

Grasping the Dice by Dicing the Grasp

Ch. Borst, M. Fischer and G. Hirzinger
Institute of Robotics and Mechatronics
German Aerospace Research Center (DLR), Wessling, Germany
Christoph.Borst@dlr.de, Max.Fischer@dlr.de

Abstract—Many methods for generating and analyzing grasps have been developed in the recent years. They gave insight and comprehension of grasping with robot hands but many of them are rather complicated to implement and of high computational complexity.

In this paper we study if the basic quality criterion for grasps, the force-closure property, is in principle easy or difficult to reach. We show that it is not necessary to generate optimal grasps, due to a certain quality measure, for real robot grasping tasks where an average quality grasp is acceptable. We present statistical data that confirm our opinion that a randomized grasp generation algorithm is fast and suitable for the planning of robot grasping tasks.

I. INTRODUCTION

Mobile robots, either wheeled or legged, with integrated arms and hands have become more and more important in the last years. Especially the community that deals with humanoid robots has grown vastly. The purpose of these systems is to use them as personal service assistants in quite unstructured environments primarily built for humans like an industrial workbench, ones kitchen or living room.

Therefore one very important capability of such systems is the autonomous handling of objects. At least these robots should be able to safely grasp objects, carry them where the user wants them and then put them safely down again. The environments these robots should work in are changed permanently by humans, so *offline* pre-planning of grasping and manipulation actions seems no proper way to ensure this capability. *Online* grasp and manipulation planning systems are needed. These systems may not generate optimal grasps and manipulation actions. Their most important goal is to find a sufficiently good solution as fast as possible.

Many different approaches have been presented to plan and analyze different grasp types for robot hands. Approaches that try to adapt a set of generic grasp shapes to a given grasp candidate are promising for the human like power grasp where the whole hand can have contact with the grasp object [12] [14]. Whereas most interest was given to the simpler precision grasp, where only the fingertips are in contact with the object. This type of grasp is well understood. Many criteria as force- and form-closure and different quantitative quality indices have been developed to rate the quality of precision grasps [7] [9].

At our institute we are most interested in algorithms for the fast and autonomous generation of precision grasps. Many methods to calculate precision grasps have been




Fig. 1. The DLR-Hand II playing dice.

presented. Most deal with the construction of so called force-closure grasps (see definition below) and some try to get optimal grasps. Almost all have in common that their computational complexity is dependent on the number of surface patches of the objects to be grasped. The example objects presented are seldom composed of more than 20 faces. For real world objects generated by sensor information from cameras or laser scanners the number of faces will normally range from 10^3 - 10^5 . So these algorithms are expected to run very slow.

With the work presented here we want to get an idea of how hard it is to get a suitable grasp for everyday tasks. It is obvious that finding the optimal grasp is of high complexity but we show that finding a fairly good force-closure grasp is easy as there exist many good grasps. We randomly generated 10^6 grasp candidates on a set of general test objects, performed a force-closure test and evaluated their quality according to our selected quality measure. From these results we can derive a ratio of force-closure grasps to all grasps and show the distribution of grasp quality on the test objects.

For better understanding what we calculate and measure for each grasp candidate we first outline some basic quality criteria for precision grasps and show how to formalize the grasp with appropriate models. Then we show the statistics

of force-closure grasp on the set of objects and present a simple way to significantly increase the ratio of force closure grasps generated. Last we present some results with a quality measure and show the “almost optimal” grasps on some objects.

II. BASIC GRASP THEORY

We define a grasp as a set of contacts on the surface of the object. The forces or torques the manipulator can exert in these contacts depend on the contact model and on the abilities of the manipulator. Here we only consider precision grasps, where only the fingertips are in contact with the object, so we can focus on the contact model and neglect manipulator constraints. After defining the contact model one can formalize commonly used closure properties of a grasp and we give a quality measure for grasps proposed by Ferrari and Canny [3]. These are all prerequisites for our statistical analysis of grasps on a set of generic objects.

A. Grasp contact

The mostly used fingertip contact models are hard-finger contacts with and without friction and soft-finger contacts [8] [11]. All these contacts can be modeled as single point contact. In the first case, hard-finger without friction, only forces against the surface normal in the contact point can be exerted on the object (fig. 2 A). With friction all forces that lie within the friction cone around the surface normal can be exerted. The cone angle is defined as $\gamma = \text{atan}(\mu)$ where μ is the friction coefficient. (fig. 2 B). With soft-finger contacts a torque around the normal can also be applied in the contact point (fig. 2 C).

There are also more realistic but more complicated models, especially for soft-finger contacts [15]. For the following considerations the simple ones are sufficient.




Fig. 2. Finger contact models: A - hard-finger frictionless, B - hard-finger with friction, C - soft-finger.

In the following sections we assume to have hard finger contacts with friction (Fig. 2 B). The other contact types are either unrealistic, as the one without friction, or unnecessary complex as the soft finger contact where the effect of the torque around the surface normal can almost be neglected for more than two contacts.

B. Closure Properties

The effect of forces applied in each contact point can be displayed as a so called 6-dimensional *wrench* with a force and torque component. A wrench is always related to

a freely selectable reference point (eg. center of mass). A force F_a for example that is acting in contact A results in the wrench $w_a = \begin{pmatrix} F_a \\ (A-R) \times F_a \end{pmatrix}$ with a reference point R .

To prevent a grasp with N contacts from slipping the forces in contacts i and their corresponding wrenches w_i and any disturbing forces or wrenches w_{ext} have to be in equilibrium:

$$\sum_{i=1}^N w_i + w_{ext} = 0. \quad (1)$$

Two commonly used closure properties have been defined for grasps where the equilibrium can always be achieved regardless of the direction of the counter force. The forces that can be applied by the manipulator are supposed to be unbound.

1) *Force-Closure*: A grasp is called *force-closure* if any disturbing wrench w_{ext} can be balanced by the wrenches applied at the contacts.

With a force-closure grasp the direction of the force applied by the manipulator in a contact point may vary within the friction cone. To find the right wrenches to keep a grasp in equilibrium is a task for the grasp controller[4].

2) *Form-Closure*: A grasp is called *form-closure* if it is force-closure with no friction at contacts present.

Intuitively form-closure means that the grasp is fully immobilized in a passive manner. The manipulator can be set absolute stiff and with any disturbing force acting on the object equilibrium is self-appearing.




Fig. 3. Sample grasps on a dice with projection into a plane to illustrate closure (A) and non closure (B).

C. A Geometrical View

For the analysis of the closure properties in vector space we simplify the friction cones by approximating them as polyhedral cones. Also we take the simple planar grasp of figure 3 A/B for illustration and only discuss the 3-dimensional wrench space. All main aspects can be seen in this example and transfer to 3-dimensional grasps with 6-dimensional wrench space is straight forward.

From each contact we can exert forces that lie within the friction cone. The cone C_i is spanned by the two spanning vectors u_{1c_i} and u_{2c_i} (fig. 4). The space that is spanned by all contacts is called the *grasp wrench space (GWS)* and characterizes the ability of the grasp to balance disturbance forces. With this geometric model the question if a grasp

is force-closure can easily be answered: A grasp is force-closure if its GWS contains a small environment around the origin. A common way to test the force-closure property is to approximate the GWS with the convex hull over all friction cone spanning vectors and check if the origin is inside the hull (Fig.4).




Fig. 4. GWS approximation of the grasps on the dice in force domain (2D) (A - with 3 cones C_1, C_2, C_3 contains the origin, B - with 2 cones C_1, C_2 does not contain origin).

For the grasps in figure 4 the GWS and the convex hull of the spanning vectors can be written as:

$$GWS = \left\{ \sum_{i=1}^N \sum_{j=1}^2 \alpha_{i,j} u_{jC_i} \mid \alpha_{i,j} \geq 0 \right\}. \quad (2)$$

$$ConvHull = \left\{ \sum_{i=1}^N \sum_{j=1}^2 \alpha_{i,j} u_{jC_i}, \mid \alpha_{i,j} \geq 0 \wedge \sum_{i,j} \alpha_{i,j} \leq 1 \right\}. \quad (3)$$

The $\alpha_{i,j}$ are two scalars per cone that allow to describe any linear combination of the spanning vectors u_j of cone C_i which result in all vectors that lie within the cone.

It holds true that $ConvHull \subset GWS$. Therefore if the origin lies in the convex hull it must also lie in GWS. Conversely it holds true that if a small neighborhood of the origin is contained in the GWS it must also be contained in the convex hull, because for small neighborhoods the upper bounds of the coefficients α are of no relevance.

D. Polyhedral Convex Cones Theory

Another type of modeling the force-closure grasp problem is convex cones theory, introduced for grasping by S. Hirai in 1991 [5]. Afore we modeled the GWS and found a way to approximately compute it. Now we describe the space of wrenches the grasp cannot resist. It is clear that with a force-closure grasp this space should be empty. The computation of that space is as complex as calculating the approximation of GWS but to test if a vector lies in this space is easy to calculate, so we have a fast verification that the grasp is not force-closure, if we can guess such a vector.

We approximate the friction cones with polyhedral (convex) cones. A polyhedral convex cone can be written in

$span$ or in $face$ form where u_j are its spanning and n_i are its face (face normal) vectors (see fig. 5).

$$\begin{aligned} C &= span\{u_1, u_2, \dots, u_k\} \\ &= face\{n_1, n_2, \dots, n_k\}. \end{aligned} \quad (4)$$

The so called *polar* of a set of vectors is defined as:

$$C^* = \{y \mid x^T y \leq 0, \forall x \in C\}. \quad (5)$$

Therefore we can write C^* also in $span$ and $face$ form:

$$\begin{aligned} C^* &= face\{u_1, u_2, \dots, u_k\} \\ &= span\{n_1, n_2, \dots, n_k\}. \end{aligned} \quad (6)$$




Fig. 5. A polyhedral convex cone C with its span (u) and face (n) vectors and a planar view on the polar C^* .

The space of vectors W that may break the grasp can then be written as the intersection of all polars of the friction cones C_i .

$$W = C_1^* \cap C_2^* \cap \dots \cap C_N^*. \quad (7)$$

As already mentioned the construction of this space is as complex as the construction of the GWS. But it is easy to check, if a given external force vector F_{ext} (the same holds in the 3D case for a 6D external wrench vector) lies in W by simply calculating the scalar product of F_{ext} with all spanning vectors. If the angles α_i between the force vector F_{ext} and the normals of all the contacts i is bigger than $90^\circ + atan(\mu)$, where μ is again the friction index, then F_{ext} is an element of W (see fig. 6 for illustration).

To guess such a vector F_{ext} , which breaks a given grasp defined by the contacts with higher probability than a randomly chosen F_{ext} , we tried two heuristics. The first selects the two contact normals with the largest angle between them and takes the opposite direction of the bisector of the normals as the guess for F_{ext} . The second calculates the average of all contact normals and takes the opposite direction as guess for F_{ext} . Although the two heuristics give different vectors, we found that their performance in finding a grasp breaking vector is almost the same.

In our grasp planning algorithm we randomly generate grasp candidates by selecting contacts on the object surface

with equal probability for each surface point. Then we test if the candidate is force-closure. This procedure can be sped up by pre-filtering the random candidates using the check and heuristics described before: only for the candidates for which the guessed F_{ext} does not break the grasp a full force-closure test is performed. The pre-filtering is conservative, so we do not reject any force-closure candidate. In section III we present results on the effectiveness of this fast pre-filtering of grasp candidates.




Fig. 6. The geometric construction of the polar space of the union of two convex cones $C_1 \cup C_2 = C_1^* \cap C_2^*$. For a vector F_{ext} the angles α_1 and α_2 have to be checked.

For detailed description of polyhedral convex cones and proof of theorems we refer to [5] and [6].

E. Quality index

For a good grasp it is a minimal requirement to be force-closure, but in the set of force-closure grasp there are still grasps with very different quality. If we assume that nothing special about the task for which the grasp should be generated is specified, a good grasp should be able to resist wrenches in any direction equally good. We use a quality measure as proposed in [3]. There the quality of a grasp is defined as the length of the smallest wrench (in any direction) that breaks the grasp, when in every contact a force with unit strength is applied. This measure can be computed by calculating the largest inscribing ball in the GWS around the origin (see fig. 7).




Fig. 7. Grasp-Wrench-Space (GWS) approximation and largest inscribing ball for two sample force-closure grasps (Grasp B has better quality index than grasp A).

For details on the definition, the efficient calculation and a discussion about friction issues of this measure we refer to the original papers [1] [3] [13]. Here we only want to give an idea about the quality index we use in section IV to calculate the histogram.

III. RATIO OF FORCE-CLOSURE GRASPS ON A SET OF TEST OBJECTS

A. The test scenario

We selected a set of basic geometrical objects (fig. 8) as they allow an easy interpretation of the results and furthermore many real world objects can be constructed with these basic primitives. A second set of “real world objects” (fig. 9) should prove that the results also hold true for our aimed setting in service robotics.




Fig. 8. The set of primitive geometric test objects.




Fig. 9. The set of real world objects.

B. Results

On each object we randomly generated 10^6 3-, 4-, and 5-finger grasp candidates and tested them for the force-closure property. In the figures 10 - 12 the ratio of force-closure grasps for all generated candidates and the ratio of force-closure grasps for the candidates, which passed the pre-filtering as described in section II-D, is shown. For the friction coefficient we chose $\mu = 0.5$, which is a typical value for the rather adverse case that our robot hand with silicon coated finger tips grasps an object with a metal surface.

It is evident that the computationally simple heuristic pre-filtering increases the ratio of force-closure grasps, especially for 3- and 4-finger grasps.

The figures show, that for the tested not extraordinary complex objects the force closure ratio for 4- and 5-finger grasps is not less than 20%. That means to get with a




Fig. 10. Force-closure ratio for 3-finger grasps with friction $\mu=0.5$.




Fig. 11. Force-closure ratio for 4-finger grasps with friction $\mu=0.5$.

probability of 99.9% a force-closure grasp one has to generate not more than 31 grasp candidates. All computations, including the force closure test for all the candidates, takes about 0.6 seconds on a Pentium III/900MHz.

IV. DISTRIBUTION OF GRASP QUALITY FOR A SET OF TEST OBJECTS

In our grasp planning system we first randomly generate grasp candidates, then we use the heuristic pre-filter to reduce the number of non force-closure candidates and at the end we rate the quality of the candidates with the measure outlined in section II-E. In this section we show that this




Fig. 12. Force-closure ratio for 5-finger grasps with friction $\mu=0.5$.

procedure allows us to find a sufficiently good grasp with a reasonable computational effort. By “sufficiently good” we mean here a grasp, which has a similar quality, in terms of our quality measure, as the one humans would typically perform for the given task.

For different objects we generated 10^6 force-closure grasps and calculated their quality measure. Figures 13–15 show the normalized histogram (area is equal to one), which can be fitted very well by a gamma distribution.




Fig. 13. Frequency distribution over the quality measure for grasps on a rectangular box. The inset shows a grasp typically performed by a human. The vertical bar marks the corresponding quality value.

At a first view there are many grasps with weak quality measure and only a few very good grasps. Therefore it seems difficult to generate a good grasp at random. But the question for us is, what is a good or suitable grasp for real everyday grasping tasks? To tackle this question we select for each object a grasp that humans typically perform on the object and compare its quality measure with the randomly generated grasps. These human grasp samples are derived from the grasp taxonomy proposed by Cutkosky and Howe [2] and the sample grasps showed in Napiers book on hands [10]. Where more than one grasp type stands to reason we choose the one with the best quality index to be conservative with the following conclusions. Each candidate with a higher quality measure than the human like grasp we can accept from our planner. The probability to get such a grasp ranges from 8% (Cone2) to 55% (Sphere). So after generating 12 (Sphere) upto 83 (Cone2) candidates the probability to get a grasp with better quality than the human preferred one is more than 99.9%. The calculation of 100 grasp candidates with measuring their quality takes about 1.8 seconds on a Pentium III/900MHz.

V. THE OPTIMAL GRASPS

It is also interesting to see which are the optimal grasps we received from our calculations. The figures show that the quality measure seems to rate the grasps in a correct physical manner but many of these optimal grasp cannot be realized by usual robot hands.

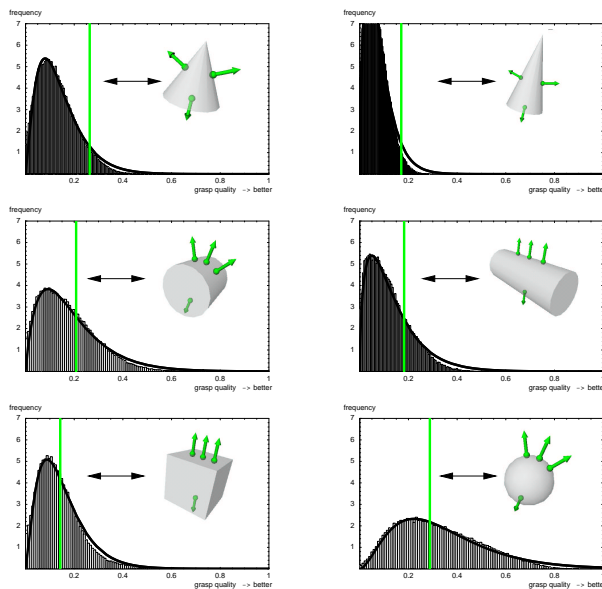


Fig. 14. Frequency figures for the basic object set with manually chosen grasp and its corresponding quality measure marked

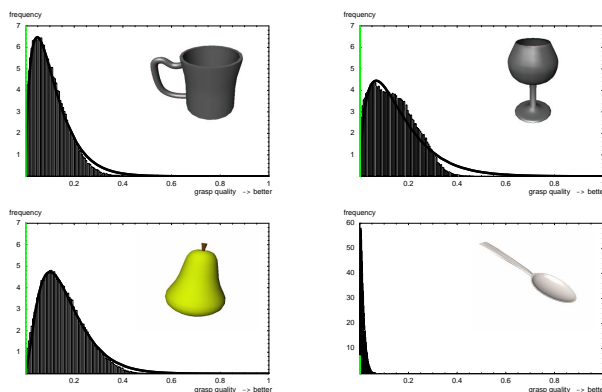


Fig. 15. Frequency figures of real world objects

VI. CONCLUSION

We generated many grasp candidates and determined the force-closure ratio within these grasps and prepared frequency distributions for the grasp quality on a set of test objects. We can summarize the results:

With a strategy to randomly generate grasps and filter them with simple heuristics the calculation of force-closure grasps can be done very fast. Even the generation of quality rated grasps can be done with such a strategy and




Fig. 16. Some optimal grasps on basic objects (best of 10^6 candidates)

lead to sufficiently good grasps after generation of only few candidates. Furthermore this method is quite easy to implement. The geometrical form of the object to grasp affects the number of force-closure grasps one can generate for the object so the calculation complexity of any grasp planner should depend on the objects form. In difference to many other approaches the complexity of our random grasp planning strategy is only determined by the objects form and not by the number of faces the objects surface is composed of. This is an important fact to consider if online planning of autonomous sensor built object models is intended. The only thing to criticize is that for some extraordinary object geometries the method will take a lot of time when only a few force-closure grasps are possible on the objects surface. The fact that no optimal grasps are generated is not of practical relevance for most objects as these optimal grasps are not realizable for most robot hands.

VII. REFERENCES

- [1] Ch. Borst, M. Fischer, and G. Hirzinger. A fast and robust grasp planner for arbitrary 3d objects. In *Proc. IEEE Conf. on Robotics and Automation*, pages 1890–1896, Detroit, Michigan, May 1999.
- [2] Mark R. Cutkosky and Robert D. Howe. Human grasp choice and robotic grasp analysis. In Subramanian T. Venkataraman and Thea Iberall, editors, *Dextrous Robot Hands*, chapter 1. Springer Verlag, 1990.
- [3] C. Ferrari and J. Canny. Planning Optimal Grasps. In *Proceedings of the IEEE Int. Conference on Robotics and Automation*, pages 2290–2295, Nice, France, May 1992.
- [4] S. Haidacher, T. Schlegl, and M. Buss. Grasp evaluation based on unilateral force closure. In *Proceedings of the IEEE/RSJ Intl. Conference on Intelligent Robots and Systems 1999*, Kyongju, Korea, pages 424–429, 1999.
- [5] Shinichi Hirai. *Analysis and Planning of Manipulation Using the Theory of Polyhedral Convex Cones*. PhD thesis, Kyoto University, 1991.
- [6] Shinichi Hirai and Haruhiko Asada. Kinematics and Statics of Manipulation Using the Theory of Polyhedral Convex Cones. *The Int. Journal of Robotics Research*, pages 434–447, October 1993.
- [7] Yun-Hui Liu. Qualitative Test and Force Optimization of 3-D Frictional Form Closure Grasps Using Linear Programming. *IEEE Transactions on Robotics and Automation*, 15(1):163–173, February 1999.
- [8] M. T. Mason and J. K. Salisbury. *Robot Hands and the Mechanics of Manipulation*. The MIT Press, Cambridge, MA, 2 edition, 1986.
- [9] Brian Mirtich and John Canny. Easily Computable Optimum Grasps in 2-D and 3-D. In *Proc. IEEE Conf. on Robotics and Automation*, pages 739–747, San Diego, CA USA, May 1994.
- [10] John Napier. *Hands*. Princeton Science Library. Princeton University Press, Princeton, New Jersey, 1993. revised by Russell H. Tuttle.
- [11] Van-Duc Nguyen. Constructing Force-Closure Grasps. *Int. Journal of Robotics Research*, 7(3):3–16, June 1988.
- [12] Nancy S. Pollard. Synthesizing Grasps from Generalized Prototypes. In *Proc. IEEE Conf. on Robotics and Automation*, pages 2124–2130, Minneapolis, Minnesota, April 1996.
- [13] Marek Teichmann and Bud Mishra. The Power of Friction: Quantifying the "Goodness" of Frictional Grasps. In A. K. Peters, editor, *Algorithms for Robotic Motion and Manipulation*, pages 311–320. A. K. Peters, Wellesley, MA, USA, 1997.
- [14] David Wren and Robert B. Fisher. Dextrous Hand Grasping Strategies Using Preshapes and Digit Trajectories. In *Proc. IEEE Int. Conf. on Systems, Man and Cybernetics*, 1995.
- [15] Nicholas Xydias and Imin Kao. Modelling of contact mechanics and friction limit surfaces for soft fingers in robotics, with experimental results. *The International Journal of Robotics Research*, 18(8):941–950, September 1999.