January 1994

An Approach to Autonomous Robot Assembly

Daniel E. Koditschek

University of Pennsylvania, kod@seas.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/ease_papers

Recommended Citation

Daniel E. Koditschek, "An Approach to Autonomous Robot Assembly", . January 1994.

Reprinted from *Robotica*, Volume 12, February 2004, pages 137-155.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ease_papers/350
For more information, please contact repository@pobox.upenn.edu.

An Approach to Autonomous Robot Assembly

Abstract

Assembly problems require that a robot with fewer actuated degrees of freedom manipulate an environment containing a greater number of unactuated degrees of freedom. From the perspective of control theory, these problems hold considerable interest because they are characterized by the presence of non-holonomic constraints that preclude the possibility of feedback stabilization. In this sense they necessitate the introduction of a hierarchical controller. This paper explores these issues in the simple instance when all of the pieces to be assembled are constrained to lie on a line. A hierarchical controller is devised for this problem and is shown to be correct: the closed loop system achieves any desired final assembly from all initial configurations that lie in its connected component in configuration space; the generated sequence of motions never causes collisions between two pieces. Further examination of this approach interprets the controller's mediation of conflicting subgoals as promoting an M-player game amongst the pieces to be assembled.

Keywords

autonomous assembly, robots, hierarchical controller, closed loop system

Comments

Reprinted from *Robotica*, Volume 12, February 2004, pages 137-155.

An approach to autonomous robot assembly

Daniel E. Koditschek*

Artificial Intelligence Laboratory, EECS Department, College of Engineering, University of Michigan, Ann Arbor, Michigan 48109 (USA)

SUMMARY

Assembly problems require that a robot with fewer actuated degrees of freedom manipulate an environment containing a greater number of unactuated degrees of freedom. From the perspective of control theory, these problems hold considerable interest because they are characterized by the presence of non-holonomic constraints that preclude the possibility of feedback stabilization. In this sense they necessitate the introduction of a hierarchical controller. This paper explores these issues in the simple instance when all of the pieces to be assembled are constrained to lie on a line. A hierarchical controller is devised for this problem and is shown to be correct: the closed loop system achieves any desired final assembly from all initial configurations that lie in its connected component in configuration space; the generated sequence of motions never causes collisions between two pieces. Further examination of this approach interprets the controller's mediation of conflicting subgoals as promoting an M-player game amongst the pieces to be assembled.

KEYWORDS: Autonomous assembly; Robots; Hierarchical controller; Closed loop system.

1. INTRODUCTION

A general problem of widely acknowledged importance in robotics concerns the coordination of multiple degrees of freedom in the face of an environment possessed of fewer. Whether formulated as the "redundant manipulator" problem, the "cooperating robots" problem, or the "grasping" problem, many researchers have attempted to deploy underconstrained and actuated joints to gain better performance than could have otherwise obtained.¹ Consider, however, how commonly the opposite situation prevails. The paradigm of the monkey and banana quickly comes to mind. A hungry monkey stares at a ripe banana hanging over its head out of reach. On the ground at its feet lies a scattered assortment of sturdy boxes. The monkey quickly realizes it can stack up the boxes, climb the pile, and yank the banana off the limb. The banana and boxes each possess (at least) six degrees of unactuated freedom. The actuated joints that the monkey can bring to bear on the problem are far fewer but it engages its environment in a planned sequence of manipulations that achieves the

desired goal state: banana-in-mouth. This is an example of an assembly problem.

This paper explores robot assembly in a much more prosaic setting: we command a force actuated "bead" robot on a wire to move a collection of unactuated "beads" on a parallel wire into some desired final configuration from arbitrary initial conditions. This problem is simple enough to admit a complete and provably correct solution but, upon a more than cursory examination, not nearly so trivial as might initially seem when we require the task to be planned and executed by a feedback controller. One degree of freedom assembly tasks, although uninteresting in their own right, seem to incorporate to a significant extent the same features that make the more general problems truly confounding – a combination of holonomic and nonholonomic constraints. This paper is offered, then, in the spirit of an extended simple example of a complex general problem. The solution and methods of analysis it introduces hold the promise of generalization across the horizon of the general case where the pieces have higher degrees of freedom,² where the pieces are truly dynamical (rather than governed by generalized damper dynamics as in this paper,³ and where the robot is itself located within the same cluttered workspace as the pieces.⁴

Roughly speaking, there are two ideas at work. First, a refined notion of artificial potential functions⁵ that we have termed *navigation functions*⁶ is employed to encode the subgoals characterizing the participation of each piece in a completed assembly. A standard method of deriving stable feedback controllers for mechanical systems from potential fields⁷ is then used to associate a correct closed loop with each of the encoded subgoals. Since, as will be demonstrated, no single closed loop can result in a completed assembly, one requires some "higher level" organizing principle to switch between the alternative closed loops. This second idea, autonomous scheduling of conflicting subgoals, is motivated in large measure by independent work in robot juggling that we have pursued over the last few years.^{8,9} A straightforward encoding of "urgency" is used to select the most deserving lower level alternative. The discrete time closed loop dynamics presented by the higher level switching process can be shown to converge by appeal to standard ideas of nonlinear programming. In summary, the present approach to assembly might be characterized as "juggling a navigation function."

The paper is organized as follows. This introduction concludes with a more detailed overview of the problem

* This work was supported in part by the National Science Foundation under grant DMC-8505160, and, in part by a Presidential Young Investigator Award.

setting, related literature, and present contribution. Section 2 examines carefully and presents a feedback controller that correctly solves the "unit" assembly problem wherein there is only one bead for the robot to place. This solution is successively generalized to incorporate two and general multi-body one degree of freedom assemblies in Section 3. The paper concludes with a brief discussion concerning the larger prospects for generalization of these ideas. An appendix presents certain constructive techniques from our previous work, along with the computational details required to make them fit here.

1.1 Assembly: manipulating many using few degrees of freedom

Loosely speaking, say that a robotic task involves an assembly problem if the environment to be manipulated possesses more degrees of (unactuated) freedom than are available to the (actuated) robot system, and the specified goal state is to be achieved starting from arbitrary initial configurations (that is, all unactuated degrees of freedom must be exercised, in general, to complete that task). The unavailability of actuated degrees of freedom might result from limitations inherent in the robot's design (e.g., the PUMA has only six joints and the widget has twenty parts) or as a function of natural constraints imposed by the environment (e.g., the monkey's twenty degree of freedom hand has no bearing upon the banana's six degree of freedom state unless there is contact). A successful assembly plan must develop a sequence of manipulations none of whose single steps can achieve the goal yet each of whose concluding states brings the environment to a more favorable situation than the prior.

Surely, any reasonably interesting task to be carried out in the real (unstructured) world by a solitary robot will have the character of an assembly problem. For example, this paradigm underlies the warehouseman's problem:¹⁰ in a large hall filled to the ceiling with storage crates lies (in the far corner under a pile of heavy cartons) a box of back issues of *Robotica*; the task is to retrieve a copy of the issue. It appears reasonable, moreover, to represent the excavation problem – e.g., "robot bulldozer, go clear out the following dimensioned cavity so that the foundation robot can pour in the footings" – as a version of the warehouseman's problem that incorporates some independent dynamics in the workspace. Undersea or outerspace robot assembly operations will add still stronger a dynamical character to the environment. But even in the most structured factory setting, it is hard to imagine that "design for manufacturability" will obviate the need to assemble widgets with more parts than agents to manipulate them.

There is, of course, a large and growing assembly literature. Perhaps most notable has been that generated by the "handey" project of Lozano-Perez and colleagues,¹¹ – arguably the first integrated system capable of perceiving, planning, and controlling a general class of objects within a reasonably complex unstructured environment in support of arbitrary user

specified pick and place operations. Subsequent work Kak and colleagues has focussed experimental¹² and analytical¹³ attention upon integrating sensing and manipulation strategies in a similarly unstructured environment. In contrast to these very important early advances in managing the full spectrum of problems that confront robots in realistic settings, the notion of assembly addressed in the present paper includes a much narrower range of issues having to do with the planning and implementation of motions. For example, de Mello and Sanderson have recently reported progress in developing an effective representation¹⁴ and computational apparatus¹⁵ for enumerating all correct binary manipulation sequences toward a specified composite object. This work represents a substantial contribution to the assembly reasoning literature. In comparison, this paper addresses the problem of generating merely one correct sequence for each initial arrangement of parts. However, here, by "correct sequence" is meant an algorithm for generating an explicit profile of actuator torques and forces that result in physical robot and object motion. Thus, within the more focussed arena of assembly as a branch of motion planning, the work reported here might be said to bridge the ground between reasoning and control.

A second literature that has strong bearing on the present work concerns the control of nonholonomically constrained mechanical systems. Fine tutorials describing how such considerations arise in problems of mobile robotics,¹⁶ fingered manipulation,¹⁷ or robotic spacecraft¹⁸ are readily available. That such issues arise in the present setting is implicit in the work of Dorst¹⁹ and Laumond⁴ and their colleagues. Techniques for generating dynamical plans – open loop controllers – in the face of nonholonomic constraints generally appeal to an optimality criterion in order to achieve a computationally well posed problem and may be implemented using established numerical procedures.¹⁷ In contrast, the construction of feedback laws – closed loop controllers – for this class of nonlinear dynamical systems is in its infancy. While this paper derives great advantage from certain general results in the surrounding nonlinear stabilizability literature, there are only a few other authors who have attempted the actual construction of feedback controllers for nonholonomically constrained mechanical systems.^{20,21} As it turns out, the non-analytical nature of the nonholonomic constraints arising in assembly problems further distinguishes the present work from any previously considered to the best of the author's knowledge.

1.2 Contributions of this paper

This paper poses the multi-body one degree of freedom assembly problem. Because of the insistence upon feedback controllers, the solution entails a necessary appeal to a hierarchical controller. The proper choice from a family of low-level force policies by a high-level scheduling algorithm results in convergence to the goal state.

1.2.1 Statement of the problem. The problem setting (see Figure 1) presents a collection of rigid bodies (let us say M of them), each constrained to translate along the same horizontal axis. Each body has a specified goal location. This ensemble of specified goals comprises the completed assembly. None of the bodies can move independently. No two bodies can overlap. A body's motion in response to external forces is modeled as a "generalized damper:" all motion ceases when the external force is removed. A nearby parallel axis constrains a single rigid body robot. The robot has a desired nest position along the axis of motion. It is actuated by a source of bounded but instantaneous force and obeys Newton's laws of motion. It has a gripper as well which, when closed upon a body, is capable of exerting forces upon it. Most critically, the robot can exert no force on any body not proximal to it.

The problem is now as follows. We are required to find a feedback controller, that is, a policy for asserting bounded actuator forces and gripper positions as a function of the robot's and bodies' positions and velocities, which causes the robot to eventually reach its nest position after first bringing the collection of bodies into the desired assembly. The robot is not permitted to let any two bodies touch each other (although an initial configuration where some are touching is legal) during the manipulations.* The robot must end in its nest configuration. The bodies may be left in a disassembled state if and only if reaching the fully assembled configuration would entail passing one through the other.

1.2.2 Motivation: planning via feedback. Why closed loop controllers? In a recent paper,²² the author reviews the principles underlying a program of robotics research that seeks to develop autonomous planning and control procedures by recourse to dynamical systems theory. Since physical machines (whether operating in a dynamical or a quasi-state environment) are ultimately force or torque controlled dynamical systems, the specification of input torques must result in certain classes of vector fields. In this light, it makes sense to specify plans in the form of appropriately constructed sensor based feedback controllers whenever possible. Such specifications make explicit the resulting (closed loop) dynamical system, and afford the application of well developed mathematical analysis when attempting proofs of correctness. Moreover, feedback controllers, unlike open loop plans, are designed to work over large classes of initial configurations (tolerance to state uncertainty) and often succeed even when the underlying dynamics are imperfectly modeled (tolerance to parametric uncertainty).²³ Further, this approach to planning encourages the design of "canonical" procedures for

"model" problems which may then be instantiated in particular settings by a change of coordinates.^{24,25} Finally, if the robot's execution of a task in a specified environment may be represented as a dynamical system on an appropriate space, and if the criterion of success is the achievement of some distinguished goal set in that space, one is in a position to assess the "autonomy" of the resulting behavior with respect to standard ideas from dynamical systems theory. Namely, if autonomy connotes an ability to contend with the full spectrum of logically possible circumstances that arise in completion of a task, systems whose goal states are globally attracting represent autonomous behavior.²²

1.2.3 Nature of the solution. The proposed controller begins with a navigation function⁶ on each of the disjoint connected components (corresponding to the possible orderings of the M bodies) of the M -dimensional configuration space. If the pieces were "animated" – that is, if each body were equipped with its own individual actuator – then assigning a force law to each degree of freedom defined by the corresponding entry in the gradient vector field induced by the navigation function would result in the bodies converging toward the desired final assembly with no collisions along the way. This animated choreography would cease short of the desired final assembly if and only if the goal lay in a component of configuration space disjoint from the initial conditions. The particular navigation function takes a form presented in Appendix A below which also presents the bulk of the computational details required to prove that it is indeed a navigation function.

Since the pieces are not animated, at most only one may move at any time. The controller operates in principle by moving each piece down the corresponding component of the navigation function gradient evaluated with all the other variables held fixed. The piece is halted at a relative minimum and the next piece to move is chosen by virtue of having the largest component of the navigation function gradient at the point in question. The global convergence of this component-wise descent procedure is established in Section 3. In practice, the robot is "summoned" to approach and grasp the designated next piece. Once grasped, the piece is controlled toward the intermediate goal point by a nonlinear version of proportional-derivative control. This intermediate maneuver is described and its correctness demonstrated in the context of a single body assembly in Section 2.

1.2.4 Implications: hierarchical feedback control. How can such a seemingly trivial problem give rise to such a welter of symbols and formulae as are to be found in the next two sections? Despite their virtues, feedback controllers have inherent limitations. There is an important distinction between attraction – convergence to a limit set – and stability – informally, the ability to effectively resist small but arbitrary disturbances away from the limit set. Conventionally, feedback is designed to achieve asymptotic stability – the coincidence of both

* The presumption here is that a desired assembly with contact may be approximated to within any required degree of sensory resolution by a close but slightly separated configuration. In point of fact, in order to obtain convergence in finite time we choose in Section 3 the expedient (but not necessary) course of relaxing the termination condition in favor of an " ϵ -close" realization.

properties – with respect to a specified goal set. In the present case this is not possible and the reasoning in the sequel is complicated accordingly.

It is intuitively clear that one robot, since it cannot possibly be in two places at the same time, cannot simultaneously defeat arbitrary disturbances acting on multiple distributed objects. That this introduces real planning difficulties in a dynamical environment (e.g. undersea or in space) is self-evident. It is perhaps less clear but will be shown below, that this precludes the possibility of achieving asymptotic stability with respect to isolated goal points. Thus, even within the static generalized damper model adopted for ease of exposition in this paper, the absence of stability precludes the possibility of making monotone “progress” (as measured by any scalar means with respect to arbitrary initial conditions) toward the goal.

Such a situation necessitates the introduction of a hierarchy of control action. One can at best introduce a family of feedback laws that stabilize with respect to some corresponding family of subgoals. If properly conceived, the intersection of these “subgoals” corresponds to the desired goal (e.g. each piece of an assembly knows its own desired destination and regards the other pieces as merely so many obstacles). The subgoals will at times conflict (e.g. piece 1 might need to move away from its destination in order to give piece 2 enough room to get out of the way of piece 1’s progress toward that destination)* The family members must be sequentially deployed online according to the judgement of a higher level switching algorithm in such a fashion as to cause convergence. The basis for this judgement must be some “coarsened” measure of progress.

In the present context, the subgoals are encoded by the coordinate slice of the navigation function corresponding to each piece. The mass of formal statements in the sequel simply shows how traversing each slice in turn leads to progress in the higher level sense without compromising the boundedness or safety (obstacle avoidance) properties of the lower level at any step. The generality of this straightforward but detailed scheduling policy appears to be sufficient that it should work in the general case. Thus, given a general n -degree of freedom multibody static assembly problem, it seems to be the case that the only ingredient of the solution missing from this paper is a navigation function.** The generality of the present solution seems sufficient, moreover, that slight modifications seem to yield similar results for dynamical environments.³ Finally, if, as in ref. 4 the robot is more realistically placed within the workspace

with the pieces, then there is some hope that results from noncooperative game theory when added to this work may provide a solution to that class of problems as well.

2 THE ONE DEGREE OF FREEDOM UNIT ASSEMBLY PROBLEM

To reduce the problem to simplest terms, consider the situation depicted in Figure 1.

A rigid body of length ρ is restricted to a single horizontal axis of motion, $b \in \mathcal{B} = \mathbb{R}$. A unit mass robot moves along a parallel horizontal axis, $r \in \mathcal{R} = \mathbb{R}$. The robot must start from an arbitrary location, relocate the body, and then return to a specified nest location, $n \in \mathcal{R}$. We require an autonomous feedback control strategy for the robot that will enable it to move toward the piece, “grab it”, and place it in the arbitrarily designated nest location, and proceed to its nest. Somewhat surprisingly this is not as easy as it sounds. The problem is entirely trivial if open loop strategies are permitted (a two line example is offered below in Paragraph 2.1.2.a). If motivated by the discussion in the introduction, a closed loop strategy is desired, then the situation changes dramatically: now, rather than being trivial, the problem is unsolvable in the traditional sense.

In Section 2.1 we will show that the nature of the nonholonomic constraints that characterize this problem preclude feedback stabilization. That is to say, if a closed loop controller results in convergence to a pair of nest-goal locations that are not nearby each other then that configuration cannot comprise a stable equilibrium state of the closed loop system. The impossibility of asymptotic stability has two implications – one practical and one theoretical. Speaking practically, the former result simply confirms the obvious fact that small perturbations of the object’s location will require extensive efforts from the actuators in order to be right if the object is not within reach. Speaking theoretically, this result precludes the possibility of even using Lyapunov methods directly as a guide to constructing feedback laws: no controller whose closed loop system admits a Lyapunov function can result in convergence to an arbitrary (n, d) pair; no controller that does result in convergence can admit a Lyapunov function. This is particularly unfortunate since Lyapunov methods are one of the practicable tools available for assessing the stability properties of nonlinear dynamic systems.

Fig. 1. A one degree of freedom unit assembly problem.

* But, note, unfortunately, it is a central trivializing feature of the one degree of freedom problem engaged here that any conflict that can be resolved (i.e., the goal is reachable from the present configuration) involves at least one of the pieces making direct progress toward its goal. This obviously need not be the case in higher degrees of freedom,² and is arguably one way of thinking about why the navigation function that arise will not be convex as they are in the present paper.

**One is always guaranteed to exist.⁶ Their construction for particular classes of configuration spaces is in progress.^{2,24}

In Section 2.2 we take advantage of the present problem's utter simplicity to develop a successful feedback controller. When the robot is not required to nest then the formal obstruction to feedback stabilization is removed, and a modified proportional-derivative controller brings the robot to the body and the robot-body pair to the goal. When the robot is not required to manipulate the body then a standard proportional-derivative controller brings the robot to its nest. Clearly, the successive application of these two controllers results in the completed assembly. Moreover, just as clearly, the conditions under which the switch from the first to the second should occur can be expressed entirely as a function of state: apply the first unless the body is at its goal position. This defines a two level feedback controller whose generalization in the next section constitutes the central contribution of the paper.

Evidently, there is a technical problem concerning time of convergence here. The limit set of a smooth vector field is never reached in finite time, thus the switching scheme proposed above requires "two infinities" of time to converge as matters presently stand. Two remedies come to mind. First, one might adjust the low level algorithms to give finite time convergence. For example, Kawski is developing a systematic theory of Hölder Continuous feedback controllers (they are in fact smooth away from the limit sets) whose resulting closed loop systems satisfy Lyapunov functions with fractional powers.²⁶ Such systems have unique solutions forward in time even though their trajectories arrive at the attracting states in finite time. Second, as a practical matter, one might relax the switching criterion and settle for " ϵ -approximate" convergence to within some small neighborhood of the nominal switching set. Since it can be simply shown that this results in convergence to within a small neighborhood of the desired final assembly we adopt the latter remedy at present.

2.1. Assembly: another source of nonholonomic constraints

For purposes of the present investigation into the qualitative control properties of the situation depicted in Figure 1 it seems preferable to adopt as unconstrained a model of contact as possible. The coupling force, c , between the body and the robot will be modeled as some function of their relative position and the "gripper tightness," considered as a command input and denoted u_2 . The least we require of $c(r-b, u_2)$, is that it satisfy the following assumptions:

Coupling: $c(0, 1) \neq 0$

Releasing: $c(x, 0) = 0$ for all $x \in \mathbb{R}$

Proximity: $|x| \geq \rho$ implies $c(x, u) = 0$, for all $u \in \mathbb{R}$.

Roughly speaking, these conditions assure that manipulation is possible (**Coupling**) if and only if the robot is sufficiently close to the body (**Proximity**) and the gripper is engaged (**Releasing**). To avoid any technical difficulties with existence or uniqueness of solutions, it is expedient

to let the coupling law be smooth, $c \in C^\infty[\mathbb{R}^2, \mathbb{R}]$. The only regularization that cannot be introduced without destroying the essence of the problem at hand is to let c be analytic. This would violate the **Proximity** assumption, allowing the robot to "beckon" the object from arbitrarily far away (with however small a force). In the next section, when attention turns to the problem of actually constructing feedback controllers, it will prove useful to adopt a more concrete instance of this large class of possible coupling rules.

2.1.1 A dynamical model. In order to write down the control system this problem presents we require a dynamical model of the robot and its environment. Assume that the robot is controlled by a force we are free to command, u_1 , and its motion is governed by Newton's laws, $\ddot{r} = u_1$. In contrast, for the sake of simplicity, assume the body is governed instead by "generalized damper" dynamics²⁷ as is common in the assembly planning literature.²⁸ That is, assume sufficiently large friction forces are present that externally applied forces applied to this point mass will result in proportional changes in velocity. In particular, zero applied force is associated with zero velocity. It is worth emphasizing that these trivial dynamics represent a convenient vehicle for the present exposition – all the results obtained in Section 2.1.2 below continue, to hold under analogously loose assumptions when the body is posited to have Newtonian dynamics.³

Denoting by $q = [r, b]^T \in \mathcal{Q} = \mathbb{R} \times \mathbb{B}$ the configuration space vector, our control system now takes the form of a dynamical model along with an algebraic constraint,

$$M(q)\ddot{q} = Bu_1$$

$$J(q, u_2)\dot{q} = 0.$$

This is a common structure that characterizes many problems in physics and engineering,²⁹ and arises as well in important aspects of robotic manipulation as discussed in Section 1.1. Here, M is the constant (diagonal) mass matrix, $B = [1, 0]^T$ denotes the fact that all controlled forces act upon the body only through the robot, and J expresses the operating constraints:

$$J(q, u_2) = [c(r-b, u_2), -1]$$

i.e., the puck can be moved by the robot if and only if they are "touching" each other and the gripper is engaged.

The general structure of this model conforms to the well known but still poorly understood class of mechanical systems with *nonholonomic constraints*. Such systems impose constraints upon the state variables which are non-integrable, that is, the constraints can only be expressed in terms of the velocity variables of the problem. To best appreciate what this means intuitively, we follow Bloch and McClamroch,²¹ and consider the "reduced" state space representation of this system

$$\dot{p} = f(p, u) \quad f(p, u) = \begin{bmatrix} u_1 \\ p_1 \\ c(p_2 - p_3, u_2)p_1 \end{bmatrix}. \quad (1)$$

where $p = (\dot{r}, r, b) \in \mathcal{P} = \mathbb{R} \times \mathbb{R} \times \mathbb{B}$, and $u \in \mathbb{R}^2$.

Fig. 2. The nonholonomic nature of system (1).

Figure 2 depicts the nature of the nonholonomy operating here. While the configuration space has two degrees of freedom, the infinitesimal motions—that is, the instantaneous velocity vectors available at any configuration—are constrained in a configuration dependent manner. Denoting the “contact” set by

$$\mathcal{C} \triangleq \{(r, b) \in \mathcal{Q} : |r - b| < \rho\},$$

we note that the system is limited to “horizontal” velocities (robot motion) away from contact, $q \notin \mathcal{C}$ and “diagonal” or “horizontal” velocities in contact $q \in \mathcal{C}$. This figure necessarily resembles closely those found in Laumond’s assembly work.⁴

2.1.2 Qualitative properties. We are now ready to present and justify the observation that even though assembly control systems are completely controllable, they fail to be stabilizable. This observation follows from Paragraph 2.1.2.a, Paragraph 2.1.2.b, and Paragraph 2.1.2.c, below. For convenience of presentation, it is first worth introducing some notation and recalling a few definitions from the control and dynamical systems literature.

Given a general system of the form (1), and a control input, $u(t)$, denote the solution at time t to the initial condition problem p_0 by $f'_u(p_0)$. Although more restricted definitions are sometimes profitable, it is traditional to allow as an admissible control input any Lebesgue integrable function of time that takes values in some bounded set of controls.³⁰ We simply assume here that the existence of solutions with respect to the notion of admissible controls has been assured so that $f'_u(\cdot)$ is always well defined. If for every pair of initial and final states, p_0, p_1 , there exists an admissible input u such that $p_1 = f'_u(p_0)$ for some time t then the system is *completely controllable*.³¹ A feedback controller is a rule for determining a control input as a function of the present state of the system: $u(t) = g(p(t))$. We will denote the “closed loop” trajectories resulting from application of this control input by $f'_g(\cdot)$. Consider now the steady state properties of some closed loop system, $f'_g(\cdot)$. Denote by $\mathcal{N}(p)$ an open neighborhood of some point, p . An

equilibrium state, p , is *attracting* if it is the limit set of some encompassing open neighborhood,

$$p = \lim_{t \rightarrow \infty} f'_g(\mathcal{N}(p)).$$

It is *stable* if every neighborhood, $\mathcal{N}(p)$ contains smaller $\mathcal{N}'(p) \subset \mathcal{N}(p)$ such that trajectories originating in $\mathcal{N}'$ remain within $\mathcal{N}$,

$$f'_g(\mathcal{N}') \subset \mathcal{N}.$$

It is *asymptotically stable* if it is attracting and stable. A control system is (continuously) *stabilizable* at p if there can be found a (continuous) feedback controller whose resulting closed loop system p is an asymptotically stable equilibrium state.

Paragraph 2.1.2.a indicates that system (1) is not completely controllable. Every completely controllable linear system is smoothly stabilizable, but the same is not at all true for nonlinear systems. Brockett launched what has become a thriving literature on nonlinear stabilizability properties by offering one of the early examples of a nicely controllable but not continuously stabilizable system.³² It is now understood that the nonholonomically constrained mechanical systems fail Brockett’s test,²¹ and thus fail to admit continuous stabilizing feedback controllers. This holds true in the present case according to Paragraph 2.1.2.b.

Yet the conclusion of Paragraph 2.1.2.c regarding the impossibility of any stabilizing feedback has not been reached in any classical nonholonomic setting to the best of the author’s knowledge. If for all initial conditions p_0 , the set of reachable states from p_0 in time s ,

$$R_s(p_0) \triangleq \bigcup_{0 \leq t \leq s} \bigcup_u f'_u(p_0),$$

contains p_0 in its interior for all $s > 0$, then the system is *small time locally controllable*. This property is known to be desirable from the point of view of stabilization theory. For example, Kawski has recently shown that a large class of planar systems which enjoy this property are “almost” smoothly stabilizable.²⁶ Bloch and McClamroch have shown that the analytic nonholonomically constrained systems arising in certain space-craft and common robotic manipulation problems^{21,33} satisfy this property. They have constructed discontinuous stabilizing feedback controllers in some instances. It is perhaps symptomatic of the non-classical nature of assembly constraints that according to Paragraph 2.1.2.a system (1) does not enjoy the small time local controllability property.

In summary, system (1), is completely controllable (Paragraph 2.1.2.a) but not continuously stabilizable (Paragraph 2.1.2.b) as is the case in the “classical” nonholonomically constrained mechanical systems which include all of the robotic control problems examined to date. In contrast to these, the non-analyticity of the nonholonomy exhibited by (1) results in a robot control problem which is not stabilizable in any sense (Paragraph 2.1.2.c).

2.1.2.a Controllability The one degree of freedom

assembly system (1) is completely controllable but not small time locally controllable. From any initial configuration, $p_0 = (\dot{r}_0, r_0, b_0)$, there exists a control policy $(u_1(t), u_2(t))$ that brings the system to the desired final configuration, $p_d = (0, n, g)$. Such an open loop policy might be rendered as follows. Open the gripper, and move the robot (which, having the dynamics of a standard double integrator, is completely controllable) to the zero velocity position $r = b_0$. Now, close the gripper and move the robot to the zero velocity position $r = d$ — the body must follow according to the coupling rule. Finally, open the gripper, and move the robot to its nest position, $r = n$, also at zero velocity.

This system fails, however, to be small time locally controllable. Since admissible inputs are bounded, $R_s(p_0)$ must lie in smaller and smaller balls around p_0 as s gets smaller. Thus, taking a non-contact configuration, $(\dot{r}_0, b_0) \notin \mathcal{C}$, we may assume that $R_s(\dot{r}_0, r_0, b_0) \cap \mathcal{C} = \emptyset$. But this implies $b = b_0$ for all $(\dot{r}, r, b) \in R_s(p_0)$. In other words, $R_s(p_0)$ has an empty interior.

2.1.2.b Stability The one degree of freedom unit assembly system (1) is not continuously stabilizable since, as an example, $(0, 0, \varepsilon)$ is not in the image of f for any $\varepsilon \neq 0$, and the system fails Brockett's test.^{32,34 *}

2.1.2.c Convergence implies instability If a non-contact goal $q_d \notin \mathcal{C}$ is an attracting equilibrium state, then it cannot be stable in the sense of Lyapunov.

If $(0, n, d) \in \mathcal{P}$ is an attracting equilibrium state then it comprises the limit set of all initial conditions in arbitrarily small neighborhoods, $(\dot{r}_0, r_0, b_0) \in \mathcal{N}(0, n, d)$. If $\mathcal{N}(0, n, d)$ is a sufficiently small neighborhood then it has an empty intersection with the contact set, $\mathcal{C}$. It follows that all trajectories from initial conditions in $\mathcal{N}(0, n, d)$ with $b_0 \neq d$ must leave that neighborhood.

2.2 Feedback controllers for unit assembly

To permit a discussion of concrete feedback controllers we now add the following assumptions to our list of attributes of c :

Passivity: $0 \leq c \leq 1$

Gripping: There exists a smooth function $c^+(x, c_0)$ such that $c_0 = c(x, c^+(x, c_0))$, for all $c_0 \in [0, 1]$, $x \in [0, \rho]$.

Note that consistency with **Proximity** above, requires $0 < \rho < \rho$, which is now assumed. Roughly speaking, these additional assumptions guarantee that the gripper cannot by itself accelerate the body to be going more quickly or in the opposite direction from the robot (**Passivity**) and that eventually after the robot has come into range of the body it has available the full scope of coupling strengths (**Gripping**). This latter assumption is really added for convenience to stay within the category

* This result seems to depend in some measure upon the dynamical model introduced above. For example, when the body's dynamics are Newtonian and the grasping model involves an exchange of forces between the robot and body then the system fails Brockett's test only away from the contact set, $\mathcal{C}$.³

of smooth vector fields so that all discontinuities are seen to arise from the control policy. One might more simply hypothesize a Boolean "glue" model of coupling whereby $c = 1$ if and only if $r - b < \rho$ and $u_2 = 1$ (e.g. as in [4]). However, the chosen prescription might be said to add a slightly greater touch of reality to the model of force interaction since it admits the possibility of the body slowly (or quickly) slipping away from too weak a robot's grasp as their common motion is opposed by the frictional forces of the environment. Note that the "glue" model can be recovered by letting ρ approach ρ .

As an example, let σ be a scalar function that smoothly interpolates the values

$$\sigma(x) = \begin{cases} 1/U_{\max} & |x| < \rho \\ 0 & |x| > \rho \end{cases}$$

where $U_{\max}$ represents the full scale of the gripper input range, $u_2 \in [0, U_{\max}]$. Then we might have

$$\begin{aligned} c(r - b, u_2) &= \sigma(r - b)u_2; \\ c^+(r - b, c_0) &= c_0/\sigma(r - b). \end{aligned}$$

2.2.1. Stabilization to a submanifold Suppose it is merely desired that the robot come to rest in an arbitrary location after placing the body at some goal point, d . None of the foregoing theoretical obstructions to stabilizability apply in principle in this case since the requirement is for convergence to a submanifold — a smooth continuum of points (that is in fact a linear subspace in this situation) — rather than an isolated equilibrium state. However, straightforward application of the task encoding methods outlined in Appendix A lead to a goal specification taking the form of an isolated point. Thus, as dictated by Paragraph 2.1.2.b, we are left with a discontinuous controller (note that there is no contradiction with Paragraph 2.1.2.c since the goal is in $\mathcal{C}$). In later sections of this paper that treat multi-body assembly problems the same encoding gives rise to a true continuum of goal states. Although we will continue to use the controller devised here in those situations, it is not clear whether some more effective encoding might result in a continuous feedback law.

Since there are no obstacles in the present setting, it is straightforward to devise a cost function that encodes this task. We simultaneously require $\dot{r}^2 = 0$ and $\varphi_d(b) \triangleq (b - d)^2 = 0$. According to the **Proximity** assumption, the latter will not be possible unless the system is brought into the contact region, $\mathcal{C}$, a condition that is obviously encoded by making $(r - b)^2$ small. Thus, the task at hand is readily encoded as the zero set of the scalar valued function,

$$v_d \triangleq \frac{1}{2}\dot{r}^2 + \frac{1}{2}\varphi_d(b) + \frac{1}{2}(r - b)^2. \quad (2)$$

Since there is no *a priori* obstruction to stabilization around the zero level set of v_d , we follow Bloch and McClamroch²¹ and attempt to find a feedback controller, g , with respect to those closed loop system, $f'_g(\cdot)$, the positive definite function v_d is a Lyapunov function. The time derivative of v_d along the motion of

the system (the inner product of the gradient of v_d with the vector field $f(1)$),

$$\dot{v}_d = Dv_d \cdot f = \dot{r}[u_1 + D\varphi_d c + (r-b)(1-c)],$$

can be made negative semi-definite by choosing

$$u_1 \triangleq -\dot{r} - D\varphi_d c - (r-b)(1-c), \quad (3)$$

since this results in

$$\dot{v}_d = -\dot{r}^2.$$

According to LaSalle's invariance principle,³⁵ the limit set of the resulting closed loop is contained within the largest positive invariant set lying in $\dot{v}_d = 0$ —that is, the plane $\dot{r} = 0$. But the closed loop vector field restricted to this set,

$$f|_{\dot{r}=0} = \begin{bmatrix} 0 \\ -D\varphi_d c - (r-b)(1-c) \end{bmatrix}$$

is only tangent when the second entry is zero. Thus,

$$D\varphi_d c(r-b, u_2) + (r-b)[1-c(r-b, u_2)] = 0 \quad (4)$$

specifies the candidate limit set of the control policy as presently defined. According to the **Proximity** assumption, the tangent condition (4) is never satisfied outside the contact region, $(r, b) \notin \mathcal{C}$. It remains to specify a gripper policy inside $\mathcal{C}$ that results in convergence of the robot to the goal, d .

This is readily accomplished, for example, by adopting the particular gripper rule, c_d , that acts as follows

$$u_2 = \begin{cases} 0 & |r-b| > \rho \\ c^+(r-b, c_d(r, b)) & |r-b| \leq \rho \end{cases} \quad (5)$$

$$c_d(r-b) \triangleq \frac{|r-b| [1 + (D\varphi_d)^2]}{|D\varphi_d| + |r-b| [1 + (D\varphi_d)^2]}.$$

According to the **Gripping** assumption one may substitute c_d for c in (4) under this policy to get

$$(1 + D\varphi_d^2) \frac{D\varphi_d}{|D\varphi_d|} = \frac{(r-b)}{|r-b|},$$

and this can only hold true when $D\varphi_d = 0$ in which case $(r, d) = (d, d)$. Thus, the controller causes the body to come to rest at the desired goal location. It may be noted that (5) is discontinuous at configurations where $|r-b| = \rho$.

We will denote the feedback controller defined by equations (3) and (5) by the symbol g_d in the sequel and summarize its effect upon the resulting closed loop system by reiterating that $(0, d, d)$ is an asymptotically stable equilibrium state whose domain of attraction is global,

$$\lim_{t \rightarrow \infty} f'_{g_d}(\mathcal{P}) = (0, d, d). \quad (6)$$

2.2.2 Convergence to a point. Now consider the full one degree of freedom assembly problem that requires the robot to relocate the body at the goal point, d , and then return to its "nest" position, n . According to Paragraph 2.1.2.b, no controller that accomplishes this task can be smooth. Moreover, according to Paragraph 2.1.2.c, there is no single Lyapunov function that can play a role here

analogous to that of (2) in the partial version of the assembly problem: Lyapunov theory only works for stable equilibrium states.

The obvious answer is to adopt a discontinuous two-stage control policy that switches from the manipulation task over to nesting task when the former is completed. This can be done perfectly well with feedback, since the conditions to be tested have to do with the geometry of state space.

The manipulation task is accomplished via the controller g_d , above. It is a consequence of standard linear systems theory that the controller

$$g_n(p) = \begin{bmatrix} u_1(p) \\ u_2(p) \end{bmatrix} \triangleq \begin{bmatrix} -\dot{r} - (r-n) \\ 0 \end{bmatrix}$$

yields a globally attracting stable submanifold,

$$\lim_{t \rightarrow \infty} f'_{g_n}(r, \dot{r}, b) = (0, n, b),$$

that corresponds to the desired nesting behavior.

Assume now that some "higher level" automaton is designed to schedule

$$u = \begin{cases} g_d & \varphi_d > \varepsilon \\ g_n & \varphi_d \leq \varepsilon \end{cases}$$

It is clear that this "hierarchical" controller results in approximate (" ε -close") convergence to the desired goal state from all initial conditions, and as $\varepsilon \rightarrow 0$, results in exact convergence albeit after longer and longer transients.* Note for future reference that since there is only one possible outcome (that is, the goal is always reachable), the switching logic tests the value of φ_d directly.

3 A HIERARCHICAL CONTROL SYSTEM FOR SCALAR ASSEMBLIES

In the previous section it was established that there is a contradiction between convergence and stability for assembly problems. However the force of this contradiction may have seemed somewhat tangential to the matter at hand since, in the case of a single body, one might be satisfied to let the robot "hover" around waiting to return to its nest (stability) rather than retire to its nest (convergence). The contradiction is sharpened by examination of the multiple body case. Here, the robot cannot possibly "hover" around all the bodies, and there is no stable means of bringing all of them to their destinations and keeping them there. This necessitates a hierarchical controller. The present section generalizes and expands upon these insights.

Section 3.1 treats the two body assembly problem which adds to the nonholonomic constraints explored previously the holonomic constraints arising from the prohibition against collisions. Two low-level feedback laws are developed that result in the robot bringing one or the other body to an intermediate subgoal location without collision. A notion of assembly *stage* is proposed

* Refer to the discussion in the beginning of this section concerning the alternatives to ε -close finite time convergence.

along with termination conditions. A proposed high level scheduler that operates within these stages alternates between the two competing lower level controllers and assigns gains to the selected alternative based upon the present configuration of the bodies. Its gain assignment strategy guarantees a uniform upper bound on input magnitudes over all possible stages. The scheduler brings the robot to its nest state only if both the low level laws would result in "negligible" progress.

Section 3.2 passes to the general case of M bodies and demonstrates that the high level scheduler developed previously imposes what amounts to a standard numerical descent technique with respect to the navigation function from which the lower level controllers are formed. Convergence is seen to result from well known principles of nonlinear programming. Thus, interaction of the two levels solves the scalar multi-body assembly problem. From arbitrary initial conditions of the robot and body, the algorithm causes the robot (i) to nest after having brought each body to its designated goal; and (ii) to stop short of the specified task if and only if the goal is unachievable.

In Section 3.3, the higher level scheduling decisions are re-interpreted in the context of game theory. The present scheme represents a purely cooperative game played by the pieces. More realistic assembly settings wherein the robot is itself included in the workspace occupied by the pieces will require a non-cooperative game formulation. A simple non-cooperative game is proposed for the two body problem. While graphical analysis shows that the resulting hierarchical control scheme works better than the previously proposed and provably correct version, the formal convergence properties of general non-cooperative M -player games in the absence of strict convexity assumptions remains unclear.

3.1 The one degree of freedom dual assembly problem

Consider the problem depicted in Figure 3 where there are two bodies, b_1, b_2 that the robot is required to place at two distinct desired goal locations, d_1, d_2 . The only new feature we introduce now is a stricture against collisions: whatever else the robot may do to the two bodies it may not attempt to make them overlap or even allow them to touch except, possibly, at zero relative velocity.

3.3.1 Dynamics and qualitative properties. The configuration space is now $\mathcal{Q} = \mathcal{R} \times \mathcal{B} \subset \mathbb{R}^3$, where the legal body configurations are comprised of two disjoint

Fig. 3. A one degree of freedom dual assembly problem.

components, $\mathcal{B} = \mathcal{B}_+ \sqcup \mathcal{B}_-$,

$$\mathcal{B}_+ \triangleq \{(b_1, b_2) \in \mathcal{B} : b_2 - b_1 \geq \rho_1 + \rho_2\};$$

$$\mathcal{B}_- \triangleq \{(b_1, b_2) \in \mathcal{B} : b_2 - b_1 \leq -(\rho_1 + \rho_2)\}.$$

Thus the present configuration space and has a more complicated topology than that in Section 2, since it has been disconnected by the " $\mathcal{B}$ -obstacle", $\mathcal{O} = \mathbb{R}^2 - \mathcal{B}_+ - \mathcal{B}_-$.

We adopt as the dynamical model the obvious extension of (1) from Section 2.1.1,

$$\dot{p} = f(p, u) \quad f(p, u) = \begin{bmatrix} u_1 \\ p_1 \\ c_1(p_2 - p_3, u_2)p_1 \\ c_2(p_2 - p_4, u_2)p_1 \end{bmatrix}, \quad (9)$$

where $p = (\dot{r}, r, b_1, b_2) \in \mathcal{P} = \mathbb{R} \times \mathcal{Q}$. Here, c_i smoothly satisfies the **Coupling**, **Releasing**, and **Passivity** assumptions introduced in Section 2 with respect to its two arguments, obeys the **Proximity** assumption with respect to ρ_i , and the **Gripping** assumption with respect to ρ_i . It is apparent that this problem inherits the same nonholonomic constraints depicted in Figure 2: velocities are limited to be parallel to the x -axis unless the configuration lies in the "free" contact set, $\mathcal{C}_2 \triangleq \mathcal{A} \cap \mathcal{C}$, where

$$\mathcal{C} = \mathcal{C}_1 \cup \mathcal{C}_2; \quad \mathcal{C}_i \triangleq \{(r, b_1, b_2) \in \mathbb{R}^3 : |r - b_i| < \rho_i\}$$

The only novelty lies in the addition of a holonomic constraint arising from the configuration space obstacle, $\mathcal{O}$.

3.1.1.a A controllable and unstabilizable system System (9) is completely controllable but an isolated attracting equilibrium state resulting from any feedback control policy is unstable.

The controllability argument of Paragraph 2.1.2.a is only slightly complicated by the presence of $\mathcal{O}$. It is intuitively obvious (and follows formally, for example, from the arguments in [6]) that a completely controllable mechanical system can be brought to any interior point of its configuration space from any source point in the same connected component without leaving that component (that is, with no collisions with any boundary). Supposing, for example, that $g \in \mathcal{B}_+$, as in Figure 3, then $b \in \mathcal{B}_+$ implies that either $b_2 > g_1$ or $b_1 < g_2$ (the situation in the Figure), thus one of the bodies is in the component of $\mathbb{R}$ connected to its goal point, and the arguments of Paragraph 2.1.2.a show that it may be brought there directly by an appropriate control that incurs no collision with the other. The remaining body must now also be in the component of $\mathbb{R}$ connected to its goal point, to which it is just as readily transferred. On the other hand, if $b \notin \mathcal{B}_+$ then there is no way of bringing both bodies to their goals without one passing through the other. Clearly, then, (9) is completely controllable on either of its two disjoint components.

Paragraph 2.1.2.b carries over directly in the present case. The more restrictive result of Paragraph 2.1.2.c holds more practical significance now. First, note that no

$p_d = (\dot{r}, n, g_1, g_2) \notin \mathbb{R} \times (\mathcal{C}_1 \cap \mathcal{C}_2)$ (that is, a phase characterized by some robot velocity at a nest state in a situation where the two bodies are not touching) can be both attracting the stable. For example, supposing for all $t > 0$, $f'_u(p_0) \notin \mathbb{R} \times \mathcal{C}_2$, we have $[0, 0, 0, 1]^T f = 0$, thus $f'_u(p_0) \rightarrow p_d$ implies $b_2 = g_2$. The argument now proceeds as in Paragraph 2.1.2.c. In contrast to that situation, notice next that $\mathcal{Q} \cap \mathcal{C}_1 \cap \mathcal{C}_2 = \emptyset$ – that is, there can be found no feedback controller that will stabilize any valid configuration.

3.1.2 Convergence to a point via a hierarchical controller. Consider the objective function

$$\varphi_d(b) = \frac{\|b - d\|^{2k}}{(b_1 - b_2)^2 - (\rho_1 + \rho_2)^2} \quad (10)$$

defined on the disjoint half planes, $\mathcal{B}_+$, $\mathcal{B}_-$. It will be demonstrated in greater generality in Paragraph 3.2.2.b that φ_d is convex (and strictly convex when $b \neq d$) whenever $k > 5/2$, which we now assume to be the case. As a consequence, φ_d has exactly two extrema – one minimum in each of the two disjoint components, $\mathcal{B}_+$, $\mathcal{B}_-$. If, as we now assume, $d \in \mathcal{B}_+$, then straightforward computation indicates that the other minimum in $\mathcal{B}_-$ is

$$\hat{d} = \frac{1}{4(k-1)} \left(\begin{bmatrix} 2k-3 & 2k-1 \\ 2k-1 & 2k-3 \end{bmatrix} d + \begin{bmatrix} -1 \\ 1 \end{bmatrix} \sqrt{(d_1 - d_2)^2 + 4k(k-1)(\rho_1 + \rho_2)^2} \right).$$

Now define the “squashed” version of this function to be

$$\tilde{\varphi}_d(b) \triangleq \frac{\varphi_d^{1/k}(b)}{[1 + \varphi_d(b)]^{1/k}},$$

a composition of φ_d with $\sigma_k(x)$ as discussed in Appendix A. Since σ_k is a smooth bijection between $(0, \infty)$ and $(0, 1)$ whose inverse is also smooth with respect to the intervals in question it is straightforward to show that $\tilde{\varphi}_d$ takes exactly the same extrema in $\mathcal{B}$ as does φ_d .⁶ The essential difference between the two functions is that $\tilde{\varphi}_d$ is defined on the boundary of the obstacle, $\partial\mathcal{O} = \{|b_1 - b_2| = \rho_1 + \rho_2\}$ where it attains its maximal value of unity. From these observations, it follows that $\tilde{\varphi}_d$ is a navigation function with respect to d on $\mathcal{B}_+$ and with respect to $\hat{d}$ on $\mathcal{B}_-$.

The two different “encodings” of the same task, φ_d , $\tilde{\varphi}_d$ will be used in the definition and analysis of the two different levels of control to be introduced below. The navigation function, $\tilde{\varphi}_d$ is an “admissible function”^{6,7} and can be turned to purposes of low level control directly as shown in Paragraph 3.1.2.a. Since it fails to be convex, however, we find it most simple to use the variant φ_d in the analysis of Paragraph 3.2.2.c.

Paragraph 3.1.2.a concerns the use of one or the other “coordinate slice” of φ_d as an encoding of where and how to move one body while the other remains fixed. The robot will approach the designated body, grasp it, and eventually reach a relative minimum (**Convergence**) without ever colliding with the other body (**Safety**). The largest force required along the way is bounded as a

function of initial conditions (**Boundedness**). While the form of the proposed feedback controller is analogous to that developed in Section 2.2.1, and the stability arguments are identical, the collision avoidance arguments represent a line of reasoning developed in some generality in ref. 7.*

3.1.2.a A bounded, safe and convergent “low level” controller. Denote by $\tilde{\varphi}_{d|b_2}$, the function defined by letting $\tilde{\varphi}_d$ vary over its first variable, b_1 , as b_2 is held fixed. Letting γ be a constant scalar gain, denote by

$$g_{d|b_2}(p) = \begin{cases} -\dot{r} - \gamma D_{b_1} \tilde{\varphi}_{d|b_2}(b_1) c_1 - (r - b_1)(1 - c_1) & \text{if } |r - b_1| > \rho_1 \\ 0 & \text{if } |r - b_1| \leq \rho_1 \\ c^T(r - b_1, c_d(r, b_1)) & \text{if } |r - b_1| \leq \rho_1 \end{cases} \quad (11)$$

the feedback controller resulting from a substitution of $\gamma \tilde{\varphi}_{d|b_2}$ for φ_d in equations (3) and (5). Then

Convergence:

$$\lim_{t \rightarrow \infty} f'_{g_{d|b_2}}(\mathcal{P}) = (0, \hat{d}_1(b_2), \hat{d}_1(b_2), b_2)$$

where

$$\hat{d}_1(b_2) \triangleq \arg \min_{(x, b_2) \in \mathcal{B}_-} \varphi_d(x, b_2) \cup \arg \min_{(x, b_2) \in \mathcal{B}_+} \varphi_d(x, b_2)$$

Safety: There exists a function, $v_{d|b_2}: \mathcal{P} \rightarrow \mathbb{R}^+$ such that $\gamma > v_{d|b_2}(p_0)$ implies

$$\tilde{\varphi}_d \circ f'_{g_{d|b_2}}(p_0) < 1 \quad \text{for all } t > 0.$$

Boundedness: $\|g_{d|b_2} \circ f'_{g_{d|b_2}}(p_0)\| \leq G_0 \circ v_{d|b_2}(p_0)$ for some continuous monotone scalar function G_0 .

Proof: Since $\mathcal{C}_1 \cap \mathcal{C}_2 \neq \mathcal{Q}$ and $r \notin \mathcal{C}_1$ implies $u_2 = 0$, that is, since the gripper is open when the robot is within reach of the second body, it is clear that b_2 is constant along the trajectories of the closed loop system. Substituting

$$v_{d|b_2}(\dot{r}, r, b_1, b_2) \triangleq \frac{1}{2} \dot{r}^2 + \frac{1}{2} \gamma \tilde{\varphi}_{d|b_2} + \frac{1}{2} (r - b_1)^2,$$

for v_d in (2) shows that $v_{d|b_2}$ is a Lyapunov function $f'_{g_{d|b_2}}(\cdot)$, hence,

$$\lim_{t \rightarrow \infty} f'_{g_{d|b_2}}(\mathcal{P}) \subset \dot{v}_{d|b_2}^{-1}(0)$$

$$= \{(0, x, x, b_2) \in \mathcal{P} : D_{b_1} \tilde{\varphi}_{d|b_2}(x) = 0\},$$

as argued in Section 2.2.1. Since φ_d is strictly convex $\mathcal{B} - d$, $\varphi_{d|b_2}$ must be convex as well, hence $\tilde{\varphi}_{d|b_2}$ has two extrema – one minimum in both of the disjoint infinite intervals on which it is defined – and the **Stability** conclusion follows.

Because $v_{d|b_2}$ is a Lyapunov function for the closed loop system and is at least as great as its constituent term, $\gamma \varphi_{d|b_2}$, we may conclude that

$$\begin{aligned} \tilde{\varphi}_d \circ f'_{g_{d|b_2}}(p_0) &= \tilde{\varphi}_{d|b_2} \circ f'_{g_{d|b_2}}(p_0) \\ &\leq \frac{v_{d|b_2}(p_0)}{\gamma}, \quad \text{for all } p_0 \in \mathcal{P}, \quad t > 0, \end{aligned}$$

from which the **Safety** conclusion follows.

* Readers who find this too abbreviated a discussion of arguments supporting Paragraph 3.1.2.a too cursory to follow might wish to consult that more leisurely reference.

s). While the analogous to the stability avoidance argument in some "low level" defined by as b_2 is held denote by

To show **Boundedness** note that φ_d is radially unbounded (see Paragraph 3.2.2.b), hence $v_{d|b_2}$ is as well. Thus, the trajectory remains in a compact set wherein $v_{d|b_2}(p) < v_{d|b_2}(p_0)$. Now $\|g_{d|b_2}\|$ is bounded above by a continuous function (formed, for instance by substituting 1 for c_1). The sup norm magnitude, G_0 , of this majorizing function grows with the extent of its domain, $v_{d|b_2}(p_0)$ as claimed. $\square$

Similarly, we may define a second Low Level Controller, $g_{d|b_1}$ in an analogous fashion to (11) and obtain the corresponding conclusions.

The question now at hand concerns the construction of a higher level scheduling algorithm to mediate between the two strategies. Given an initial condition, $p_0 \in \mathcal{P}$, define a *stage* as a new contact event: an event beginning at a time s around which there is some open interval, $s \in \mathcal{J}$ with the property

$$\begin{aligned} f'_g(p) &\in \mathbb{R} \times \mathcal{C} \quad \text{for all } t \in \mathcal{J}, \quad t \geq s; \\ f'_g(p) &\notin \mathbb{R} \times \mathcal{C} \quad \text{for all } t \in \mathcal{J}, \quad t \leq s. \end{aligned}$$

Define $s_0 \triangleq 0$ and let s_{k+1} denote the time of the next stage following that which began at $t = s_k$. Since $\mathcal{C}_1 \cap \mathcal{C}_2$ has no intersection with $\mathcal{Q}$, each stage, s_k is associated with a particular index, $l(k) \in \{1, 2\}$. A stage terminates at time $t \triangleq s'_k$ under the condition

$$0 < \kappa(k) \triangleq \frac{1}{2} \dot{r}^2(s'_k) < \kappa(k-1), \quad (14)$$

and

$$0 < \varphi(k) \triangleq \frac{1}{2} \varphi_d(b(s'_k)) < \varphi(k-1)$$

These definitions form the basis of a series of discrete events that take place on the contact set, $\mathcal{C}$, whose control may now be effected by appropriately scheduling the low level algorithms treated by Paragraph 3.1.2.a. There are two important questions to address in the analysis of the "high level" scheduler: First, does the "contact schedule" generate an admissible sequence of low level feedback controllers? Second, can it be shown that the resulting "contact schedule" achieves the desired final configuration, or terminates when that configuration is found to be unreachable? Paragraph 3.1.2.b addresses the first question. The second will be treated in Section 3.2.2.

Paragraph 3.1.2.b defines a straightforward extension of the trivial schedule (8). A high level scheme alternates back and forth between the low level controllers (11) corresponding to the two bodies at each stage until they form a configuration which is δ -close to an extremum of φ_d . At this point, the robot is directed to nest according to the low level controller (7). Since φ_d has only minima, one in each component of $\mathcal{B}$, it is intuitively clear that the second question has an affirmative answer. In any event, this will be demonstrated in greater generality in Section 3.2.2. Paragraph 3.1.2.b assures that the robot nests if and only if a global extremum has been reached (**Termination**), that no collisions between the bodies may occur at any stage (**Gain**), and that there will never be a time when arbitrarily large input magnitudes are required (**Growth**).

3.1.2.b An admissible contact schedule Consider the

scheduler that chooses a controller, g_k , at time s'_k according to the rule

$$g_0 = g_{d|b_1}$$

$$g_{k+1} = \begin{cases} g_n & \text{if } \|D_b \varphi_d(b(s'_k))\| \leq \delta \\ g_{d|b_{l(k)}} & \text{if } \|D_b \varphi_d(b(s'_k))\| > \delta \end{cases}$$

with gain, γ , chosen as

$$\gamma(k+1) \triangleq [\kappa(k) + R(k)][1 + \varphi(k)].$$

where $R(k)$ is chosen so that $\varphi_d(b) \leq \varphi(k)$ implies*

$$2 \|b\|^2 + \rho_1 + \rho_2 < R(k).$$

Then

Termination: $g_{k+1} \neq g_n$ implies $s'_{k+1} < \infty$

Gain: $\gamma(k+1) > v_{d|b_{l(k+1)}}(p(s'_k))$

Growth: $\|g_k\| \leq G_0$ for some $G_0 < \infty$

Proof: **Termination** follows directly from **Stability** of Paragraph 3.1.2.a according to condition (14). **Stability** also implies $q(s'_k) \in \mathcal{C}_{l(k)}$ which gives

$$\begin{aligned} (r(s'_k) - b_{l(k+1)}(s'_k))^2 &< (r(s'_k) - b_{l(k)}(s'_k))^2 \\ &\quad + (b_{l(k+1)}(s'_k) - b_{l(k)}(s'_k))^2 \\ &\leq \rho_1 + \rho_2 + 2 \|b\|^2 < R(k), \end{aligned}$$

and, in turn, **Gain**, since

$$\begin{aligned} &\frac{v_{d|b_{l(k+1)}}(p(s'_k))}{\gamma(k+1)} \\ &= \frac{\kappa(k) + \frac{1}{2}(r(s'_k) - b_{l(k+1)}(s'_k))^2 + \frac{1}{2}\gamma(k+1)\bar{\varphi}_d(s'_k)}{\gamma(k+1)} \\ &= \frac{\kappa(k) + \frac{1}{2}(r(s'_k) - b_{l(k+1)}(s'_k))^2}{[\kappa(k) + R(k)][1 + \varphi(k)]} + \frac{\varphi(k)}{1 + \varphi(k)} \\ &< 1. \end{aligned}$$

To reach the **Growth** conclusion note that since $\varphi(k+1) \leq \varphi(k)$ we have $R(k+1) \leq R(k)$, hence (recall from (14) that the robot's kinetic energy decreases at each termination),

$$\gamma(k+1) = \frac{\kappa(k+1) + R(k+1)}{1 - \varphi(k+1)} \leq \frac{\kappa(k) + R(k)}{1 - \varphi(k)} = \gamma(k),$$

thus

$$\begin{aligned} v_{d|b_{l(k+1)}}(p(s'_k)) &< \kappa(k) + \rho_1 + \rho_2 + \gamma(k+1)\varphi(k) \\ &/ (1 + \varphi(k)) < \kappa(0) + \rho_1 + \rho_2 + \gamma(0)\varphi(0)/(1 + \varphi(0)) \triangleq G_1 \end{aligned}$$

On the other hand, since $\gamma(k) > \rho_1 + \rho_2 > 0$ by construction, we have

$$v_{d|b_{l(k+1)}}(p(t)) > \frac{1}{2} \dot{r}^2(t) + (\rho_1 + \rho_2)\bar{\varphi}_d(b(t)).$$

Thus, for all time $t > 0$ we have $p(t)$ contained within the compact set wherein

$$\frac{1}{2} \dot{r}^2(t) + (\rho_1 + \rho_2)\bar{\varphi}_d(b(t)) < G_1$$

and

$$(r(t) - b_1(t))^2 + (r(t) - b_2(t))^2 < 2(R(0) + G_1).$$

But $\|g_k\|$ is bounded above by the continuous function

$$|\dot{r}| + \gamma(0) \|D\varphi_d\| + |r - b_1| + |r - b_2|,$$

* The existence of such a bound follows from the radial unboundedness of φ_d - see Paragraph 3.2.2.b.

which is itself bounded according to the **Boundedness** conclusion of Paragraph 3.1.2.a. $\square$

According to Paragraph 3.1.2.a **Safety**, $g_{d|b_i}$ preserves each connected component of $\mathcal{B}$. Thus, $f'_{g_{d|b_i}}(\cdot)$ chooses a distinguished member of the set $\hat{d}_i(b_i)$ corresponding to the component of $\mathcal{B}$ within which it began, and in this sense (that is, under the agreement to consider a particular component) we may abuse notation and consider $\hat{d}_i$ to be a function. In this sense the contact schedule of Paragraph 3.1.2.b represents an approximation to the difference equation

$$\begin{aligned} b_1(j+1) &= \hat{d}_1(b_2(j)) \\ b_2(j+1) &= \hat{d}_2 \circ \hat{d}_1(b_2(j)). \end{aligned}$$

Since φ_d , $\varphi_{d|b_2}$, $\varphi_{d|b_1}$ are all convex (and strictly convex when $b \neq d$) the dynamics of this difference equation specifies a contraction on both components of $\mathcal{B}$, and must converge to the unique fixed points of each component, for example by application of Li and Basar's contraction arguments³⁶ (Thm. 1). Paragraph 3.2.2.c will show that the approximate algorithm discussed here works as well.

3.2 General one degree of freedom assemblies

Suppose there are M bodies, each with a radius $\rho_i > 0$, located at $b = [b_1, \dots, b_M]^T \in \mathbb{R}^M$, and a desired goal state, $d = [d_1, \dots, d_M]^T \in \mathbb{R}^M$, toward which it is desired to move all the pieces. As before, we impose the restriction against two bodies colliding, and ask for a feedback controller to bring arbitrary initial conditions to the desired goal or terminate if the goal is not reachable.

3.2.1 Qualitative properties. Each of the $q \triangleq M(M-1)/2$ functions,

$$(b_i - b_j)^2 - (\rho_i + \rho_j)^2$$

vanishes twice – when body i touches body j on the left and when it touches body j on the right – defining two disconnected open half-spaces in $\mathbb{R}^M$ where the factor is positive. These correspond to the two different orderings of the two bodies in $\mathbb{R}$. There are 2^q different mutually disjoint subsets of $\mathbb{R}^M$ whereon each of these is positive that together comprise the free configuration space, $\mathcal{B}$. Since it is defined by the finite intersection of convex sets (in this case, the q different open half-spaces), $\mathcal{B}$ is open and convex. Only $q!$ of these convex components are physically meaningful in the sense that they represent a valid ordering of the M bodies. We now assume that d is located in one of these – the **domain**, $\mathcal{B}_+$. It follows that there are $q! - 1$ other possible orderings of the bodies from which the goal is not reachable.

The dynamical model is the obvious extension of (9) and it is clear that the arguments of Paragraph 3.1.1.a carry over directly. That is, there is no possibility of a closed loop algorithm that would stabilize any isolated desired goal.

3.2.2 Nonlinear programming. We now describe a relaxation method for numerical minimization that goes back at least as far as Gauss and has been particularly

attractive to those employing hand computation in the descent steps.³⁷ Namely, given a scalar valued cost function, φ that is to be minimized, one chooses at any point the most promising coordinate direction in which to move and performs a scalar search for the local minimum. We follow Luenberger in referring to this as the Gauss–Southwell Descent Algorithm.³⁸

Suppose, now, the existence of a cost function $\varphi: \mathbb{R}^M \rightarrow \mathbb{R}^+$. Adopt the notation

$$\bar{b}_i = [b_1, \dots, b_{i-1}, b_{i+1}, \dots, b_M]^T,$$

to denote the projection of b onto the $M-1$ dimensional subspace on which its i th component vanishes. Denote by

$$\varphi_{d|\bar{b}_i}(x) \triangleq \varphi_d(b_1, \dots, b_{i-1}, x, b_{i+1}, \dots, b_M).$$

the restriction of φ_d to the one dimensional affine subspaces of $\mathbb{R}^M$ that this $(M-1)$ dimensional subspace parametrizes – that is, all the lines parallel to its normal so that $D\varphi_{d|\bar{b}_i}(b_i) = D_{b_i}\varphi_d(b)$. Associated with this cost function is the low level controller $g_{d|\bar{b}_i}$ defined exactly in (11), and the closed loop trajectory $f'_{g_{d|\bar{b}_i}}(p)$. We are interested in generalizing the higher level scheduling algorithm. The only new problem is to decide which the best piece to move next after a previous piece has come to rest near a relative local minimum. Let $l(b)$ be an index valued function with the property

$$l(b) \in \arg \max_{i \leq M} |D_i \varphi(b)|,$$

that is, a function which picks out a component of b whose direction of descent with respect to φ is greatest.

Using this notation, define now a transformation of $\mathbb{R}^M$

$$\begin{aligned} T\left(\begin{bmatrix} b_{l(b)} \\ \bar{b}_{l(b)} \end{bmatrix}\right) &= \begin{bmatrix} \hat{d}_{l(b)}(\bar{b}_{l(b)}) \\ \bar{b}_{l(b)} \end{bmatrix}; \\ \hat{d}_{l(b)}(\bar{b}_{l(b)}) &\triangleq \arg \min \varphi_{d|\bar{b}_{l(b)}}(x) \end{aligned}$$

that leaves $M-1$ components unchanged, and brings the component whose gradient with respect to φ_d is largest to its minimum relative to the other fixed components. Again, disregarding the abuse of notation, it is understood that T leaves invariant the disjoint basins around the various minima of φ_d so that it is a well defined function on each component of $\mathcal{B}$.

While T represents the ideal higher level behavior of the scheduling algorithm, concrete implementation of the kind we are concerned with in the present paper as described in Paragraph 3.1.2.a can provide only an approximation of this scheme. In consideration of the departure from the ideal we will consider an extended family of algorithms within which T takes its place. Paragraph 3.2.2.a is essentially a re-statement of a discrete version of LaSalle's Invariance Principle (Thm. I.7.9).

3.2.2.a A convergence criterion Suppose that the discrete dynamical system

$$b_{n+1} = \hat{T}(b_n)$$

defined by the map $\hat{T}: \mathcal{B} \rightarrow \mathcal{B}$ has bounded trajectory

and enjoys the two properties

(i) φ_d is a Lyapunov function for $\hat{T}$, that is

$$\Delta\varphi_d \triangleq \varphi_d \circ \hat{T} - \varphi_d$$

is never positive

(ii) the fixed points of $\hat{T}$ are exactly the isolated local minima of φ_d .

Then the minima of φ_d are asymptotically stable with respect to the iterates of (16).

3.2.2.b A convex function The objective function,

$$\varphi_d(b) \triangleq \frac{(\sum_{i=1}^M (b_i - d_i)^2)^k}{\prod_{i=1}^M \prod_{j=i+1}^M (b_i - b_j)^2 - (\rho_i + \rho_j)^2},$$

is convex on its domain, $\mathcal{B}$, when

$$k > (M(M-1) + 1)/2,$$

in which case

$$\bar{\varphi}_d \triangleq \frac{\varphi_d^{1/k}}{[1 + \varphi_d]^{1/k}}$$

is a navigation function on each of the 2^q disconnected components of $\mathcal{B}$.

Proof. Since φ is the composition, $\eta \circ h \circ a$ of $\eta \triangleq v^k/\delta$ and

$$h(x, y) = \left[\frac{\sum_{i=1}^q x_i^2}{\prod_{i=1}^q y_i^2 - c_i^2} \right]$$

with the injective affine map,

$$a(b) \triangleq \begin{bmatrix} b - d \\ Vb \end{bmatrix}$$

(where $V \in \mathbb{R}^{q \times M}$ is appropriately defined), it has a positive definite matrix of second derivatives (for $b \neq d$) as long as

$$k > (2q + 1)/2,$$

according to Corollary 3 and Lemma 4 in Appendix B (this relationship between k and M will be assumed throughout the paper). It follows that φ_d is a strictly convex function on the domain $\mathcal{B} - d$ and every extremum must be a minimum. Moreover, since the Hessian matrix is full rank except at $b = d$, there can be no more than one extremum in any connected component on which φ_d is defined. The function is radially unbounded as well, that is $\lim_{\|b\| \rightarrow \infty} \varphi_d(b) = \infty$, according to the assumption relating k and M , so that

$$\mathcal{B}_R \triangleq \{b \in \mathcal{B} : \varphi_d(b) \leq R\}$$

is compact. Thus, each component of $\mathcal{B}$ contains one minimum of φ_d and no other extrema. In particular, d is the unique minimum in the domain, $\mathcal{B}_+$. It now follows from [6] that $\bar{\varphi}_d$ is a navigation function on each connected component of $\mathcal{B}$. $\square$

3.2.2.c Convergence of the high level controller Consider the scheduler of Paragraph 3.1.2.b with $g_k \triangleq g_d|_{\bar{b}_k(k)}$ where the index, $l(k) \triangleq l(b(s'_k))$ is chosen according to the Gauss-Southwell logic (15) and the switching tolerance, δ is to be chosen below. Under this

contact schedule the robot must at some finite time approach its nest asymptotically. Moreover, for any $\varepsilon > 0$ there can be found a $\delta > 0$ such that

$$\|b - d\| \leq \varepsilon,$$

unless the initial condition was disconnected from the goal $b_0 \notin \mathcal{B}_+$.

Proof: Paragraph 3.1.2.b defines a discrete dynamical system (16) on $\mathcal{B}$ that is allowed to iterate on $\mathcal{B} - \mathcal{T}_\delta$, where $\mathcal{T}_\delta \triangleq \{b \in \mathcal{B} : \|D\varphi_d(b)\| \leq \delta\}$. Paragraph 3.2.2.b shows that $\|D\varphi_d\|$ is a positive definite function on $\mathcal{B}$ with respect to the isolated compact set $\arg \min \varphi_d$. Thus, for any $\varepsilon > 0$ there can be found a $\delta > 0$ such that $\|D_b\varphi_d(b_0)\| < \delta$ implies

$$\inf_{b \in \arg \min \varphi_d} \|b_0 - b\| < \varepsilon$$

and $\mathcal{T}_\delta$ is the disjoint union of arbitrarily small neighborhoods of the minima of the convex function, φ_d . Thus, if the conditions of Paragraph 3.2.2.a are met, every iterate of (16) reaches $\mathcal{T}_\delta$ at some finite time and the result follows.

Condition (ii) obtains from (15).

Since φ_d is radially unbounded,

$$\mathcal{B}_R \triangleq \{b \in \mathcal{B} : \varphi_d(b) \leq R\},$$

is a compact set. It is positive invariant, $\hat{T}(\mathcal{B}_R) \subset \mathcal{B}_R$, since φ_d is a Lyapunov function. It follows that every trajectory $\hat{T}^n(b_0)$ remains in the bounded set $\mathcal{B}_{\varphi_d(b_0)}$. Paragraph 3.2.2.a now implies that the minima of φ_d are asymptotically stable.

The compact positive invariant set, $\mathcal{B}_R$, being the inverse image of a convex function, is also convex³⁹ (Thm. 1.4.6). Since, φ_d is a convex function on $\mathcal{B}_R$ (and has a positive definite Hessian matrix everywhere except $b = d$) it has one and only one extremum—a minimum—on each component of that set. Thus, the minima of φ_d constitute the forward limit set of the entire system, $\lim_{k \rightarrow \infty} \hat{T}^k(\mathcal{B}) = \arg \min \varphi_d$.

According to Paragraph 3.1.2.b **Termination** the low level closed loop system “computes another iterate” of $\hat{T}$ in finite time unless $\|D_b\varphi_d(b_0)\| < \delta$. But since $\arg \min \varphi_d$ is a globally attracting asymptotically stable set of fixed points, that condition must be reached at some finite state. $\square$

3.3 Assembly as a game of its pieces

By interpreting

$$\Theta \triangleq [\varphi_d|_{\bar{b}_1}, \dots, \varphi_d|_{\bar{b}_M}]^T$$

as a vector of “payoff functions” that describe the individual objectives of each of the M pieces to be assembled, we may interpret the high level automaton whose convergence was demonstrated in Paragraph 3.2.2.c as refereeing an M -player game. Since each component of Θ is a coordinate slice of the same global function, this is a purely cooperative game. It is perhaps unsurprising in this light that the assembly procedure developed above succeeds.

3.3.1 Non-cooperative games arising from more realistic settings. The game interpretation is much more than an intellectual digression. In the most relevant settings of the assembly problem it quickly becomes clear that purely cooperative games will not suffice to solve the problem. For example, the essential difference between the problem considered by Laumond and colleagues⁴ and those treated above lies in the separation of the robot from the environment to be manipulated. Once the robot is included as a body with physical extent within the workplace, the "obstacles" presented by the ungrasped pieces will have a very different configuration space geometry and even topology depending upon which piece the robot is grasping. To mention an extreme case, suppose one robot-piece mating forms a spherically symmetric shape in the workplace and a second mating does not. The first mating gives rise to a configuration space which is a simple cross product of a punctured Euclidean vector space with the full set of rotations. The second mating gives rise to a configuration space wherein different rotations might or might not cause a collision at the same relative point of translation. This situation does not arise in one degree of freedom problems, since without a "pushing" model, a robot on the same wire as the beads can manipulate only the two closest beads and the configuration space is trivial.

One way to proceed in the more general situation is to assign a new vector of payoffs,

$$\Theta' = [\varphi_1, \dots, \varphi_M]^T,$$

where, for example, the obstacles encoded by φ_1 correspond to the configuration space boundary formed with respect to the $M - 1$ motionless pieces that confront the mated robot-piece-number-one partial assembly. The low level control procedure would work exactly as in Paragraph 3.1.2.a. But it is not clear how to measure progress—that is, there is no longer an obvious termination condition for the high level scheduler corresponding to (14).

3.3.2 The one degree of freedom dual assembly problem revisited. Although a non-cooperative version of the one degree of freedom "assembly game" is not necessary, there nevertheless may be some benefits relative to the original cooperative formulation above. For example, return to the assembly problem with two bodies depicted in Figure 3 and consider the two objective functions

$$\varphi_1(b) \triangleq \frac{\|b_1 - d_1\|^{2k}}{(b_1 - b_2)^2 - (\rho_1 + \rho_2)^2},$$

$$\varphi_2(b) \triangleq \frac{\|b_2 - d_2\|^{2k}}{(b_1 - b_2)^2 - (\rho_1 + \rho_2)^2},$$

defined on the disjoint intervals $\mathcal{B}_i(b_j)$, $i, j = 1, 2$. Noting that φ_i represents a special case of φ (where d_j is equated to b_j), it follows that φ_i is convex as well and takes two extrema—one minimum in both of the disjoint components of $\mathcal{B}_i(b_j)$. Once again, the relation

$$\hat{d}_1(b_2) \triangleq \arg \min_{(x, b_2) \in \mathcal{B}} \varphi_1(x, b_2)$$

becomes a function upon selection of one or the other disjoint component of $\mathcal{B}$.

Graphical analysis shows that an iterated repetition of these low level controllers analogous to that analyzed in Paragraph 3.1.2.b results in convergence to the desired goal configuration. Moreover, it seems to be the case that the transients are much shorter: convergence always takes place in one or two steps when the goal and the initial configuration are in the same component of $\mathcal{B}$. Unfortunately, it is not presently clear how to generalize the analysis of this algorithm to the multi-body case. Clearly, Paragraph 3.1.2.a and Paragraph 3.1.2.b continue to hold when φ_i is substituted for $\varphi_d|_{\mathcal{B}_{i(b_i)}}$ in (11). Yet following the gradient of φ_i does not always lower the value of φ_d . It is intriguing to consider the effects in general of retaining (14) as a termination condition in this setting.

We might try to follow Li and Basar,³⁶ who show that such an algorithm gives rise to a contraction when

$$\begin{bmatrix} D_{11}^2 \varphi_1 & D_{12}^2 \varphi_1 \\ D_{21}^2 \varphi_2 & D_{22}^2 \varphi_2 \end{bmatrix}$$

has a positive determinant. Unfortunately, φ_i have degenerate Hessian matrices at isolated points of $\mathcal{B}$, and the determinant test will fail in general. On the other hand, while the contraction result seems to depend upon the convexity of the cost structure, the nondegenerate cost function, $\tilde{\varphi}_i$, is not convex.

4 CONCLUSION

This paper has proposed a new version of the assembly problem. Instead of requiring an abstract description of all mating sequences, only one is required from each possible problem configuration. However, the specification function (the map from initial configuration to mating sequence) is required to take the form of a feedback controller. That is, the mating is specified implicitly as a force law for bringing many unactuated degrees of freedom into a completed configuration with a single actuated robot possessed of far fewer degrees of freedom.

The strategy of this paper has been to focus careful attention upon what might be considered the simplest instance of the problem. This is the situation where each piece to be manipulated has only one degree of freedom and where the robot does not occupy the same workplace as the pieces. In this simple setting, the paper offers a complete algorithm along with a demonstration of its correctness. That is to say, the robot is guaranteed to arrive at its final nest destination. It will do so after repeating the operation of approaching, grasping, towing, and releasing various of the pieces. It will not leave an incomplete assembly if and only if the configuration representing the final assembly was a path connected to the configuration represented by the initial location of the pieces. It will never cause collision between two pieces in the course of manipulations. It will require only bounded forces.

This study of the simplest instance reveals sev-

general features of the problem. First, because of the nonholonomic nature of the accompanying kinematic constraints, assembly systems share with many other interesting control systems that arise in robotics the property that although completely controllable, they are not continuously feedback stabilizable. However, in contrast to the more usual classical (analytic) constraints, the impossibility of exerting even a very small influence upon arbitrarily distant bodies precludes the possibility of stabilizing assemblies with any feedback law. In assembly, *a convergent closed loop system can never be stable*. The second general observation concerns the motivation for considering hierarchical controllers. In the usual case of analytic nonholonomic constraints, if stability is to be assured then no single continuous feedback controller can be used. The necessity for some kind of discontinuous switching between the continuous pieces of a feedback law having been established, one is in a position to offer a rational justification for the introduction of hierarchy (a presently popular but poorly motivated construct within the contemporary intelligent controls literature). In assembly, the situation is still more compelling. No stability of any kind is possible within the original terms of the problem, so any solution mechanism that employs standard means of convergence (for example, with Lyapunov-like principles in mind as is almost always in the nonlinear context), must be designed and proven correct with respect to a "coarser" setting of the problem wherein the obstructions are "factored out." Finally, there appears to be in manipulation tasks a natural mode of higher level analysis which achieves this "factoring out" by removing any consideration of the robot other than as a very abstracted agent of environmental indexing and re-ordering. Namely, identifying the discrete stages of the higher level dynamics with respect to events on the contact set in the configuration space hides the dynamical irregularities and permits a new appeal to standard notions of dynamical systems theory. In our earlier juggling studies the significance of the contact set was intuitively apparent and made a formal appearance in the familiar garb of a Poincaré section.⁸ In the present study, since there is no underlying periodic phenomenon in sight, the contact set makes a muddier appearance but appears to be the simplest means of state dependent indexing.

As matters stand, one obvious task in a program of generalization beyond the present simplistic setting is the construction of navigation functions on more complex bodies moving in spaces of higher dimension. Indeed, this is a matter of active investigation in our lab at present.^{2,24} Another clear need is for the introduction of more realistic dynamical models and contact models, for instance as in Ref. 3. Yet there is at least one issue not encountered in the present study whose treatment will be essential to a truly practicable incarnation of these ideas. Although assembly in one degree of freedom raises many of the problems to be found in general, it does not provide a convenient arena for investigating the critical aspect of the robot itself occupying space in the

workplace. This is an aspect of assembly whose complications have been addressed by Laumond and colleagues.⁴ In the present context, it is intriguing to note the game theoretic interpretation of the last section and to attempt the development of a descent methodology for non-cooperative multi-player problems whose cost structure may not admit any convexity properties.

ACKNOWLEDGEMENTS

The author thanks Dr. Ernest Kent for pointing out the nonholonomic nature of assembly control tasks. Conversations with Prof. Isil Bozma, Prof. Tamer Basar, Prof. Martin Bühler, Prof. N. Harris McClamroch, Prof. Marc Railbert, Prof. Elon Rimon, Mr. Alfred Rizzi, Mr. Sanjiv Singh, and Dr. Louis Whitcomb have helped clarify the proper scope and appropriate tools for this paper. Thanks to Brian Clendennin for helping rid the revised version of typos.

APPENDIX A: TASK ENCODING: TRANSLATING GOALS INTO CONTROLLERS

The investigation of task encoding principles arises from the effort to render abstract goals into feedback laws. This section offers a brief review of techniques for encoding two rather different classes of robotic tasks, both of which have been developed in previous independent research. Perhaps fortuitously, the control policy developed in this paper represents a combination of both.

A.1 Navigation

Let a fully actuated mechanism move in a cluttered but perfectly known workplace. There is a particular location of interest and it is desired that the "robot" move to this location from anywhere else in the workplace without colliding with the obstacles present. Initiated by Khatib a decade ago,⁴⁰ the idea of using artificial potential functions for robot task description and control was adopted or re-introduced independently by a number of researchers.⁴¹⁻⁴³ Gradually, there seems to have emerged a common awareness of several fundamental problems with the potential function methodology. Spurious local minima seemed unavoidable, and unrealizable infinite torques were thought to be required at the obstacle boundaries. In fact, an artificial potential function need satisfy a longer list of technical conditions in order to give rise to a bounded torque feedback controller that guarantees convergence to the goal state, from almost every initial configuration. This list comprises the notion of a *navigation function* introduced to the literature two years ago.⁴⁴ Roughly speaking, a navigation function attains its maximal value of unity on the boundary of connected domain (if the domain is not compact then one considers the point at infinity to be part of the boundary), and has one and only one minimum on that domain.

The question immediately arises whether such desirable features may be achieved in general. In fact, the answer is affirmative: smooth navigation functions

exist on any compact connected smooth manifold with boundary.⁶ Thus, in any problem involving motion of a mechanical system through a cluttered space (with perfect information and no requirement of physical contact) if the problem may be solved at all, we are guaranteed that it may be solved by a navigation function. There remains the engineering problem of how to construct such functions.

In previous work we have found the following simple "encoding" of holonomic constraints to result in navigation functions. Let v be the euclidean distance to the goal state, and let δ provide some measure of the distance away from a forbidden region. Typically, δ will not satisfy any formal metric properties, but it must vanish on the boundary and remain positive within the interior of the allowable region. Encode the objective "go to $v=0$ and do not go to $\delta=0$ " via the quotient, v/δ . Strengthen the force of the attracting objective by taking the power,

$$\varphi \triangleq \frac{v^k}{\delta}.$$

Since φ is unbounded, normalize it via the composition with a smooth "squashing" function,

$$\sigma_k(x) \triangleq \frac{x^{1/k}}{(1+x)^{1/k}}.$$

Note that the composition

$$\tilde{\varphi} \triangleq \sigma_k \circ \varphi = \frac{v}{(v^k + \delta)^{1/k}}$$

is a kind of analytic "switch" that vanishes on the goal state, goes to unity exactly on the bad states, and varies smoothly in between. It is shown in ref. 6 that this construction indeed satisfies the conditions of a navigation function on spaces of a particularly simple topology and geometry, the "sphere worlds", whose obstacles are disjoint Euclidean balls in a Euclidean vector space. In this paper it will be seen that the construction results again in a navigation function in a geometrically simple but topologically distinct setting.

A.2 Juggling

Consider a frictionless plane inverted into the earth's gravitational field. Two pucks are allowed to slide freely on this plane except when batted by a simple "robot" – a revolute motor with a bar attached to it whose axis of rotation is orthogonal to the plane. The robot has one degree of actuated freedom (perhaps one and a half, if one considers the recourse to "whole arm" manipulation⁴⁵ as adding freedom) while the environment possesses two for each puck. The robot is given the task of repeatedly batting the two pucks so that each one attains a periodic trajectory whose apex lies at a specified vertical and horizontal locus on the inverted plane.

The same nonholonomic constraints are present in juggling as in assembly:³ no imposition of control upon the environment is possible until contact has been made. We have shown⁴⁶ that the "vertical one-juggle" task – batting a single puck on an inclined plane so that it

eventually attains a repeated purely vertical motion at a specified horizontal position – may be encoded as a fixed point of a certain discrete dynamical system. Moreover, we have shown how to construct a sensor based feedback strategy for the robot that accomplishes the task⁸ and is provably correct as well.^{22,47} Let b denote the puck's position and r denote the robot's joint angle. The effective strategy calls for the robot to track a "mirrored" reflection of the puck's trajectory, $r = \mu(b, \dot{b})$. The correctness proof involves passage to the discrete dynamical system (now interpreted as the return map on a Poincaré Section of the limit cycle) with respect to which the encoded goal point is shown to be globally asymptotically stable.⁴⁸ In the present context it is useful to interpret the discrete dynamics as a "higher level" control process brought about by appropriately abstracting away the "lower level" details of the robot's torque control strategy.

The primitive one-juggle solution extends rather simply to the case of juggling two pucks simultaneously. Although the procedure is heuristic, extensive experimental study reveals that it works remarkably well in practice. There are two "low level" strategies, μ_l, μ_r that solve the vertical one juggle for the left and right hand pucks, respectively. Given only one robot (with one degree of freedom), the question arises as to how these two strategies should be "assigned" to the robot. Consider an emergency situation, when both pucks are falling toward the bar nearly at the same time. It is imperative to service the nearest first. Moreover, it is well worth sacrificing any nominally desirable one-juggle performance to the work of keeping both aloft and restoring phase separation between them. This intuition can be readily implemented by use of an "analytic switch" that triggers on the good and bad events interpreted in each puck's phase space.⁹ The same underlying approach can be brought to bear in mediating between the conflicting subgoals of difficult pieces in assembly.

APPENDIX B: DETAILS OF COMPUTATION

It will prove useful to have a general formula for the Hessian matrix of a composed scalar valued function. Thus, supposing $h: \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $\gamma: \mathbb{R}^m \rightarrow \mathbb{R}$, denoting Dh by H , $D\gamma$ by g^T and $D^2\gamma$ by G , we have

$$D(\gamma \circ h) = (g^T \circ h) \cdot H$$

and

$$\begin{aligned} D^2(\gamma \circ h) &= D(H^T \cdot (g \circ h)) \\ &= [(g^T \circ h) \otimes I] D(H^T)^s + H^T \cdot D(g \circ h). \end{aligned}$$

It follows that

$$D^2(\gamma \circ h) = H^T \cdot (G \circ h) \cdot H + [(g^T \circ h) \otimes I] D(H^T)^s$$

B.1 Convexity

Lemma 1: Fix an arbitrary constant $c \in \mathbb{R}^q$. For

Autonomous assembly

variables $n \in \mathbb{R}^p$, $d \in \mathbb{R}^q$, consider the map

$$h: \mathbb{R}^{p+q} \rightarrow \mathbb{R}^2: (n, d) \mapsto \begin{bmatrix} v \\ \delta \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^p n_i^2 \\ \prod_{i=1}^q d_i^2 - c_i^2 \end{bmatrix}.$$

If $\gamma: \mathbb{R}^2 \rightarrow \mathbb{R}$ is (strictly) convex, $D_1\gamma > 0$, $D_2\gamma < 0$, and

$$\left([D^2\gamma]((v, \delta)) + \begin{bmatrix} D_1\gamma \frac{1}{2v} & 0 \\ 0 & D_2\gamma \frac{(2q-1)}{2q\delta} \end{bmatrix} \right) > 0$$

then $\gamma \circ h$ is (strictly) convex as well.

Proof: We have

$$D_n v = 2n^T; \quad D_n^2 v = 2I_p,$$

and, letting N denote an orthogonal ($p \times p$) matrix whose first column is $n/\|n\|$, the Hessian may be re-written in the form

$$D_n^2 v = 2NN^T = \frac{2}{v} nn^T + P; \quad P \triangleq N \begin{bmatrix} 0 & 0_{p-1}^T \\ 0_{p-1} & I_{p-1} \end{bmatrix} N^T,$$

where P is a positive semi-definite matrix whose kernel is in the subspace spanned by n . Fixing the notation,

$$\Delta \triangleq \begin{bmatrix} d_1 & 0 \\ & \ddots \\ 0 & d_q \end{bmatrix}; \quad \Xi \triangleq \begin{bmatrix} c_1 & 0 \\ a & \ddots \\ 0 & c_q \end{bmatrix};$$

$$\Gamma \triangleq (\Delta^2 - \Xi^2)\Delta; \quad d^{-1} \triangleq \Gamma o; \quad o \triangleq \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix},$$

yields

$$D_d \delta = 2\delta d^{-T}; \quad D_d^2 \delta = -2\delta \Gamma (\Xi^2 + I_q - 2oo^T) \Gamma.$$

Similarly letting Θ denote an orthogonal ($q \times q$) matrix whose first column is $o/q^{1/2}$, the Hessian of δ may be re-written in the form

$$\begin{aligned} D_d^2 \delta &= -2\delta \Gamma \Theta \begin{bmatrix} 1-2q & 0_{q-1}^T \\ 0_{q-1} & I_{q-1} \end{bmatrix} \Theta^T \Gamma - 2\delta \Gamma \Xi^2 \Delta^{-2} \Gamma \\ &= \frac{2\delta(2q-1)}{q} d^{-1} d^{-T} - Q \end{aligned}$$

$$Q \triangleq 2\delta \Gamma \left(\Theta \begin{bmatrix} 0 & 0_{q-1}^T \\ 0_{q-1} & I_{q-1} \end{bmatrix} \Theta^T + \Xi^2 \Delta^{-2} \right) \Gamma$$

Again, Q is a positive (semi)-definite matrix (whose kernel lies in the subspace spanned by o if Γ is singular).

Noting that

$$H = Dh = \begin{bmatrix} Dv \\ D\delta \end{bmatrix} = \begin{bmatrix} D_n v & 0_q^T \\ 0_p^T & D_d \delta \end{bmatrix} = \begin{bmatrix} 2n^T & 0_q^T \\ 0_p^T & 2\delta d^{-T} \end{bmatrix}$$

we may evaluate the expression (18), as

$$\begin{aligned} D^2(\gamma \circ h) &= \begin{bmatrix} 2n & 0_p \\ 0_q & 2\delta d^{-1} \end{bmatrix} [D^2\gamma]((v, \delta)) \begin{bmatrix} 2n^T & 0_q^T \\ 0_p^T & 2\delta d^{-T} \end{bmatrix} \\ &\quad + D_1\gamma(v) \begin{bmatrix} D_n^2 v & 0_{p \times q} \\ 0_{q \times p} & 0_{q \times q} \end{bmatrix} \\ &\quad + D_2\gamma(\delta) \begin{bmatrix} 0_{p \times p} & 0_{p \times q} \\ 0_{q \times p} & D_d^2 \delta \end{bmatrix} \end{aligned}$$

$$\begin{aligned} &= \begin{bmatrix} 2n & 0_p \\ 0_q & 2\delta d^{-1} \end{bmatrix} \left([D^2\gamma]((v, \delta)) \right. \\ &\quad \left. + \begin{bmatrix} D_1\gamma/2v & 0 \\ 0 & D_2\gamma \frac{(2q-1)}{2q\delta} \end{bmatrix} \right) \begin{bmatrix} 2n^T & 0_q^T \\ 0_p^T & 2\delta d^{-T} \end{bmatrix} \\ &\quad + \begin{bmatrix} D_1\gamma(v)P & 0_{p \times q} \\ 0_{q \times p} & -D_2\gamma(v)Q \end{bmatrix}. \end{aligned}$$

These two singular symmetric matrices have kernel that intersect at most at 0. The second matrix is positive semi-definite, thus the joint Hessian will be positive definite if the first matrix is positive semi-definite as claimed. $\square$

Lemma 2 The map

$$\gamma: \mathbb{R}^2 \rightarrow \mathbb{R}: (v, \delta) \mapsto v^k/\delta$$

is strictly convex when $\gamma > 0$ as long as $k > 2$.

Proof: We have

$$D\gamma = \gamma v^T; \quad v^T = [k/v, -1/\delta]$$

hence

$$\begin{aligned} D^2\gamma &= \gamma(vv^T + Dv) \\ &= \gamma \begin{bmatrix} k(k-1)/v^2 & -k/v\delta \\ -k/v\delta & 2/\delta^2 \end{bmatrix} \end{aligned}$$

Since the determinant of the last array is $k(k-2)/v^2\delta^2 > 0$, for $k > 2$, the condition implies that $D^2\gamma$ is positive definite when $\gamma > 0$. $\square$

Corollary 3: The function $\hat{\phi} \triangleq \gamma \circ h$ is strictly convex as long as $k > (2q+1)/2$.

Proof: We must show that

$$\begin{bmatrix} k(k-1)/v^2 & -k/v\delta \\ -k/v\delta & 2/\delta^2 \end{bmatrix} + \begin{bmatrix} \frac{k}{2v^2} & 0 \\ 0 & -\frac{(2q-1)}{2q\delta^2} \end{bmatrix} > 0$$

This is equivalent to

$$\begin{aligned} &\left| \begin{bmatrix} k(2k-1)/2v^2 & -k/v\delta \\ -k/v\delta & \frac{2q+1}{2qb^2} \end{bmatrix} \right| \\ &= \frac{k(2k-1)(2q+1) - 4qk^2}{4qv^2\delta^2} > 0, \end{aligned}$$

or $k(2k - (2q+1)) > 0$ as claimed. $\square$

Lemma 4: If $a: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is affine and $\hat{\phi}: \mathbb{R}^m \rightarrow \mathbb{R}$ is convex then $\varphi \triangleq \hat{\phi} \circ a$ is convex as well. If $n \leq m$ and a is injective and $\hat{\phi}$ is strictly convex then φ is also strictly convex.

Proof: This follows from (18) after noting that $D(A^T)^S = D^2a = 0$ according to the affine assumption. $\square$

References

1. J. Baillieul and D.P. Martin, "Resolution of kinematic redundancy" *Robotics* (American Mathematical Society, 1990).
2. Louis L. Whitcomb and Daniel E. Koditschek, "Automatic assembly planning and control via potential functions" *Proc. IEEE International Workshop on Intelligent Robots and Systems* Osaka, Japan, Nov. 1991. IEEE (1991) pp. 55-60.
3. Daniel E. Koditschek, "Assembly: Another source of nonholonomy in robotics" *Proc. American Control Conference*, Boston, MA, June 1991. American Society of Control Engineers (1991) pp. 1627-1632.
4. Rachid Alami, Thierry Simeon and Jean-Paul Laumond, "A geometrical approach to planning manipulation tasks. The case of discrete placements and grasps". In (H. Miura and S. Arimoto, editors) *Robotics Research* 453-463 (MIT Press, 1990).
5. Oussama Khatib, "Real time obstacle avoidance for manipulators and mobile robots" *The Int. J. Robotics Research* 5(1) 90-99 (Spring 1986).
6. Daniel E. Koditschek and Elon Rimon, "Robot navigation functions on manifolds with boundary" *Advances in Applied Mathematics* 11 412-442 (1990).
7. Daniel E. Koditschek, "The control of natural motion in mechanical systems" *ASME J. Dynamic Systems, Measurement, and Control* 113(4) 547-551 (Dec. 1991).
8. M. Bühler, D.E. Koditschek and P.J. Kindlmann, "A family of robot control strategies for intermittent dynamical environments" *IEEE Control Systems Magazine* 10 16-22 (Feb. 1990).
9. M. Bühler, D.E. Koditschek and P.J. Kindlmann, "Planning and Control of Robotic Juggling Tasks" In (H. Miura and S. Arimoto, editors) *Fifth International Symposium on Robotics Research* (MIT Press, 1990) pp. 321-332.
10. J.E. Hopcroft, J.T. Schwartz and M. Sharir, "On the complexity of motion planning for multiple independent objects: Pspace hardness of the 'Warehouseman's Problem'" *Robotics Report 14* (N.Y.U. Courant Institute Department of Computer Science, New York, 1984).
11. T. Lozano-Pérez et al., "A robot systems that recognizes, plans, and manipulates" *Proceedings IEEE Int. Conf. Robotics and Aut.*, (1987) pp. 843-849.
12. Susan N. Gottschlich and Avinash C. Kak, "A dynamic approach to high-precision parts mating" *IEEE Transactions on Systems, Man, and Cybernetics* 19(4) 797-810 (July, 1989).
13. Seth A. Hutchinson and Avinash C. Kak, "Planning sensing strategies in robot work cell with multi-sensor capabilities" *IEEE J. Robotics and Automation* 5(6) 765-783 (Dec., 1989).
14. L.S. Homem de Mello and A.C. Sanderson, "Representations of mechanical assembly sequences" *IEEE J. Robotics and Automation* 7(2) 211-227 (Apr., 1991).
15. L.S. Homem de Mello and A.C. Sanderson, "A correct and complete algorithm for generation of mechanical assembly sequences" *IEEE J. Robotics and Automation*, 7(2) 228-240 (Apr., 1991).
16. Jean-Claude Latombe, *Robot Motion Planning* (Kluwer, Boston, MA, 1991).
17. R. Murray and S. Sastry, "Grasping and manipulation using multifingered robot hands" *Robotics, Proceedings of Symposia in Applied Mathematics Vol. 41*, Providence, RI, 1990 American Mathematical Society (1990).
18. P.S. Krishnaprasad, R. Yang and W. Dayawansa, "Control problems on principle bundles and nonholonomic mechanics" *Proc. IEEE Conf. Dec. Contr. Brighton, UK*, Dec., 1991 IEEE (1991) pp. 1133-1138.
19. Leo Dorst, Indur Mandhyan and Karen Trovato, "The geometrical representation of path planning problems" Technical Report MS-90-054 (Philips Laboratories Briarcliff Manor, July, 1990).
20. C. Samson and K. Ait-Abderrahim, "Feedback control of a nonholonomic wheeled cart in Cartesian space" *Proc. IEEE Int. Conf. Robt. Aut.*, Apr., 1991 (1991) pp. 1136-1141.
21. Anthony M. Bloch and N. Harris McClamroch, "Control of mechanical systems with classical non-holonomic constraints" *Proc. 28th IEEE Conf. on Decision and Control, Tampa, FL*, Dec., 1989 (1989) pp. 201-205.
22. D.E. Koditschek, "Globally stable closed loops imply autonomous behavior" *Proc. Fifth IEEE International Symposium on Intelligent Control*, Philadelphia, PA, Sept. 1990 pp. 651-656.
23. Daniel E. Koditschek, "Robot control systems, In (Stuart Shapiro, editor) *Encyclopedia of Artificial Intelligence* (John Wiley and Sons, New York, 1987) pp. 902-923.
24. E. Rimon and D.E. Koditschek, "Exact robot navigation in topologically simple but geometrically complicated environments" *Proc. IEEE International Conference on Robotics and Automation, Cincinnati, OH*, May, 1990 (1990) pp. 1937-1943.
25. E. Rimon and D.E. Koditschek, "The construction of analytic diffeomorphisms for exact robot navigation on star worlds" *Transactions of the American Mathematical Society* 327(1) 71-115 (Sep., 1991).
26. Matthias Kowski, "Stabilization of nonlinear systems in the plane" *Systems and Control Letters* 12 169-177 (1989).
27. D.E. Whitney, "Force feedback control of manipulator fine motions" *ASME J. Dyn. Syst.* 98 91-97 (June 1977).
28. Tomás Lozano-Pérez, Matthew T. Mason and Russell H. Taylor, "Automatic synthesis of fine-motion strategies for robots" *The Int. J. Robotics Research* 3(1) 3-23 (1984).
29. Ju. I. Neimark and N.A. Fufaev, *Dynamics of Nonholonomic Systems Vol. 33 Translations of Mathematical Monographs* (American Mathematical Society, Providence, RI, 1972).
30. H.J. Sussmann, "A general theorem on local controllability" *SIAM J. on Control and Optimization* 25(1) 158-194 (1987).
31. E.D. Sontag, *Mathematical Control Theory* (Springer-Verlag, NY, 1990).
32. Roger W. Brockett, "Asymptotic stability and feedback stabilization" In (Roger W. Brockett, Richard S. Millman and Hector J. Sussman, editors) *Differential Geometric Control Theory*, chapter 3 (Birkhäuser, 1983) pp. 181-191.
33. M. Reyhanoglu and N.H. McClamroch, "Reorientation of space multibody systems maintaining zero angular momentum" *AIAA Guidance, Navigation and Control Conf. New Orleans, La. AIAA 1330-1340* (1991).
34. J. Zabczyk, "Some comments on stabilizability" *App. Math. Optim.* 19 1-10 (1989).
35. J.P. Lasalle, *The Stability of Dynamical Systems, volume 25 of Regional Conference Series in Applied Mathematics* (SIAM, Philadelphia, PA, 1976).
36. Shu Li and Tamer Basar, "Distributed algorithms for the computation of noncooperative equilibria" *Automatica* 23(4) 523-533 (1987).
37. G.E. Forsythe and W.R. Wasow, *Finite Difference Methods for Partial Differential Equations* (John Wiley New York, 1960).
38. David G. Luenberger, *Linear and Nonlinear Programming* (Addison-Wesley, Reading, MA, 1984).
39. R. Tyrrell Rockafellar, *Convex Analysis* (Princeton U.P. Princeton, NJ, 1970).
40. O. Khatib and J.-F. Le Maitre, "Dynamic control of manipulators operating in a complex environment" *Proceedings Third International CISM-IFTOMM Symposium, Udine, Italy*, Sept., 1978 (1978) pp. 267-282.
41. F. Miyazaki, S. Arimoto, M. Takegaki and Y. Maeda, "Sensory feedback based on the artificial potential for robot manipulators" *Proceedings 9th IFAC, volume 6*

- Budapest, Hungary, 1984) (1984) pp. 2381-2386.
42. J.R. Andrews and N. Hogan, "Impedance control as a framework for implementing obstacle avoidance in a manipulator" In (David E. Hardt and Wayne J. Book, editors) *Control of Manufacturing and Robotic Systems* (A.S.M.E., Boston, MA, 1983) pp. 243-251.
 43. V.V. Pavlov and A.N. Voronin, "The method of potential functions for coding constraints of the external space in an intelligent mobile robot" *Soviet Automatic Control* 6 45-51 (1984).
 44. E. Rimon and D.E. Koditschek, "Exact robot navigation using cost functions: The case of spherical boundaries in E". *IEEE International Conference on Robotics and Automation Philadelphia, PA, Apr., 1988* (1988) pp. 1791-1796.
 45. Kenneth Salisbury, Brian Eberman, Michael Levin and William Townsend, "The design and control of an experimental whole-arm manipulator" *Proc. Fifth Int. Symp. Rob. Res., Tokyo, Japan, Sept., 1989* (1989) pp. 261-269.
 46. M. Bühler and D.E. Koditschek, "A prelude to juggling" *26th IEEE Conference on Decision and Control Los Angeles, CA., Dec., 1987* (1987) (paper available from authors - not in proceedings).
 47. M. Bühler, D.E. Koditschek and P.J. Kindlmann, "A Simple Juggling Robot: Theory and Experimentation" In (V. Hayward and O. Khatib, editors) *Experimental Robotics I* (Springer-Verlag, Berlin 1990) pp. 35-73.
 48. M. Bühler and D.E. Koditschek, "From stable to chaotic juggling" *Proc. IEEE International Conference on Robotics and Automation, Cincinnati, OH, May, 1990* (1990) pp. 1976-1981.

oratories,

control of
ice" *Proc.*
1991) pp.

, "Control
holonomic
ision and
-205.

ops imply
ternational
PA, Sept,

In (Stuart
Intelligence
-923.

navigation
complicated
ference on
May, 1990

ruction of
ion on star
cal Society

ems in the
(1989).
manipulator
ne 1977).

Russell H.
ategies for
(1984).

amics of
Mathemati-
ty, Provid-

al control-
tion 25(1)

(Springer-

l feedback
S. Millman
Geometric
. 181-191.
entation of
o angular
id Control
l).
lity" *Appl*

ns, volume
athematics.

ims for the
Automatica

Difference
ohn Wiley,

ogramming

etron U.P.

control of
vironment"
A-IFTOMM
. 267-282.
Y. Maeda,
potential for
volume 6,