IEEE/RJS Int. Conf. Intelligent Robots and Systems, Lausanne, Switzerland, Sep. 2002, pp 1578-1584

Robust Finger Gaits from Closed-Loop Controllers

Manfred Huber[1], Roderic A. Grupen[2]

[1] *Dept. of Computer Science and Engineering, University of Texas, Arlington, TX, USA, huber@cse.uta.edu*
[2] *Dept. of Computer Science, University of Massachusetts, Amherst, MA, USA, grupen@cs.umass.edu*

Abstract

Object manipulation using multi-fingered dextrous robot hands poses formidable control challenges due to the complexity of the manipulator and its workspace limitations. To address large-scale manipulation tasks, finger gaits are often used to reposition fingers. In this paper, an approach to finger gaiting is presented which constructs behavior on-line by activating combinations of reusable feedback control laws with formal stability and convergence properties drawn from a control basis. Finger gaits are constructed as finite state control strategies in a DEDS framework, while actual contact locations and object motions are computed reactively based on local contact information. This leads to manipulation strategies that are robust with respect to a limited range of perturbations, object geometries, and manipulator kinematics. To demonstrate this, two finger gaits are constructed and applied to different object geometries and robot hands.

1 Introduction

Grasping and object manipulation using a multi-fingered hand poses many control challenges due to the complexity of the manipulator, its inherent workspace limitations, and the wide range of available object geometries and manipulation tasks. Moreover, such systems often have to operate in uncertain and unstructured environments.

Over the last 20 years a number of approaches for grasp formation and object manipulation have been proposed. Grasp formation has been focused mainly on the calculation of contact points for form closure grasps (grasp configurations which constrain the geometry of the object), and force closure grasps (configurations that enable the contacts to withstand limited perturbations in the force domain). Object manipulation generally requires to move and reorient the object while maintaining a stable grasp configuration. A number of approaches have been proposed to implement and execute a known object trajectory with a dextrous hand [18, 17, 5]. For larger-scale manipulations of the object, the workspace limitations of the fingers impose additional constraints. In particular, it is often necessary to reposition the fingers throughout the manipulation task in order to maintain the object in an area that permits fine manipulations. This repositioning has to occur while a stable grasp is maintained and is generally achieved using finger gaiting [16, 8].

In a number of approaches [13, 11], finger gaits are used mainly to replace contacts that are established by fingers that reached their workspace limits. The driving force behind the manipulation strategy is a trajectory planner for the object, and finger relocation occurs only if a finger reaches its workspace limit. These strategies therefore require that an additional finger is available and that this finger can reach a location on the object that establishes a stable grasp without one of the limiting contacts. Furthermore, the success of the fine motion strategies and grasping algorithms used relies heavily on precise geometric information about the object. Similarly, approaches that make use of geometric grasp primitives [19] require rather precise information about the object geometry *a priori* to permit the precomputation of sequences of the primitives that can lead to the desired object manipulation.

Another approach to finger gaiting is to preplan complete finger trajectories for the object manipulation that already include finger relocations [12, 2]. This permits finger replacements even when none of the fingers has reached its workspace limit, thus increasing the chances of finding a successful finger gait for the desired object manipulation. On the other hand, it results in a very complex planning problem and relies on the availability of precise information about the object and the manipulator. Moreover, uncertainties and perturbations at run-time can easily lead to the need for frequent replanning.

Discrete Event Systems have also been used to achieve object manipulation [20]. Here, a finite number of legal contact locations are predetermined and a discrete event supervisor determines which contact to establish or to remove. Again, the object geometry has to be known *a priori* in order to determine a finite set of valid grasp points.

The approach proposed here provides a more reactive control scheme for finger gaits. Here, manipulation strategies are generated in the general tradition of behavior-based control approaches [1, 21] using the control basis approach. Manipulation gaits take the form of sequences of reactive closed-loop controllers where contact configurations on the object are computed on-line, making it possible to compensate for run-time errors. Moreover, contact configurations and large-scale object motions are driven by explicit stability and workspace criteria rather than by a fixed, preplanned object trajectory. The actual large-scale motion arises partly as a byproduct of the chosen gait sequence. This largely geometry-independent character of the finger gaits and the careful design of the basis controllers permits, within the limitations of the manipulator kinematics and sensory resources, to transfer finger gaits to different manipulators and varying object geometries. In the following we will demonstrate this by means of an object reorientation and an object translation task on two different manipulation platforms and object geometries.

2 The Control Basis Approach

The approach presented in this paper addresses the control of manipulation tasks in a "bottom-up" fashion. Task-specific behavior is composed on-line from a basis set of reusable, declarative feedback control laws, $\{\Phi_i\}$, with formal stability and convergence results. Declarative in this context implies that elements of this control basis are designed to reflect generic control strategies, largely independent of the task geometry and robot kinematics. Careful design of the control basis leads to a small set of control laws that span a broad class of tasks and are robust with respect to perturbations. Control is derived on-line by associating input resources (sensors or sensor abstractions), $\{\sigma_i\}$, and output resources (actuators), $\{\tau_i\}$, with elements of the control basis. The resulting bound controllers, $\Phi_i\frac{\sigma_j}{\tau_k}$ are activated concurrently under the "subject to" ("$\triangleleft$") constraint, resulting in composite controllers of the form $\Phi_S \triangleleft \Phi_D$. This composition mechanism, similar to pseudoinverse control [24] where lower level controllers can only operate within the nullspace of the higher level controller, limits control output of "subordinate" controllers, Φ_S, to actions which do not counteract the control objectives of "dominant" controllers, Φ_D. This ensures that the composite controllers inherit the formal properties directly from the stability and robustness of the underlying control basis. A complete manipulation policy that includes finger gaiting is derived as a sequence of concurrent control activations. The stability and convergence properties of the underlying controllers, in turn, permit transitions to be limited to convergence events, facilitating the representation of finger gaits as finite state supervisors in a discrete event dynamic system (DEDS) framework [23, 14]. Due to the predictable performance of the control basis, the supervisor can be modeled on a logical level with all sensory patterns interpreted within the elemental controllers. States in this supervisory automaton are characterized by patterns of controller equilibria achieved upon convergence of the active controller. To accomplish complex tasks, the composition policy leads the system through a set of favorable equilibria. Figure 1 illustrates the complete policy formation process used to generate task-specific manipulation policies from a small set of general control laws and the available resources of the robot system.

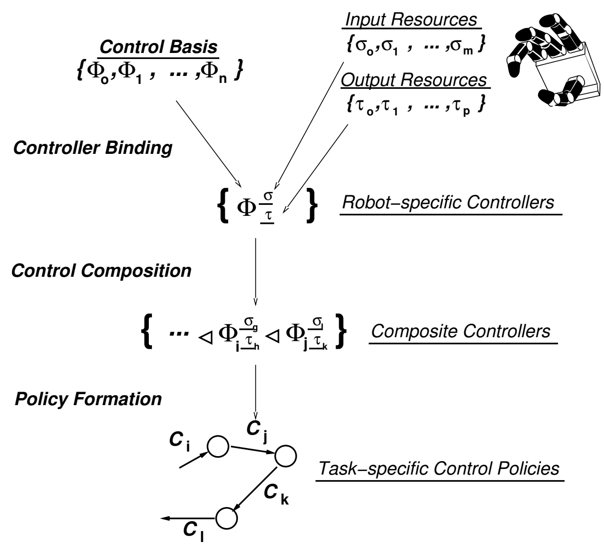

Figure 1: *Control Policy Formation in the Control Basis Approach*

Different tasks are achieved by different composition functions over the same basis controllers, rather than by the design of new task-specific behaviors. In addition, the kinematics-independent nature of the underlying control basis allows the same set of controllers to be used on different robot platforms.

2.1 A Control Basis for Dextrous Manipulation

The choice of initial control laws has a large influence on the expressiveness of the overall control system. In particular, the set of basis control laws largely determines the range of tasks that can be addressed. The control basis proposed here follows the general tradition of hybrid position/force control approaches [22] which decomposes generic robot control tasks into independent force and position control components. Here, three feedback control laws are used which address generic robot control objectives, namely configuration space position control, contact force control, and kinematic conditioning. Although a detailed description of the controllers is beyond the scope of this paper, this section will give a short overview of each along with a brief discussion of the composition mechanism.

Φ_p: Configuration Space Motion Control

The objective of this position control task is for the robot to achieve a given reference configuration while avoiding all obstacles present in its workspace. Here, harmonic function path control [7] is used to generate robust motion through the configuration space of the robot. This potential field approach provides a reactive feedback control law that is complete and correct. Harmonic functions satisfy Laplace's equation

$$\nabla^2 \phi_p = \sum_{i=1}^{n} \frac{\partial^2 \phi_p}{\partial q_i^2} = 0$$

which constrains the second derivatives of the potential surface ϕ_p with respect to the configuration space parameters q_i. This results in important formal properties, the most relevant of which are that the control law, derived by monotonically descending this surface, minimizes the probability that the robot will hit a mapped obstacle, making it robust with respect to local perturbations [7, 6].

Φ_f: Contact Configuration Control

The second control objective addressed in the control basis is to achieve a reference interaction force with the environment by distributing multiple contacts over environmental geometries. This is addressed by locally optimizing contact locations and normals and is used to form candidate grasps on an object. The contact configuration control law [3, 4] used here is based on the net wrench residual ϵ:

$$\epsilon = \sum_{cont} w_{\vec{cont}} - w_{\vec{ref}} \ , \ \vec{w_i} = \begin{pmatrix} \vec{f_i} \\ \vec{m_i} \end{pmatrix}$$

From this, the potential surface of the control law is derived by the quadratic

$$\phi_f = \epsilon^T \epsilon$$

For certain environmental geometries, this control law has been proven to converge to a globally optimal solution [3]. In general, a controller can be derived from this law which converges to a force or moment solution which is locally robust with respect to a range of uncertainties and perturbations. At the same time, this control law only requires local contact information without the need for complete information about the object geometry.

Φ_k: Posture Control

The third control law addresses the kinematic configuration of a robot system. While not strictly required in the tradition of hybrid control, this is an important subtask for dextrous manipulators which contain redundant degrees of freedom. The objective here is to allow the robot to maintain favorable postures while it is engaged in interactions with the environment. The manipulability-based control law [10] used here combines joint limit information with the manipulability index to obtain a control law which optimizes the ability of the mechanism to produce forces and motions. As opposed to the first two control objectives which are concerned with the robot's interaction with the environment, this control law manages the manipulator's internal configuration.

From these control laws, controllers can be constructed which take the form $\Phi_i \frac{\sigma}{\underline{\tau}}$. In the most general terms, the input and output resources σ and τ are sensors and actuators associated with the fingers of the robot hand or abstractions derived from the stream of sensor data such as the location of the manipulation frame. Resolving individual degrees of freedom within a series kinematic chain is often not necessary. Figure 2 shows the set of controllers (Φ_p, Φ_f, Φ_k) and possible robot resources (fingers $0, 1, 2, 3$ and object parameters x, y, φ) used for the finger gaits presented in the next sections.

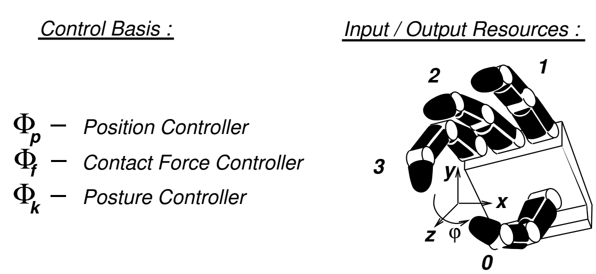

Figure 2: *Control Laws and System Resources for a Four-Fingered Robot Hand*

Concurrent combinations of bound controllers can be constructed using the "subject to" constraint. For example, $\Phi_k \frac{0,1,2,3}{\underline{\varphi}} \triangleleft \Phi_f \frac{0,1,2}{\underline{0}}$, uses a dominant contact configuration controller to attempt to achieve a stable grasp with fingers 0, 1, and 2 by moving finger 0. At the same time, the subordinate kinematic conditioning controller optimizes the configuration of all four fingers by modifying the orientation, φ, of the object. Complete manipulation strategies are constructed by composing these control elements into a context-dependent gait. Figures 3 and 7 present examples of the type of finite state supervisors employed in this work. The composite controller active at each point in time is indicated on the associated transition. In this approach, a finger gait is therefore not defined in terms of geometric contact locations but rather as a sequence of stability, kinematic conditioning, and navigational objectives. Grasp stability is established explicitly through the activation of corresponding controllers. Throughout execution, the stability of the object can thus be ensured within the limitations of the controller's capabilities and the sensory and motor resources of the manipulator.

3 Experiments

To demonstrate the applicability of the proposed approach and to illustrate the flexibility and robustness of the resulting control strategies, two different finger gaits for large-scale object manipulations are presented here. In addition, two different four-fingered robot platforms were used to illustrate the robustness and generality of finger gaits derived in the proposed control approach.

3.1 Object Reorientation

The first manipulation task chosen here is the rotation of an object within the grasp of a four-fingered dextrous hand. This task requires the construction of a control gait which coordinates the control of all degrees of freedom of the robot in order to achieve rotation while preserving a stable grasp. Figure 3 shows a finite state control policy for the object rotation task. In this gait, the object is initially grasped with fingers 0, 1, and 3 and the finger gait subsequently cycles through a sequence of stable, three-fingered grasps corresponding to the four equilibrium states in the finite state automaton. Each controller in this cycle achieves a specific three-fingered grasp by moving a single finger. This also leaves the system in a suboptimal kinematic configuration which, in turn, is compensated through changes in the orientation of the object by a concurrent, subordinate instance of the kinematic conditioning controller.

Rotating a Cylindrical Object In a first experiment the finger gait of Figure 3 was applied to the Utah/MIT dextrous hand and a cylindrical object. The corresponding finger gait and the resulting object rotation are shown in Figure 4.

These pictures show the Utah/MIT hand rotating a soda can in freespace. This illustrates the success of the control approach proposed here in achieving the desired reorientation while constantly maintaining a stable force-closure grasp.

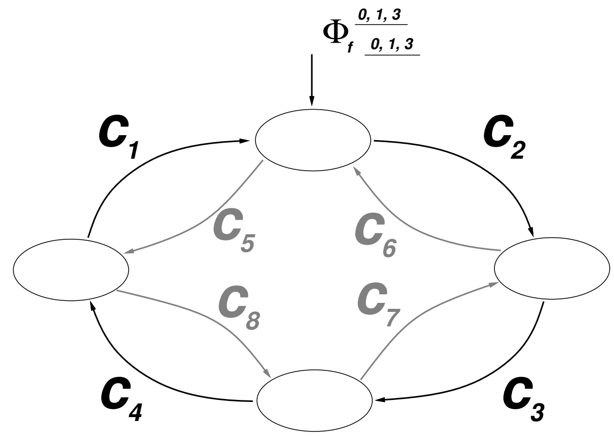

Counterclockwise:

$C_1 = \Phi_k{}^{\frac{0,1,2,3}{\varphi}} \triangleleft \Phi_f{}^{\frac{0,1,3}{1}}$

$C_2 = \Phi_k{}^{\frac{0,1,2,3}{\varphi}} \triangleleft \Phi_f{}^{\frac{0,1,2}{2}}$

$C_3 = \Phi_k{}^{\frac{0,1,2,3}{\varphi}} \triangleleft \Phi_f{}^{\frac{1,2,3}{3}}$

$C_4 = \Phi_k{}^{\frac{0,1,2,3}{\varphi}} \triangleleft \Phi_f{}^{\frac{0,2,3}{0}}$

Clockwise:

$C_5 = \Phi_k{}^{\frac{0,1,2,3}{\varphi}} \triangleleft \Phi_f{}^{\frac{1,2,3}{1}}$

$C_6 = \Phi_k{}^{\frac{0,1,2,3}{\varphi}} \triangleleft \Phi_f{}^{\frac{0,2,3}{2}}$

$C_7 = \Phi_k{}^{\frac{0,1,2,3}{\varphi}} \triangleleft \Phi_f{}^{\frac{0,1,3}{3}}$

$C_8 = \Phi_k{}^{\frac{0,1,2,3}{\varphi}} \triangleleft \Phi_f{}^{\frac{0,1,2}{0}}$

Figure 3: *Control Gait for Rotation Tasks. The Solid Arcs and Controllers Correspond to a Counterclockwise Rotations Gait While the Shaded Ones Represent a Clockwise Rotation*

Rotating a Pentagonal Prism To demonstrate the robustness of the control policy with respect to limited variations in the object geometry, Figure 5 shows a simulation of the Utah/MIT hand rotating a pentagonal prism. This image sequence shows the grasp configurations corresponding to the four equilibrium states of the finger gait and the corresponding control transitions.

It should be noted here that the contact configurations achieved on this object are not optimal grasps. Due to the workspace limitations of the robot some of the contact locations have to be placed close to the corners of the object. However, all grasps on the object establishes force closure and would thus ensure a stable grasp. Much of the flexibility and robustness of the gait results from the characteristics of the contact configuration controller which does not require complete information about the object geometry but only local information about contact locations and normals.

Changing the Manipulator Kinematics The same gait used previously on the Utah/MIT hand

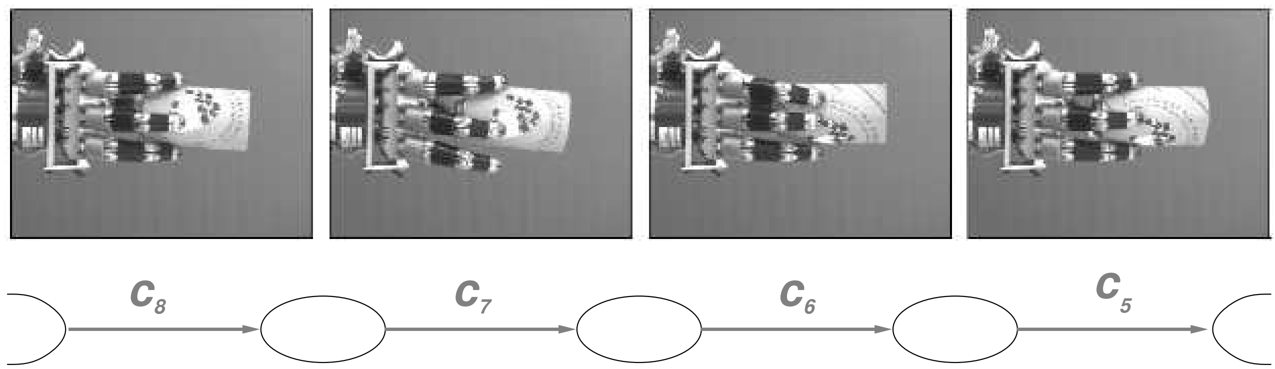

Figure 4: Utah/MIT hand Rotating a Soda Can (top), and the Corresponding Gait Transitions (bottom)

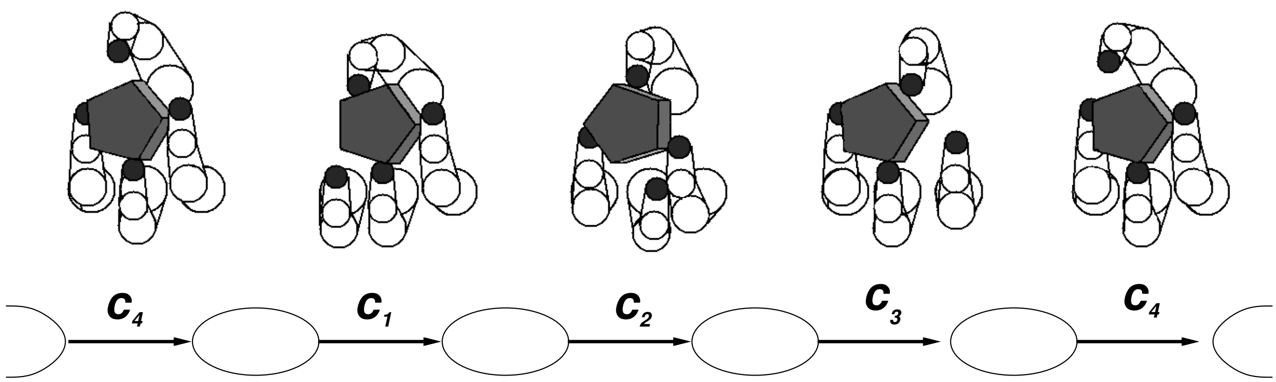

Figure 5: Rotating a Pentagonal Object (top), and the Corresponding Gait Transitions (bottom)

can also be transferred seamlessly to a different four-fingered manipulator. The only changes are the kinematic functions of the mechanism. The following shows the robot "Thing" rotating a globe using the same finger gait. The object here is larger in part because of the different kinematic structure of the mechanism. Figure 6 shows the resulting object rotation.

This further illustrates the robustness of finger gaits and control policies constructed in the control basis approach. The reactive control substrate and the DEDS supervisor here permit not only to compensate for local perturbations but also make it possible to transfer policies across a limited range of object and manipulator geometries.

3.2 Object Translation Gait

To achieve different overall task objectives, the control approach proposed here uses different composition policies over the same set of basis controllers. To address a new task it is therefore only necessary to modify the finite state supervisor. To demonstrate this, a second finger gait has been designed that translates the object within the hand while maintaining a stable grasp configuration. The location of the manipulation frame of the object, x, y, is computed in this gait as the centroid of the contact locations involved in the current grasp. Figure 7 shows the control gait for this object translation task.

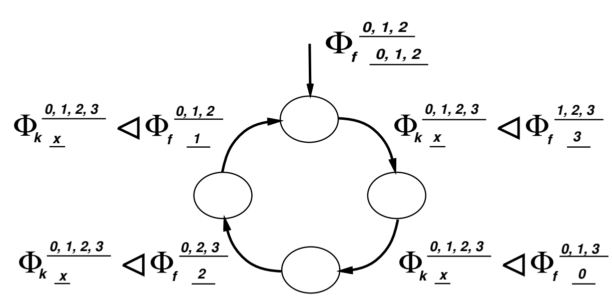

Figure 7: Control Gait for Shifting Task

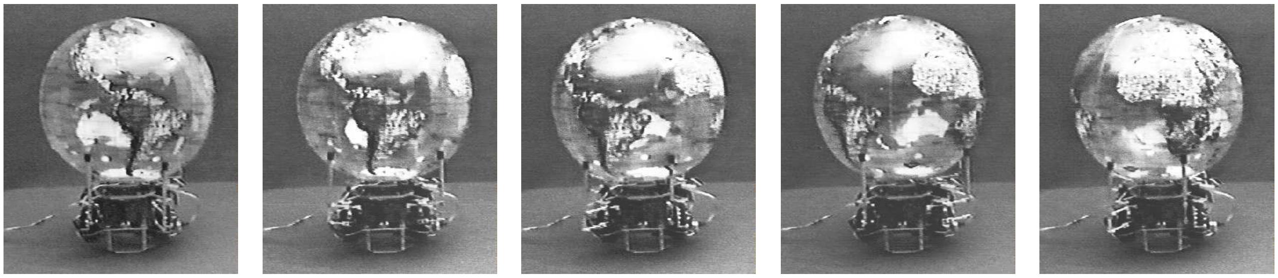

Figure 6: *Four-fingered Robot "Thing" Rotating a Beach Ball*

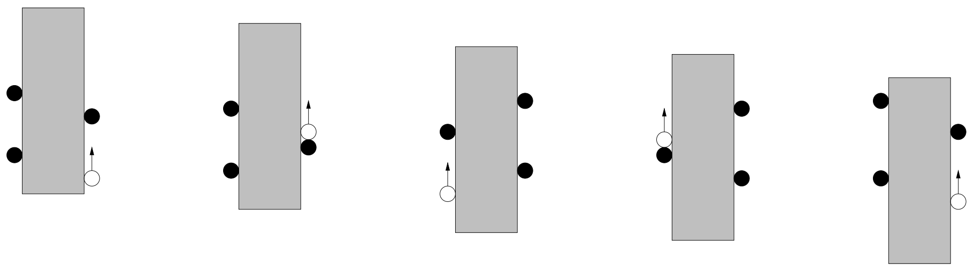

Figure 8: *Contact Configurations During Object Translation*

Similar to the rotation gait discussed previously, the object is initially grasped with three fingers and the finger gait subsequently cycles through a sequence of stable, three-fingered grasps. However, due to the different sequence of grasps, the object motion resulting from the kinematic conditioning of the manipulator is a translation through the fingers. Figure 8 shows the sequence of contact configurations achieved during the execution of one gait cycle.

This further illustrates the potential of the control basis approach for large-scale object manipulation by means of finger gaits. The same set of controllers is used here to accomplish a different manipulation task. Again, the reactive control basis computes contact locations on-line without the need for complete information about the object geometry and thus leads to a manipulation gait that is robust with respect to limited variations in object geometry and robot kinematics.

It should be noted here that the applicability and success of a given finger gait still depends on the workspace limitations of the robot. The translation gait in Figure 8, for example, could be executed successfully on the four-fingered hand "Thing" but would not achieve the desired task objective on the Utah/MIT dextrous hand. The reason is that the kinematic configuration of the Utah/MIT hand with its opposing thumb does not permit fingertip grasps with two fingers on either side of the object. This implies that the second grasp controller would not be able to establish its grasp objective. However, even on this case, the initial stable grasp would not be compromised. Rather, the finger gait would terminate with the initial grasp intact.

4 Beyond Finger Gaits

In this paper, the control basis approach has been applied to two different dextrous manipulation tasks. However, the careful choice of the control basis and the flexible composition mechanism make it possible to apply the same control basis to a wider range of tasks and robot platforms. In particular, its applicability is not limited to dextrous hands but can also be used to control other robot manipulators such as robot arms and walking robots. This has already been demonstrated in a number of additional manipulation tasks using a complete hand/arm system [9], as well as in locomotion tasks using a walking robot [15]. These experiments further illustrate the expressive power of the compact set of control laws within the control basis approach.

5 Conclusions

Object manipulation with a dextrous robot hand is a complex task due to the high number of degrees of freedom, the manipulator's workspace limitations, and the uncertainty present. Finger gaits are one way to address the workspace limitations for large-scale manipulations within a robot hand. This paper proposes an approach to designing robust and flexible manipulation gaits which construct contact placements and object motions on-line from a composition of control elements drawn from a carefully constructed basis of reusable feedback control laws with formal stability and convergence properties. To illustrate this approach, a control basis for manipulation tasks and two different finger gaits were designed. The resulting control policies were applied to different object geometries and manipulator kinematics and were able to robustly perform the task objectives.

Overall, the experiments suggest that the approach scales well to the complexity of manipulation tasks. Furthermore, the generic character of the control basis allows it not only to generalize across task domains but also to improve generalization across a variety of hardware platforms.

Acknowledgments

This work was supported in part by NSF IRI-9503687, IRI-9704530, and CDA-9703217.

References

[1] R. A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, March 1986.

[2] M. Cherif and K. K. Gupta. Planning for in-hand dextrous manipulation. In *Third Workshop on the Algorithmic Foundations of Robotics*, Houston, TX, March 1998.

[3] J. A. Coelho Jr. and R. A. Grupen. Effective multifingered grasp synthesis. Technical Report 94-12, CMPSCI Dept., Univ. of Mass., Amherst, February 1994.

[4] J. A. Coelho Jr. and R. A. Grupen. A control basis for learning multifingered grasps. *J. Robotic Sys.*, 14(7):545–557, October 1997.

[5] A. Cole, P. Hsu, and S. Sastry. Dynamic regrasping by coordinated control of sliding for a multifingered hand. In *Proceedings of the 1989 Conference on Robotics and Automation*, pages 781–786, Scottsdale, AZ, May 1989. IEEE.

[6] C. I. Connolly. Harmonic functions and collision probabilities. In *Proc. IEEE Int. Conf. Robot. Automat.*, San Diego, CA, May 1994. IEEE.

[7] C. I. Connolly and R. A. Grupen. The applications of harmonic functions to robotics. *J. Robotic Sys.*, 10(7):931–946, October 1993.

[8] R. Grupen. *General Purpose Grasping and Manipulation with Multifingered Robot Hands.* PhD thesis, University of Utah, Merrill Engineering Building, Salt Lake City, UT 84112, August 1988.

[9] R. A. Grupen, M. Huber, J. A. Coelho Jr., and K. Souccar. Distributed control representation for manipulation tasks. *IEEE Expert*, 10(2):9–14, April 1995.

[10] R. A. Grupen and K. Souccar. Manipulability-based spatial isotropy: A kinematic reflex. In *International Workshop on Mechatronical Computer Systems for Perception and Action*, pages 157–163, Halmstad, Sweden, June 1993. CCA, Halmstad University.

[11] L. Han and J. C. Trinkle. Dextrous manipulation by rolling and finger gaiting. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 730–735. IEEE, May 1998.

[12] I. Harmati, B. Lantos, and S. Payandeh. Extensions of nonlinear control concept to object manipulation with finger relocation. In *Proceedings of IFAC, SYROCO 2000*, 2000.

[13] J. W. Hong, G. Lafferriere, B. Mishra, and X. L. Tang. Fine manipulation with multifingered hand. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 1568–1573. IEEE, May 1990.

[14] M. Huber and R. A. Grupen. A hybrid discrete event dynamic systems approach to robot control. Technical Report 96-43, CMPSCI Dept., Univ. of Mass., Amherst, October 1996.

[15] M. Huber, W. S. MacDonald, and R. A. Grupen. A control basis for multilegged walking. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages Vol.4 2988–2993, Minneapolis, MN, April 1996. IEEE.

[16] S. Leveroni and K. Salisbury. Reorienting objects with a robot hand using grasp gaits. In *Int. Symp. Robotics Res.*, pages 2–15. Springer Verlag, 1995.

[17] Z. Li, P. Hsu, and S. Sastry. Grasping and coordinated manipulation by a multifingered robot hand. *Int. J. Robotics Res.*, 8(4):33–50, August 1989.

[18] P. Michelman and P. Allen. Task-directed, precise manipulation with a dextrous robot hand. In *Proceedings of the International Conference on Intelligent Autonomous Systems*, pages 729–738. IOS Press, February 1993.

[19] T. Omata and M. A. Farooqi. Regrasps by multifingered hand based on primitives. In *Proc. IEEE Int. Conf. Robot. Automat.*, pages 2774–2780. IEEE, May 1996.

[20] S. Payandeh. A model for controlled planning of object manipulation using multiple agents: An overview. In *Proceedings of ASME Annual Meetings: Symposium on Intelligent Control*, pages 759–801, 1998.

[21] M. Raibert. *Legged Robots that Balance.* MIT Press, Cambridge, MA, 1986.

[22] M. H. Raibert and J. J. Craig. Hybrid position/force control of manipulators. *Journal of Dynamic Systems, Measurements, and Control*, 102:127–133, June 1981.

[23] P. J. Ramadge and W. M. Wonham. The control of discrete event systems. *Proceedings of the IEEE*, 77(1):81–97, January 1989.

[24] T. Yoshikawa. *Foundations of Robotics : Analysis and Control.* MIT Press, Cambridge, MA, 1990.