2005 American Control Conference
June 8-10, 2005. Portland, OR, USA

Application of Supervisory Control Methods to Uncertain Multiple Contact Mechanical Systems

Todd D. Murphey
Electrical and Computer Engineering
University of Colorado
Boulder, CO 80309
murphey@colorado.edu

Abstract— **Estimation of contact state is important to any multi-point interaction that involves frictional stick/slip phenomena. In particular, when there are more kinematic constraints than there are degrees of freedom, some contact interfaces must slip, leading to the need for contact state estimation. Fortunately, supervisory control techniques from adaptive control can be applied to this problem with relatively little modification. We discuss this approach in terms of a distributed manipulation experiment developed to explore overconstrained manipulation. In this context, we show in a simulated model that on-line contact state estimation dramatically improves performance over methods that estimate contact states off-line.**

I. Introduction

A manipulation system consisting of many points of contact typically exhibits stick/slip phenomenon due to the point contacts moving in kinematically incompatible manners. We call this manner of manipulation *overconstrained manipulation* because not all of the constraints can be satisfied. Naturally, uncertainty due to overconstraint can sometimes be mitigated by having backdrivable actuators, soft contacts, and by other mechanical means, but these approaches avoid the difficulties associated with stick/slip phenomenon at the expense of losing information about the state of the mechanism. This, in turn, leads either to degraded performance or to requiring additional sensors. This paper is concerned with systems that have multiple points of contact, all of which are frictional and adequately described by either constraint forces (when there is no slipping at the point contact) or by the slipping reaction force. Prototypes of this situation include distributed manipulation systems, such as those found in [12], [9], as discussed in Section II.

An important question in these systems is that of contact state estimation [13]. That is, for n kinematic constraints associated with contact interfaces, estimating at any given time which constraints are satisfied (the contact is "sticking") and which are not (the contact is "slipping"). (Note that the contact state is therefore associated with a boolean variable that is zero when the sticking and one when slipping. We will elaborate on this in the Section III.) How can one determine the stick/slip state for each point contact? Without a sensor at each point contact, the output signal must be used in some way to determine this. Moreover, the computational complexity of the solution must be considered as well, since for an n contact system there are 2^n possible stick/slip combinations. The main contribution of this paper is to show that methods from adaptive control called supervisory control [2], [6], [7] provide a reasonable, though in no sense optimal, solution to this problem.

This paper is organized as follows. Section II discusses distributed manipulation in more detail, and discusses the experimental implementation used before in [12]. Section III describes the algorithm developed in [12] and gives an example simulation for this experimental system when the contact states are assumed to be known perfectly. We then illustrate how variations in the contact state can, not surprisingly, degrade the performance of the algorithm. Section IV gives the necessary background for understanding a supervisory control system, such as that described in [2], [6], and proves the relevant stability properties. We also discuss the implementation of this adaptive control method to the distributed manipulation example and show in simulation that the original algorithm performance is indeed recovered even when the contact states are not known a priori.

II. Motivation: Distributed Manipulation

Distributed manipulators usually consist of an array of similar or identical actuators combined together with a control strategy to create net movement of an object or objects. The goal of many distributed manipulation systems is to allow precise positioning of planar objects from all possible starting configurations. Such "smart conveyors" can be used for separating and precisely positioning parts for the purpose of assembly. Distributed manipulator actuation methods ranges from air jets, rotating wheels, and electrostatics on the macro-scale, to MEMS and flexible cilia at the micro-scale.

Methods to design distributed manipulation control systems have been proposed in several works, including [12], [9], [4], [3], [5]. However, in cases where only a small number of actuators are in contact with the manipulated object or the coefficient of friction μ is very high, continuous approximations of these systems have been shown experimentally not to work well [12], [9]. In these cases, the physics of the actual array and the object/array interface must be incorporated into the control design process. In particular, the discontinuous nature of the equations of motion must be addressed.

The work in [12] describes an experimental test-bed that was designed to evaluate and validate such control systems. Our modular system can emulate a reasonably large class of distributed manipulators that generate motion through rolling and sliding frictional contact between the moving object and actuator surfaces. In such cases friction forces and intermittent contact play an important role in the overall system dynamics, leading to non-smooth dynamical system behavior. The control question is twofold in its theoretical interest. First, unlike many other control problems currently being studied, distributed manipulation problems are typified by being massively overactuated. A planar distributed manipulation problem will typically only have three outputs (x, y, θ), but it may potentially have thousands of inputs. Therefore, control schemes must scale with the number of actuators in order to be able to implement them on real devices such as MEMS arrays. Second, there is the question of physical modeling. Partly because of the aforementioned overactuation, nonsmooth effects become commonplace in distributed manipulation due to intermittent contact, friction, and kinematically incompatible constraints. When these are the dominant concerns, they must be incorporated into the modeling and therefore into the control design as well. Control laws appropriate to these systems have been successfully designed, as discussed in Section III.

(a)


(b)


Fig. 1. The Caltech Distributed Manipulation System. (a) Front View (b) Module

A photograph of the apparatus can be seen in Figure 1. The design is a modular one based on a basic cell design. Each cell contains two actuators. One actuator orients the wheel axis, while the other actuator drives the wheel rotation (see Figure 1(b)). These cells can easily be repositioned within the supporting structure to form different configurations. The system shown in Figure 1(a) is configured with a total of nine cells—though more can be easily added. The position and orientation of the manipulated object is obtained and tracked visually. To enable visual tracking, a right triangle is affixed to the moving object. For more details on the experimental setup, please refer to [12].

When experiments were performed using this device in [12], the contact state was estimated in an open loop manner. That is, based on physical principles (described momentarily), a static condition was chosen under which the contact states would change. Somewhat surprisingly, this worked in the implementation, but only because the device was in a controlled environment and it was well characterized. In most scenarios this will not be the case. Examples include outdoor slip-steered vehicles that have multiple contacts with the ground that are not well characterized because of changing ground characteristics as well as MEMS manipulation where electrostatic forces are difficult to accurately model. Therefore, we necessarily must address the problem of contact state estimation and accommodation. Even in a nine cell system, this is a nontrivial task. Each wheel has two constraints (a rolling constraint and a no-sideways-slip constraint) leading to a grand total of $2^{18} \sim 10^5$ possible contact states. Although not formally addressed here, there exist techniques to reduce the number of possible states, including kinematic reduction[11] and coordination[10].

III. Modeling and Analysis

To explicitly investigate, incorporate, and control the complex frictional contact phenomena inherent in overconstrained manipulation, one needs to develop general modeling schemes that can capture these phenomena without being intractable from a control perspective. One could resort to a general Lagrangian modeling approach that accounts for the contact effects through Lagrange multipliers. Instead, we sought to develop a general modeling scheme that captures the salient physical features, while also leading to equations that are amenable to control analysis.

To realize this goal, we use a "Power Dissipation Method" (PDM) approach to model the governing dynamics of an overconstrained mechanical system involving a discrete number of frictional contacts. One can show that this method almost always produces unique models [11] that are relatively easy to compute, are formally related to the Lagrangian mechanics, and to which one can apply control system analysis methods. This method produces first-order governing equations, instead of second-order equations that are associated with Lagrange's equations.

Assume that the moving body and actuator elements that contact the object can be modeled as rigid bodies making point contacts that are governed by the Coulomb friction law at each contact point. Let q denote the configuration of the array/object system, consisting of the object's planar location, and the variables that describe the state of each actuator element. Under these conditions, the relative motion of each contact between the object and an actuator array element can be written in the form $\omega(q)\dot{q}$. If $\omega(q)\dot{q} = 0$, the contact is not slipping, while if $\omega(q)\dot{q} \neq 0$, then $\omega(q)\dot{q}$ describes the slipping velocity.

In general, the moving object will be in contact with the actuator array at many points. From kinematic considerations, one or more of the contact points must be in a slipping state, thereby dissipating energy. The *power dissipation function* measures the object's total energy dissipation due to contact slippage.

Definition 3.1: The *Dissipation* or *Friction Functional* for an n-contact state is defined to be

$$\mathcal{D} = \sum_{i=1}^{n} \mu_i N_i \mid \omega(q)\dot{q} \mid \tag{1}$$

with μ_i and N_i being the Coulomb friction coefficient and normal force at the i^{th} contact, which are assumed known.

Assuming that the motion of the actuator array's variables are known, the *power dissipation method* assumes that the object's motion at each instant is the one that instantaneously minimizes power dissipation $\mathcal{D}$ due to contact slippage. This method is adapted from the work of [1] on wheeled vehicles. For a greater discussion of the formal characteristics of the PDM, and a discussion of the relationship between the PDM and Lagrangian approaches for such a system, see [11].

When one applies the PDM method, the governing equations that result take the form of a multiple model system.

Definition 3.2: A control system Σ evolving on a smooth n-dimensional manifold, Q, is said to be a *multiple model driftless affine system (MMDA)* if it can be expressed in the form

$$\Sigma: \qquad \dot{q} = f_1(q)u_1 + f_2(q)u_2 + \cdots + f_m(q)u_m. \tag{2}$$

where $q \in Q$. For any q and t, the vector field f_i assumes a value in a finite set of vector fields: $f_i \in \{g_{\alpha_i} | \alpha_i \in I_i\}$, with I_i an index set. The vector fields g_{α_i} are assumed to be analytic in (q, t) for all α_i, and the controls $u_i \in \mathbb{R}$ are piecewise constant and bounded for all i. Moreover, letting σ_i denote the "switching signals" associated with f_i

$$\begin{array}{rlcl} \sigma_i: & Q \times \mathbb{R} & \longrightarrow & \mathbb{N} \\ & (q, t) & \longrightarrow & \alpha_i \end{array}$$

the σ_i are measurable in (q, t).

An MMDA is a driftless affine nonlinear control system where each control vector field may "switch" back and forth between different elements of a finite set. In our case, this switching corresponds to the switching between different *contact states* between the object and the array surface elements (i.e., different sets of slipping contacts) due to variations in contact geometry, surface friction properties, and normal loading. In [11] it was shown that the PDM generically leads to MMDA systems as in Definition 3.2 and is formally equivalent to a kinematic reduction of the Lagrangian formulation of the equations of motion.

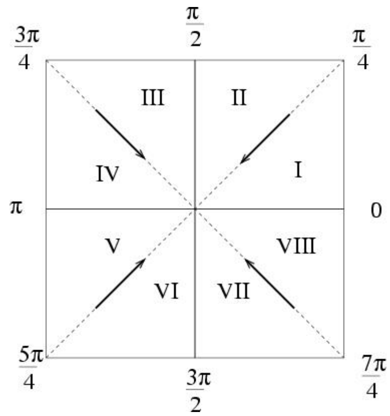

Fig. 2. A distributed manipulator with four actuators

The work in [12] showed that the PDM predicts that the governing equations for a distributed manipulation system are:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{\theta} \end{bmatrix} = f_1 u_1 + f_2 u_2 \tag{3}$$

where

$$f_1 \in \begin{bmatrix} \frac{-y_i}{(x_j - x_i)s_j + (y_i - y_j)c_j} \\ \frac{x_i}{(x_j - x_i)s_j + (y_i - y_j)c_j} \\ \frac{1}{(x_i - x_j)s_j + (y_j - y_i)c_j} \end{bmatrix} \quad f_2 \in \begin{bmatrix} \frac{s_j((x_i - x_j)c_i + y_i s_i) + c_i c_j y_j}{(x_j - x_i)s_j + (y_i - y_j)c_j} \\ \frac{-c_i c_j x_i - s_i(x_j s_j - (y_i - y_j)c_j)}{(x_j - x_i)s_j + (y_i - y_j)c_j} \\ \frac{-\cos(\theta_i - \theta_j)}{(x_i - x_j)s_j + (y_j - y_i)c_j} \end{bmatrix}$$

where $c_i = \cos(\theta_i)$, $s_i = \sin(\theta_i)$, etc. The input u_1 is the input to the closest actuator to the center of mass, and the input u_2 is the input to the second closest actuator to the center of mass. It should be noted that here the index notation should be thought of as mapping (i, j) pairs to equations of motion in some neighborhood (not necessarily small) around the i^{th} and j^{th} actuator. The transition between the equations of motion determined by actuators i and j to equations of motion determined by actuators k and l will in general be determined by the location of center of mass. This in turn leads to the state space being divided up by transition boundaries between different sets of equations of motion.

Consider Figure 2, which might represent a portion of a distributed manipulator near a desired equilibrium point. This region has four actuators (corresponding to the inputs $u_1, \ldots, u_4$ and represented in the figure by arrows) located at $(\pm 1, \pm 1)$, all pointed towards the origin. An analysis of this system using the PDM method shows that the region can be divided into 8 distinct regions, labeled $\mathbf{I} - \mathbf{VIII}$, where one contact state holds. These are separated by 8 boundaries, labeled $0 - 2\pi$ in increments of $\frac{\pi}{4}$. In each one of the regions $\mathbf{I} - \mathbf{VIII}$ a control law is calculated from the Lyapunov function $k(x^2 + y^2 + \theta^2)$ by solving $\dot{V} = -V$ for u_i, where k is some constant to be chosen

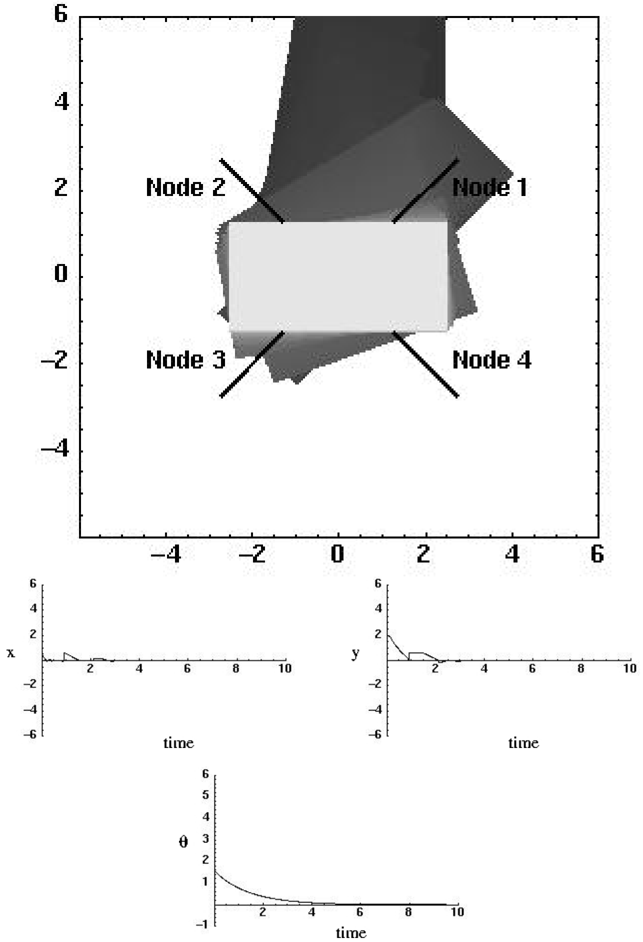

Fig. 3. Simulation of distributed manipulation when the contact state is known perfectly. The rectangle represents the center of the object which is actually in contact with all four of the actuators (Nodes 1-4). The time history progresses from dark triangles at time 0 to the light triangle at time 10. The bottom three plots are plots of the X,Y, and θ coordinates against time.

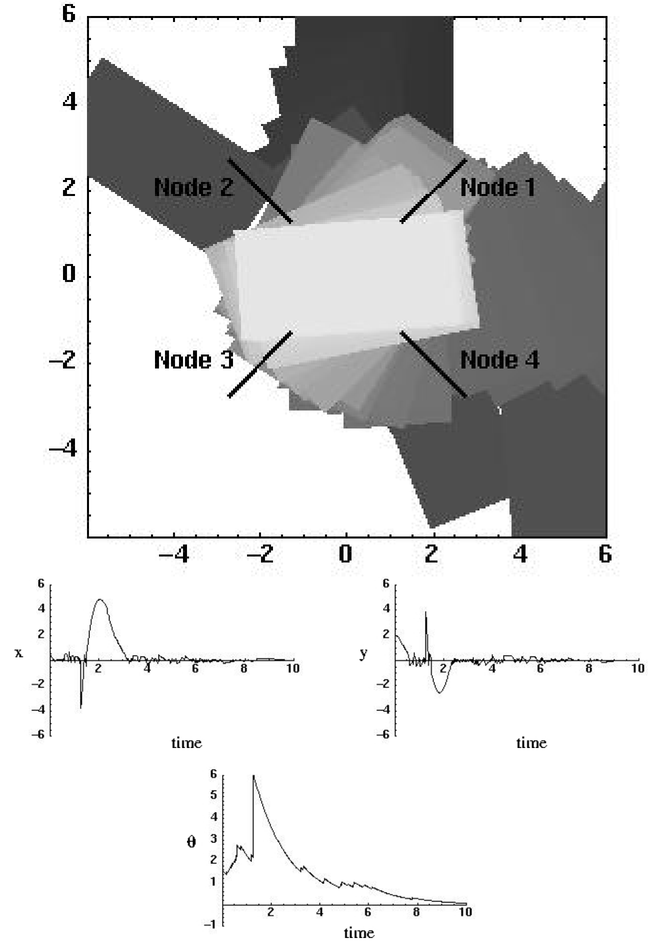

Fig. 4. Simulation of distributed manipulation when the contact state is estimated in some open loop manner but is incorrect. The object is only barely stabilized to the origin due to the contact state being varying from the nominal value. (As in Fig. 3, the bottom three plots are plots of the X,Y, and θ coordinates against time.)

during implementation. Therefore, there are eight control laws, each defined in a separate octant. These control laws can be found in [12].

If these estimated boundaries $0 - 2\pi$ are accurate, then the control laws perform quite well. Figure 3 shows a simulation of the four actuator system. The object is indicated by a rectangle, but the reader should note that although the rectangle is illustrated as being small, the actual body it represents is in contact with all four actuators at all times, which are denoted in the figure by Nodes 1-4. The initial condition is $\{x_0, y_0, \theta_0\} = \{.5, 2, \frac{\pi}{2}\}$, and progress in time is denoted by the lightening of the object. The three plots beneath the XY plot are X, Y, and θ versus time, respectively. This, and the other simulations, were all done in *Mathematica*, using Euler integration in order to avoid numerical singularities when crossing contact state boundaries. In Fig. 3, the object is stabilized to $(0, 0, 0)$ with no difficulty.

In the simulations the constraints are enforced separately from the control law, allowing the control to switch at different times from the constraints. In particular, if the boundary that determines the physical contact state is allowed to vary while the control laws only change at the estimated boundaries $0-2\pi$, then the performance degrades substantially. Starting the object at an initial condition of $\{x_0, y_0, \theta_0\} = \{.5, 2, \frac{\pi}{2}\}$, Fig. 4 shows this degradation in comparison to Fig. 3, although the system is still stable. In the case of Fig. 4, the controller is assuming that the contact state changes when the center of mass of the object crosses the line $x = 0$, whereas the contact state is actually changing when the line $x = -0.3y$ is crossed. This is precisely the difficulty fixed by estimating the contact state on-line, as shown in Section IV.

IV. Supervisory Control and Hybrid Observability

Efforts in the adaptive control community have already created a framework appropriate to addressing the problem of estimating and accommodating changes in contact state.

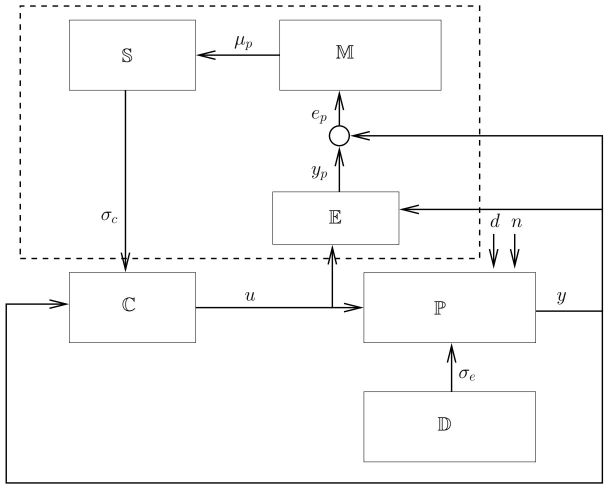

Fig. 5. A supervisory control system

In particular, *supervisory control* (as in [2], [6], [8] and elsewhere) is an effective technique to use when a system is a linear multiple model system. Fortunately, our system, when reduced to a kinematic system using the power dissipation method, is a first order system with constant vector fields. (In fact, not only is it linear, it does not even have drift.) Hence, it is a particularly trivial multiple model system. With little modification, this supervisory framework easily answers how to estimate the current contact state based on the output of the system as well as stabilize the outputs to a desired equilibrium.

The basic idea in supervisory control is that if there is a family (finite or possibly a parameterized continuum) of plants P_σ indexed by σ representing the dynamics, then one can choose controllers appropriate to each P_σ and orchestrate a "switching" between these controllers such that the resulting system is stable. Traditionally, this is a technique where σ is constant but unknown.

Consider the block diagram representation of a supervisory control system found in Fig. 5. Denote the set of possible admissible plants by $\mathbb{P}$. Each model in $\mathbb{P}$ represents a contact state of the overconstrained system. Assume that associated with each plant P_σ coming from $\mathbb{P}$ there is a known stabilizing controller C_σ. Denote the set of these controllers by $\mathbb{C}$. To determine which model in $\mathbb{P}$ most closely "matches" the *actual* model, the input-output relationships for all the plants in $\mathbb{P}$ will need to be estimated. Hence, the need for the *estimator*, denoted by $\mathbb{E}$, which will generate errors between the predicted output for each plant and the actual output of the multiple model system. These errors will then be fed into the *monitoring signal generator*, denoted by $\mathbb{M}$, which will provide monotone increasing signals μ_p determined by

$$\begin{aligned} \dot{W} &= -2\lambda W + \begin{bmatrix} x_{\mathbb{E}} \\ y \end{bmatrix} \begin{bmatrix} x_{\mathbb{E}} \\ y \end{bmatrix}^T, \quad W(0) \geq 0 \\ \mu_p &:= (c_p - 1) W (c_p - 1)^T + \epsilon_\mu, \qquad p \in \mathcal{P} \end{aligned} \tag{4}$$

where $W(t)$ is a symmetric non-negative-definite $k \times k$ matrix with $k = dim(x_{\mathbb{E}}) + 1$, $x_{\mathbb{E}}$ is the state of the estimator, ϵ_μ is a parameter determined by the monitoring signal designer, and c_p is the output one form determining y (from Figure 5) from $y = c_p x_{\mathbb{E}}$. The monitoring signal will be fed into the *switching logic*, denoted by $\mathbb{S}$, which will then determine by means of a *switching signal*, σ_c, which controller to use to control the system output. Call the triple $(\mathbb{S}, \mathbb{M}, \mathbb{E})$ the *supervisor*. Additionally, there is an environmental signal generator $\mathbb{D}$ creating σ_e. $\mathbb{D}$ represents the externally driven switches in contact state that we would like to estimate. Lastly, denote by $N_{\sigma_e}(t_0, t)$ the number of switches σ_e experiences during time $[t_0, t)$.

For our purposes it is sufficient to note that the supervisory control system, as described, is stable so long as 1) each controller C_σ stabilizes its associated plant P_σ, 2) the estimator tracks the contact state well, and 3) the supervisor switches fast enough to ensure convergence without switching so fast as to induce instability. This last requirement is formalized in the following assumption which will turn up in the proof of Proposition 4.1.

Assumption 4.1: Assume σ_e switching is "slow on the average," i.e.,

$$N_{\sigma_e}(t, \tau) \leq N_0^e + \frac{t - \tau}{\tau_{AD}^e}$$

where $N_0^e > 0$ is called the "chatter bound" and τ_{AD}^e is called the "average dwell time."
There is a similarly defined notion of "slow on the average" for the signal σ_c that is used to ensure that the control switching does not destabilize the system [6]. Lastly, let $S_{ave}[\tau_{AD}, N_0]$ be the set of all switching signals for which $N_\sigma(t, \tau) \leq N_0 + \frac{\tau - t}{\tau_{AD}}$.

Now it is possible to prove the following proposition.

Proposition 4.1: Given a linear affine multiple model system where each model is stabilized by a linear control law, for any λ_d (the stability margin) and any d_σ (the time lag induced by the estimator $\mathbb{E}$) there exists a supervisor $\mathbb{S}$ such that with τ_{AD}^e sufficiently large the resulting system is stable with stability margin λ_d.

Proof: First we note that the linear affine multiple model systems combined with their stabilizing controllers may be rewritten as $\dot{x} = A_q x$ where q is an index on the set of models. Hence, the question of stability is whether these systems will destabilize when they switch from one model to another. In the case where $d_\sigma = 0$, then the prior work in [7] directly applies, and we are done. It is when $d_\sigma \neq 0$ that we must modify their approach. In particular, when $d_\sigma \neq 0$ the controller for a given model may be destabilizing for another model. Hence, until the observer accurately tells the supervisor which model is currently governing the system evolution there may indeed be instability in the system. Therefore, in order to proceed with the proof we need the next Lemma (which is an adaptation from [7]), which provides the trade off between the average dwell time and the τ_{AD}^e and the time lag d_σ.

Lemma 4.2: Given two compact sets of $n \times n$ matrices $\mathcal{A} : \{A_p : p \in \mathcal{P}\}$, $\mathcal{A}' : \{A_q' : q \in \mathcal{P}\}$ and a positive constant λ_0 such that $A_p - \lambda_0 I$ is asymptotically stable for each $p \in \mathcal{P}$, then, for any $\lambda \in [0, \lambda_0)$, there is a finite constant τ_{AD}^* and a finite constant d_σ such that if t_i and t_{i+1} are switching times for the switching signal σ:

$$\dot{x} = \begin{cases} A_q' x & on \quad [t_i, t_i + d_\sigma) \\ A_p x & on \quad [t_i + d_\sigma, t_{i+1}) \end{cases} \tag{5}$$

is uniformly exponentially stable over $\mathcal{S}_{ave}[\tau_{AD}, N_0]$ with stability margin λ, for any average dwell time $\tau_{AD} \geq \tau_{AD}^*$ and any chatter bound $0 < N_0 < 1$.

Proof: Assume we have a family of stable plants indexed by $p \in \mathcal{P}$ with matrix representations A_p where $A_p - \lambda_0 I$ asymptotically stable and another set of plants indexed by $q \in \mathcal{Q}$ with matrix representations A_q' (potentially unstable with maximum eigenvalue across all q equal to $\lambda_{A'}$). Then fix τ_{AD}^e, the average dwell time for the signal $\sigma_e(t)$. First it will be proven that for a given dwell time, the Lyapunov functions V_p decrease along trajectories of Eq. (5) for d_σ sufficiently small. Then it will be shown that if the switching from σ_e (which determines τ_{AD}^e) is sufficiently slow, the resulting linear switched system is stable. Let

$$\{t_0, t_1, t_2, \cdots, t_{N_{\sigma_e}(t_0,T)-1}, T\}$$

be the switching times for σ_e. On an interval $[t_i, t_{i+1})$ let d_σ denote the time delay between σ_e switching and σ_c switching. Then choose d_σ such that

$$0 < d_\sigma < \frac{\lambda - \lambda_0}{\lambda_{A'} - \lambda_0}(t_{i+1} - t_i)$$

this implies that

$$\lambda_0(t_{i+1} - t_i - d_\sigma) + \lambda_{A'} d_\sigma < \lambda(t_{i+1} - t_i) \tag{6}$$

which in turn implies that

$$V_p(x(t_{i+1})) < e^{\lambda(t_{i+1} - t_i)} V_p(x(t_i)) \tag{7}$$

Moreover, this is true on any interval, and, because there are only a finite number of switching points, there exists a lower bound on the d_σ (denoted d_*) required to ensure that all the V_p decrease along the trajectories $x(t)$. If d_* is the lower bound, then Eq. (7) holds for all $p \in P$ $q \in Q$ and we have

$$\lambda_0(t_{i+1} - t_i - d_*) + \lambda_{A'} d_* < \lambda(t_{i+1} - t_i)$$

for all i. Recall from [7] that since each stable system represented by A_p has a Lyapunov function V_p, we know that there exists a positive real μ such that for some j $V_i \leq \mu V_j$. Then, following (without modification) the logic in [7], we can see that to get a stability margin of λ_d, we need to satisfy:

$$-\lambda(T - t_0) + N_{\sigma_e}(t_0, T)\log\mu < k - \lambda_d(T - t_0).$$

Due to Eq. (6), this will occur if

$$-\lambda_0(T - t_0 - d_*) + N_{\sigma_e}(t_0, T)(\log\mu - \lambda_{A'} d_*) < k - \lambda_d(T - t_0).$$

for $k > 0$. Now solve for $N_{\sigma_e}(t_0, T)$ and get the following relationship:
$N_{\sigma_e}(t_0, T)$ satisfies $N_0^* := \frac{2k}{\log\mu - \lambda_A d_*}$ and $\tau_{AD}^{e*} := \frac{\log\mu - \lambda_A d_*}{2(\lambda - \lambda_d)}$. This guarantees that

$$\|x(t)\| \leq e^{k - \lambda(T - t_0)}\|x(t_0)\| \tag{8}$$

whenever $\sigma \in S_{\text{ave}}$, $\tau_{AD}^e > \tau_{AD}^{e*}$, and $d_\sigma < d_*$. ■

This Lemma (and N_0^* and τ_{AD}^{e*} in particular) gives us a formal relationship between τ_{AD}^e and d_σ that needs to be satisfied in order for stability to hold. This ends the proof. ■

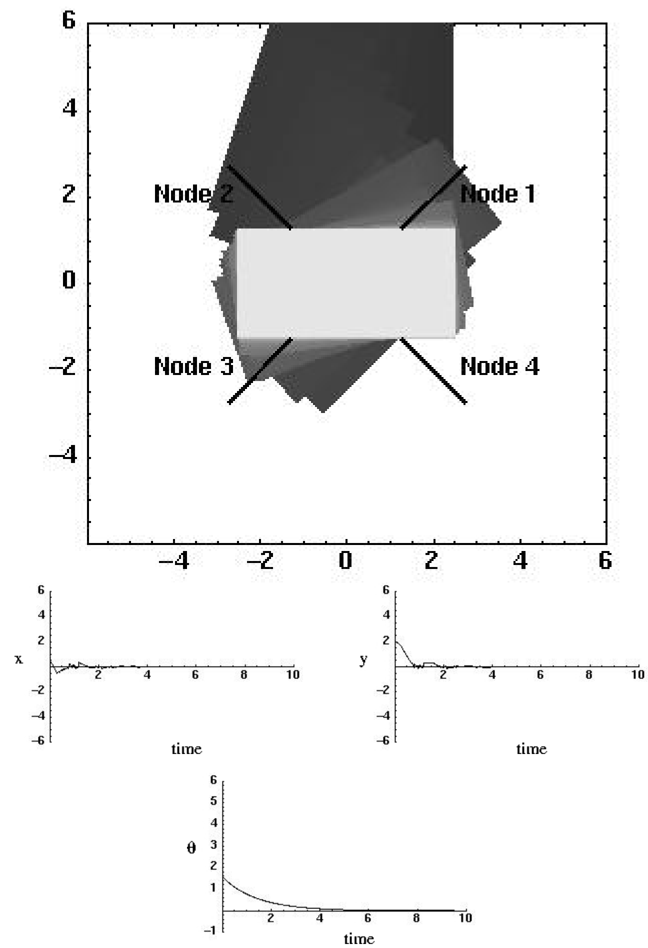

Fig. 6. Simulation of distributed manipulation when the contact state is estimated on-line, using the supervisory control methodology. Here the performance is much closer to that seen in Fig. 3, the case where our knowledge of the state is perfect. (As in Fig. 3, the bottom three plots are plots of the X,Y, and θ coordinates against time.)

Proposition 4.1 indicates that if the contact states change slowly enough (i.e., τ_{AD}^e is large) and supervisory feedback is fast enough (i.e., d_σ is small), then the system can be controlled by a supervisory controller $\mathbb{S}$ using an estimate from an estimator $\mathbb{E}$ that is estimating the contact state on-line. Among other things, this means that one does not have to concern oneself with the friction model to establish where switching occurs. Instead, the contact states can change arbitrarily, so long as they do so sufficiently slowly on the average. Also note that if there is a common Lyapunov function and $d_\sigma = 0$, then $\log \mu = \log 1 = 0$, and the system will be stable for any switching signal. In situations where this is not the case, it would be useful to know if a combination of physical geometry and controller choice can guarantee a lower bound on τ^e_{AD}, but for now we leave it as a standing assumption that it can be bounded. However, one can see that in the case of the example we have been using, the only place where switching can be "fast" is near the origin of the xy plane. Therefore, a purely local analysis should be sufficient to understand conditions for slow switching in σ_e.

Now apply this supervisory approach to the four actuator array from Section III. Replace the boundary $x = 0$ with the boundary $x = -0.3y$, and allow the estimator $\mathbb{E}$ to estimate the contact state and the supervisor $\mathbb{S}$ to orchestrate the controller. In this case (found in Fig. 6) the performance is considerably better than that found in Fig. 4 and resembles the performance found in Fig. 3. However, there are several important characteristics missing from this simulation. First, there is no noise in the output of the system, and it would be useful to know what the sensitivity of this nonsmooth system is to such output noise. Secondly, there is no time delay, which will almost certainly play a substantial role in the dynamics near the origin.

V. Conclusions

In this paper we have introduced the problem of estimating contact states for overconstrained systems and stabilizing these systems. We have offered a solution that is based on adaptive control techniques developed in [2], [6], [8]. The problem of contact state estimation and accommodation is clearly important for systems in which stick/slip phenomena play a dominant role. Indeed, for the distributed manipulation experiment described here, manipulation tasks are actually impossible without the constant trade-off between sticking and slipping. Ultimately, the analytical techniques presented here should be extended to the more geometric setting of grasping and manipulation in the presence of gravitational forces. In the meantime, these results will be implemented on a version of the experiment discussed in Section II.

Despite the validity of the estimation and stabilization techniques presented here, more work must be done to make these methods more computationally efficient. In the supervisory control approach, every model must be integrated forward in time. In the case of the four wheel manipulator there are 8 constraints, leading to 2^8 possible dynamic equations of motion. Utilizing the kinematic reduction found in [11], these can be reduced to 8 total states, a tractable number for the supervisor. However, if the number of actuators is large, or if the system does not satisfy the conditions to be kinematically reducible (as is the case in many relevant grasping problems), then the traditional supervisory approach will surely fail. A potential solution to this is to serially test models rather than testing them in parallel, provided they satisfy some sort of a priori known hierarchy. This method will require the additional characteristic of scaling the gains on the controller down whenever the current model chosen by the supervisor is not predicting the output well. Proving stability of such techniques is non-trivial and is the focus of ongoing study.

References

[1] J.C. Alexander and J.H. Maddocks. On the kinematics of wheeled mobile robots. *The International Journal of Robotics Research*, 8(5):15–27, October 1989.

[2] B.D.O. Anderson, T.S. Brinsmead, F. De Bruyne, J.P. Hespanha, D. Liberzon, and A.S. Morse. Multiple model adaptive control. I. finite controller coverings. *George Zames Special Issue of the Int. J. of Robust and Nonlinear Control*, 10(11-12):909–929, Sep 2000.

[3] K. F. Böhringer, B. R. Donald, L. E. Kavraki, and F. Lamiraux. A distributed, universal device for planar parts feeding: unique part orientation in programmable force fields. In *Distributed Manipulation*, pages 1–28. Kluwer, 2000.

[4] K.F. Böhringer and H. Choset, editors. *Distributed Manipulation*. Kluwer, 2000.

[5] M. Coutinho and P. Will. The intelligent motion surface: a hardware/software tool for the assembly of meso-scale devices. In *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 1997. Albuquerque, New Mexico.

[6] J.P. Hespanha, D. Liberzon, A.S. Morse, B.D.O. Anderson, T.S. Brinsmead, and Franky De Bruyne. Multiple model adaptive control, part 2: Switching. *Int. J. of Robust and Nonlinear Control Special Issue on Hybrid Systems in Control*, 11(5):479–496, April 2001.

[7] J.P. Hespanha and A.S. Morse. Stability of switched systems with average dwell-time. Technical report, EE-Systems, University of Southern California, 1999.

[8] Switching in Systems and Control. *D. Liberzon*. Birkhäuser, Boston, 2003.

[9] J.E. Luntz, W. Messner, and H. Choset. Distributed manipulation using discrete actuator arrays. *Int. J. Robotics Research*, 20(7):553–583, July 2001.

[10] T. D. Murphey and J. W. Burdick. Nonsmooth controllability and an example. In *Proc. IEEE Conf. on Decision and Control (CDC)*, Washington D.C., 2002.

[11] T. D. Murphey and J. W. Burdick. The power dissipation method and kinematic reducibility of multiple model robotic systems. *Transactions on Robotics and Automation*, 2004. Submitted.

[12] T.D. Murphey and J.W. Burdick. Feedback control for distributed manipulation with changing contacts. *International Journal of Robotics Research*, 23(7/8):763–782, July 2004.

[13] R. Howe T. Debus, P. Dupont. Contact state estimation using multiple model estimation and hidden markov models. In *The Eighth Int. Symp. on Experimental Robotics*, Ischia, Italy, July 2002.