actuator must produce to resist a given wrench  $\mathbf{w}_m$ . We now use Algorithm 1 as follows:

$$\text{objective: minimize } \max_j \tau_c^j \quad (26)$$

$$\text{additional constraints: } \mathbf{w} = \mathbf{w}_m \quad (27)$$

For the grasp in Fig. 9, we can compute the optimal actuator commands for a force of  $(\mathbf{w} = [0, -1.21743, 0, 0, 0, 0]^T)$ , the largest force in the negative  $y$ -direction that can be resisted when applying a preload of 0.1Nm at every proximal joint, according to our previous analysis. We find that the optimal torques at these joints are actually only  $(\boldsymbol{\tau}_c = [0, 0.042, 0.073]^T)$ . This shows that a large amount of the preload (0.1Nm at every joint) is wasted in the sense that it does not increase disturbance resistance in this particular direction.

In many practical applications it may be of interest to take into account physical limits such as the maximum torque an actuator can apply or a maximum permissible normal force in order to not break the grasped object. Such constraints can be expressed as linear inequalities and are straightforward to add to our model.

#### D. Analysis of computational performance

In Section V we state that an accurate solution to the grasp problem with discretized friction cones but without hierarchical refinement requires a very large number of friction edges. We further argue that solving such a problem becomes computationally intractable as a large number of friction edges results in a large number of binary variables in the MIP. To verify these hypotheses, we analyze the convergence of our algorithm with varying levels of refinement of the friction approximation. We use the grasp in Fig. 9 as an example, with the task of finding the largest force in the negative  $x$ -direction that the grasp can withstand. We do so at varying levels of refinement, and record the predicted force magnitude and runtime.

We compare two approaches: The first method (“full resolution”) directly uses a friction cone approximation that is at the desired level of accuracy in its entirety. The second method (“hierarchical refinement”) always starts with a coarse approximation and refines as described in Algorithm 1. Throughout

all experiments, both methods (when able to finish) produced identical solutions, but the running times varied greatly. All recorded data can be found in Table I.

We notice that, at high levels of refinement, full resolution becomes intractable, whereas hierarchical refinement finds a solution efficiently. The study of how the refinement level affects the returned solution is more complex. The exact value of the solution generally reaches a point where increasing the accuracy of the approximation (adding more friction edges) stops making a significant difference. In some cases, as in the case of the maximum wrench in the left side of Table I, this happens for accuracy levels that only hierarchical refinement can reach. In others, as in the case of the maximum robust wrench (with  $2.5^\circ$  normal uncertainty) in the right side of Table I, both methods are able to find good approximations of the final value. At the more shallow levels, full resolution will often outperform hierarchical refinement, but since we generally do not know which of these cases any specific query might fall into, only hierarchical resolution allows us to increase the accuracy without the risk of compute time exploding. These results show both that high accuracy is actually required in order to obtain meaningful results, and that solving such problems without hierarchical refinement is computationally intractable.

Another interesting finding is that the predicted maximum resistible force is exceedingly large when only four friction edges are used in the first step of the refinement process. This is because at this stage the MDP is only enforced such that the friction force and negative relative tangential contact motion lie within a  $90^\circ$  sector. This allows sufficient freedom to the solver to use the rigidity of the robot hand when backdriven along with unphysical object motions to create large contact forces - much like what was described in Section VII-A where there are no constraints on the friction direction at all. This illustrates again the need for a high accuracy solution to the grasp problem including the MDP.

In order to investigate the practical applicability of our method, we tested the runtime of our algorithm on a range of grasps that could be encountered in a grasping task (see Fig. 13.) We generated these grasps on a range of differently shaped objects using a brute force grasp planner [31]. On each of these grasps we applied our framework to perform




Fig. 13: Algorithmically generated grasps to investigate the computational performance of our method.

Grasp	Number of contacts	Stability check	Maximum disturbance		Optimal torques	
		Time (s)	Mean time (s)	Median time (s)	Mean time (s)	Median time (s)
(a)	3	0.246	$0.797 \pm 0.067$	0.754	$0.389 \pm 0.063$	0.343
(b)	3	0.234	$0.785 \pm 0.085$	0.745	$0.447 \pm 0.074$	0.408
(c)	3	0.294	$0.893 \pm 0.056$	0.907	$0.347 \pm 0.017$	0.345
(d)	3	0.259	$0.664 \pm 0.081$	0.560	$0.308 \pm 0.036$	0.281
(e)	4	1.07	$10.0 \pm 3.4$	6.39	$3.26 \pm 0.50$	3.83
(f)	4	15.7	$36.7 \pm 6.5$	39.5	$13.5 \pm 4.8$	9.58
(g)	4	2.78	$33.8 \pm 9.3$	20.8	$16.8 \pm 5.3$	12.3
(h)	4	8.02	$23.7 \pm 4.2$	23.2	$12.3 \pm 3.3$	8.10
(i)	4	1.66	$14.4 \pm 2.7$	11.4	$1.67 \pm 0.12$	1.60
(j)	5	6.06	$150 \pm 24$	105	$22.8 \pm 5.8$	19.0
(k)	6	51.0	$398 \pm 104$	344	$136 \pm 54$	80.6
(l)	6	27.5	$488 \pm 99$	330	$777 \pm 288$	626

TABLE II: Runtime analysis of our method for the three tasks demonstrated in this paper performed on a consumer desktop computer for the grasps shown in Fig. 13. Where multiple trials were performed we report the mean  $\pm$  standard error as well as the median runtimes.

the three tasks demonstrated above: Checking for stability, finding the maximum resistible wrench in a given direction and computing the optimal actuator commands. In the second task the direction along which to find the maximum resistible disturbance can have an impact on the runtime of the algorithm. We hence repeated the solution process with ten different randomly generated directions to obtain more meaningful results. Similarly, the time it takes to optimize the actuator commands depends on the external wrench applied to the object and hence we also repeat these experiments with ten randomly sampled external wrenches. The results are recorded in Table II. In order to guarantee accurate results we continued refinement until we reached a level equivalent to 2048 sectors in the friction approximation. Empirically we have found this level of refinement guarantees convergence for all grasps tested (although many converged sooner.)

We note that, as expected, the runtime of our algorithm grows with the number of contacts. For grasps with three contacts, all queries typically had sub-second runtimes; up to and including five contacts we noticed runtimes typically between 1 and 30 seconds. Also as expected, stability checks

(the most fundamental operation needed for a grasp) are faster than computations like maximum resistible disturbance or optimal joint torques. This suggests different applications of our framework for different scenarios: pruning a larger number of possible grasps using the faster stability analysis, then computing the optimal torques only on the most promising candidates. Even with six contacts, our method had a runtime on the order of minutes, suitable for example for fixturing analysis in a manufacturing line. Finally, we also note that, in the absence of our hierarchical refinement, it is altogether intractable to approach this level of accuracy for all except the simplest of grasps.

## VIII. CONCLUSIONS

In this paper, we have described a grasp stability model that allows for efficient and accurate solution methods under realistic constraints. Noting that an exact formulation of Coulomb friction includes non-convex constraints (due to the Maximum Dissipation Principle), we use a discretization method that allows the problem to be reformulated as a piecewise convex Mixed Integer Program solvable through branch

and bound. However, such discretization methods traditionally involve a trade-off: coarse discretizations provide only rough approximations of the exact constraints, while high resolutions discretizations are computationally intractable.

To address this problem, we introduce a hierarchical refinement method that progressively increases the resolution of the discretization only in the relevant areas, guided by the solution found at coarser levels. Our local refinement method remains efficient up to high discretization resolution, and also provides strong guarantees: if a solution can not be found at a coarse approximation level, the underlying exact problem is guaranteed not to have a solution either. Combined, these two features make our method efficient for problems both with and without exact solutions. It is, to the best of our knowledge, the first time that grasp stability models incorporating Coulomb friction (along with the MDP) have been solved with such high discretization resolution.

Our motivating factor for including the MDP in our grasp stability analysis is to study what we refer to as passive grasp stability. This phenomenon occurs when torques at a hand's nonbackdrivable joints arise passively in response to a disturbance, allowing the grasp to resist the disturbance without the need for any actively applied motor torque. This is a ubiquitous situation in practice, and one that greatly expands the usability of many robot hands equipped with highly geared motors. However, no current grasp model can distinguish between the disturbances that will be resisted by passive effects and those that will not.

We thus combine our friction model with models of unilateral contacts and nonbackdrivable joints (also formulated as Mixed Integer constraints). The overall model accepts many types of queries: for example, we can analyze the space of wrenches applied to an object that a given grasp can withstand, or compute optimal joint commands given a specific object wrench. Thanks to the hierarchical refinement method, these can be solved efficiently (on the order of seconds per query) even with very high resolution approximations of the MDP.

Running our analysis method on a number of example grasps, we showed that our method predicts effects both intuitive (pressing directly against contacts is passively stable, but pulling the object away requires preload torques) and more subtle (an object wedging itself in a grasp in response to a disturbance for a given contact geometry.) In contrast, grasp stability models that do not consider the MDP produce unrealistic results, and fail to predict the dependence of disturbance resistance on applied preloads.

A limitation of our method is that, while it performs well in practice, its theoretical running time remains worst-case exponential in the level of discretization for friction constraints. Furthermore, for cases where a coarse discretization yields a sufficiently accurate solution, hierarchical refinement might be outperformed by an equivalent method with uniform resolution (although these cases are generally unknowable in advance, without actually solving up to high resolutions.)

A further limitation is an arguably narrow definition of grasp stability. An initially unstable grasp may - through movement of the fingers and object and hence the contact - eventually settle in a different stable equilibrium grasp. In such a case our

framework can only determine that the initial grasp is unstable and makes no prediction on stability of the final grasp. In order to account for the motion of an initially unstable grasp we would have to model the dynamics of the grasp. As discussed in Section III using the currently available dynamics engines to this end comes with its own difficulties.

From a practical perspective, in future work we would like to explore additional applications of our approach to grasp analysis and planning. While the runtime of our algorithm currently is too large for online grasp planning with more than a few contacts we see the main practical relevance of our method in enabling practitioners to understand and utilize passive effects in grasping and in providing labels for learning-based grasp planners. The model we introduced can allow multiple types of queries, and in this paper we have only presented some of the possible applications. Furthermore, we would like to extend our own previous work on the passive stability of underactuated hands [18]. Passive stability phenomena are an important feature of underactuation and we believe the framework presented in this paper is well suited to the investigation of such problems.

We believe the framework may also be applicable to problems encountered in the field of robotic locomotion, for instance to determine the balance of a legged robot on uneven terrain [12]. From a theoretical perspective we aim to study the possibility of deriving an algorithm with similar applicability and a guaranteed polynomial running time in all cases. In previous work we achieved this for planar grasps [19] and believe this work can provide insights into how to develop a computationally efficient algorithm for three dimensional grasping.

Implementations of the grasp model and algorithm presented in this paper are publicly available as part of the open source GraspIt! simulator [32].

## REFERENCES


- [1] C.S. Adjiman, S. Dallwig, C.A. Floudas, and A. Neumaier. A global optimization method,  $\alpha$ bb, for general twice-differentiable constrained NLPs I: Theoretical advances. *Computers & Chemical Engineering*, 22(9): 1137 – 1158, 1998.
- [2] M. Anitescu and G. Hart. Solving nonconvex problems of multibody dynamics with joints, contact, and small friction by successive convex relaxation. *Mechanics Based Design of Structures and Machines*, 31(3):335–356, 2003.
- [3] M. Anitescu and F. A. Potra. Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems. *Nonlinear Dynamics*, 14(3):231–247, Nov 1997.
- [4] M. Anitescu and A. Tasora. An iterative approach for cone complementarity problems for nonsmooth dynamics. *Computational Optimization and Applications*, 47(2): 207–235, Oct 2010.
- [5] D. Baraff. Coping with friction for non-penetrating rigid body simulation. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Tech-*

- niques, SIGGRAPH '91, pages 31–41, New York, NY, USA, 1991. ACM.
- [6] D. Baraff. Fast contact force computation for nonpenetrating rigid bodies. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '94, pages 23–34, New York, NY, USA, 1994. ACM.
- [7] E. M. L. Beale and J. J. H. Forrest. Global optimization using special ordered sets. *Mathematical Programming*, 10(1):52–69, Dec 1976.
- [8] A. Bicchi. Force distribution in multiple whole-limb manipulation. In *Robotics and Automation, 1993. Proceedings., 1993 IEEE International Conference on*, pages 196–201. IEEE, 1993.
- [9] A. Bicchi. On the problem of decomposing grasp and manipulation forces in multiple whole-limb manipulation. *Robotics and Autonomous Systems*, 13(2):127–147, 1994.
- [10] A. Bicchi. On the closure properties of robotic grasping. *The International Journal of Robotics Research*, 14(4): 319–334, 1995.
- [11] S. P. Boyd and B. Wegbreit. Fast computation of optimal contact forces. *IEEE Transactions on Robotics*, 23(6): 1117–1132, Dec 2007.
- [12] T. Bretl and S. Lall. Testing static equilibrium for legged robots. *IEEE Transactions on Robotics*, 24(4):794–807, Aug 2008.
- [13] M. Buss, H. Hashimoto, and J. Moore. Dextrous hand grasping force optimization. *IEEE Trans. on Robotics and Automation*, 12:406–418, 1996.
- [14] M. Cutkosky and I. Kao. Computing and controlling the compliance of a robotic hand. *IEEE Transactions on Robotics and Automation*, 5(2), 1989.
- [15] E. Drumwright and D. Shell. Modeling contact friction and joint friction in dynamic robotic simulation using the principle of maximum dissipation. *Algorithmic Foundations of Robotics IX: Selected Contributions of the Ninth International Workshop on the Algorithmic Foundations of Robotics*, pages 249–266, 2011.
- [16] C. Ferrari and J. Canny. Planning optimal grasps. In *IEEE International Conference on Robotics and Automation*, pages 2290–2295, 1992.
- [17] LLC Gurobi Optimization. Gurobi optimizer reference manual, 2018. URL <http://www.gurobi.com>.
- [18] M. Haas-Heger, G. Iyengar, and M. Ciocarlie. Passive reaction analysis for grasp stability. *Transactions on Automation Science and Engineering*, 15(3), 2018.
- [19] M. Haas-Heger, C. Papadimitriou, M. Yannakakis, G. Iyengar, and M. Ciocarlie. Passive static equilibrium with frictional contacts and application to grasp stability analysis. In *Robotics: Science and Systems*, 2018.
- [20] L. Han, J. Trinkle, and Z. Li. Grasp analysis as linear matrix inequality problems. *IEEE Trans. on Robotics and Automation*, 16:663–674, 2000.
- [21] H. Hanafusa and I. Asada. Stable prehension by a robot hand with elastic fingers. In *Proc. of the 7th ISIR, Tokyo*, 1977.
- [22] P. C. Horak and J. C. Trinkle. On the similarities and differences among contact models in robot simulation. *IEEE Robotics and Automation Letters*, 4(2):493–499, April 2019.
- [23] D. Kaufman, S. Sueda, D. James, and D. Pai. Staggered projections for frictional contact in multibody systems. *ACM Transactions on Graphics (SIGGRAPH Asia 2008)*, 27(5):164:1–164:11, 2008.
- [24] J. Kerr and B. Roth. Analysis of multifingered hands. *International Journal of Robotics Research*, 4(4):3–17, 1986.
- [25] David Kirkpatrick, Bhubaneswar Mishra, and Chee-Keng Yap. Quantitative steinitz’s theorems with applications to multifingered grasping. *Discrete & Computational Geometry*, 7(3):295–318, Mar 1992.
- [26] Leo Liberti. Reformulation and convex relaxation techniques for global optimization. *Quarterly Journal of the Belgian, French and Italian Operations Research Societies*, 2(3):255–258, Oct 2004.
- [27] Jeffrey Mahler, Jacky Liang, Sherdil Niyaz, Michael Laskey, Richard Doan, Xinyu Liu, Juan Aparicio, and Ken Goldberg. Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. 07 2017.
- [28] M. Mason. *Mechanics of Robotic Manipulation*. MIT Press, Cambridge, MA., 2001.
- [29] R. Mattikalli, D. Baraff, and P. Khosla. Finding all stable orientations of assemblies with friction. *IEEE Transactions on Robotics and Automation*, 12(2):290–301, 1996.
- [30] Garth P. McCormick. Computability of global solutions to factorable nonconvex programs: Part i — convex underestimating problems. *Mathematical Programming*, 10(1):147–175, Dec 1976.
- [31] C. Meeker, M. Haas-Heger, and M. Ciocarlie. Computing an intuitive teleoperation mapping empirically and algorithmically. *IEEE Transactions on Robotics*, 2019 (submitted).
- [32] A. Miller and P. Allen. GraspIt!: a versatile simulator for robotic grasping. *IEEE Robotics and Automation Magazine*, 11(4):110–122, 2004.
- [33] Y. Nakamura, K. Nagai, and T. Yoshikawa. Dynamics and stability in coordination of multiple robotic mechanisms. *The International Journal of Robotics Research*, 8(2):44–61, 1989.
- [34] V.D. Nguyen. Constructing force-closure grasps. *International Journal of Robotics Research*, 8(3):3–16, 1988.
- [35] R.S. Palmer. *Computational Complexity of Motion and Stability of Polygons*. PhD thesis, Cornell University, 1989.
- [36] J.-S. Pang and J. Trinkle. Stability characterizations of rigid body contact problems with coulomb friction. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 80(10), 2000.
- [37] J.-S. Pang, J. Trinkle, and G. Lo. A complementarity approach to a quasistatic multi-rigid-body contact problem. *Computational Optimization and Applications*, 5(2): 139–154, Mar 1996.

- [38] T. Pang and R. Tedrake. A robust time-stepping scheme for quasistatic rigid multibody systems. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5640–5647, 2018.
- [39] D. Prattichizzo, J. Salisbury, and A. Bicchi. Contact and grasp robustness measures: Analysis and experiments. *Experimental Robotics IV: The 4th International Symposium, Stanford, California, June 30 – July 2, 1995*, pages 83–90, 1997.
- [40] F. Reuleaux. *The Kinematics of Machinery*. Macmillan, 1876.
- [41] J.K. Salisbury. *Kinematic and Force Analysis of Articulated Mechanical Hands*. PhD thesis, Stanford University, 1982.
- [42] J.K. Salisbury and B. Roth. Kinematic and force analysis of articulated mechanical hands. *ASME Journal of Mechanisms, Transmissions, and Automation in Design*, 105:35–41, 1983.
- [43] P. Song, J.-S. Pang, and V. Kumar. A semi-implicit time-stepping model for frictional compliant contact problems. *International Journal for Numerical Methods in Engineering*, 60(13):2231–2261, 2004.
- [44] D. Stewart. Rigid-body dynamics with friction and impact. *SIAM Rev.*, 42(1):3–39, March 2000.
- [45] D. E. Stewart and J. C. Trinkle. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *International Journal for Numerical Methods in Engineering*, 39(15):2673–2691, 1996.
- [46] A. Tasora and M. Anitescu. A matrix-free cone complementarity approach for solving large-scale, nonsmooth, rigid body dynamics. *Computer Methods in Applied Mechanics and Engineering*, 200(5):439 – 453, 2011.
- [47] M. Tawarmalani and N. V. Sahinidis. A polyhedral branch-and-cut approach to global optimization. *Mathematical Programming*, 103:225–249, 2005.
- [48] E. Todorov. Implicit nonlinear complementarity: A new approach to contact dynamics. In *2010 IEEE International Conference on Robotics and Automation*, pages 2322–2329, 2010.
- [49] E. Todorov. Convex and analytically-invertible dynamics with contacts and constraints: Theory and implementation in mujoco. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6054–6061, May 2014.
- [50] J. C. Trinkle and D. C. Zeng. Prediction of the quasistatic planar motion of a contacted rigid body. *IEEE Transactions on Robotics and Automation*, 11(2):229–246, April 1995.
- [51] J. C. Trinkle, J.-S. Pang, S. Sudarsky, and G. Lo. On dynamic multi-rigid-body contact problems with coulomb friction. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift fr Angewandte Mathematik und Mechanik*, 77(4):267–279, 1997.
- [52] Hoang Tuy. *Convex Analysis and Global Optimization*. Springer Publishing Company, Incorporated, 2nd edition, 2016. ISBN 3319314823, 9783319314822.



**Maximilian Haas-Heger** Maximilian Haas-Heger received the MEng degree in Aeronautical Engineering from Imperial College London in 2015. Since 2015, he is a PhD candidate in the Robotic Manipulation and Mobility Lab at Columbia University in New York. His research focuses on the theoretical foundations of robotic grasping; specifically on the development of accurate grasp models and their application for dexterous manipulation.



**Matei Ciocarlie** Matei Ciocarlie (S'07-M'12) earned the Ph.D. degree in Computer Science from Columbia University, New York, NY, USA. He was a Research Scientist at Willow Garage, Inc., Menlo Park, CA, USA, and a Senior Research Scientist with Google, Inc., Mountain View, CA, USA. He is currently an Associate Professor of Mechanical Engineering with affiliated appointments in the Computer Science Department and Data Science Institute at Columbia University, New York, NY, USA. His work focuses on robot motor control, mechanism and sensor design, planning and learning, all aiming to demonstrate complex motor skills such as dexterous manipulation.