



February 1993

Comparative Experiments with a New Adaptive Controller for Robot Arms

Louis L. Whitcomb
Yale University

Daniel E. Koditschek
University of Pennsylvania, kod@seas.upenn.edu

Alfred A. Rizzi
Yale University

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

 Part of the [Electrical and Computer Engineering Commons](#), and the [Systems Engineering Commons](#)

Recommended Citation

Louis L. Whitcomb, Daniel E. Koditschek, and Alfred A. Rizzi, "Comparative Experiments with a New Adaptive Controller for Robot Arms", . February 1993.

Copyright 1993 IEEE. Reprinted from *IEEE Transactions on Robotics and Automation*, Volume 9, Issue 1, February 1993, pages 59-70.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

NOTE: At the time of publication, Daniel Koditschek was affiliated with the University of Michigan. Currently, he is a faculty member of the School of Engineering at the University of Pennsylvania.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ease_papers/325
For more information, please contact repository@pobox.upenn.edu.

Comparative Experiments with a New Adaptive Controller for Robot Arms

Abstract

This paper presents a new model-based adaptive controller and proof of its global asymptotic stability with respect to the standard rigid-body model of robot-arm dynamics. Experimental data from a study of one new and several established globally asymptotically stable adaptive controllers on two very different robot arms 1) demonstrate the superior tracking performance afforded by the model-based algorithms over conventional PD control, 2) demonstrate and compare the superior performance of adaptive model-based algorithms over their nonadaptive counterparts, 3) reconcile several previous contrasting empirical studies, and 4) examine contexts that compromise their advantage.

Disciplines

Electrical and Computer Engineering | Engineering | Systems Engineering

Comments

Copyright 1993 IEEE. Reprinted from *IEEE Transactions on Robotics and Automation*, Volume 9, Issue 1, February 1993, pages 59-70.

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Pennsylvania's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org. By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

NOTE: At the time of publication, Daniel Koditschek was affiliated with the University of Michigan. Currently, he is a faculty member of the School of Engineering at the University of Pennsylvania.

Comparative Experiments with a New Adaptive Controller for Robot Arms

Louis L. Whitcomb, *Member, IEEE*, Alfred A. Rizzi, *Member, IEEE*, and Daniel E. Koditschek, *Member, IEEE*

Abstract—This paper presents a new model-based adaptive controller and proof of its global asymptotic stability with respect to the standard rigid-body model of robot-arm dynamics. Experimental data from a study of one new and several established globally asymptotically stable adaptive controllers on two very different robot arms 1) demonstrate the superior tracking performance afforded by the model-based algorithms over conventional PD control, 2) demonstrate and compare the superior performance of adaptive model-based algorithms over their nonadaptive counterparts, 3) reconcile several previous contrasting empirical studies, and 4) examine contexts that compromise their advantage.

I. INTRODUCTION

SEVERAL years ago, a flurry of activity among robotic control theorists [11], [14], [26], [28], [34], resulted in a new class of adaptive controllers for robot-arm manipulators. These algorithms comprised the first in the literature whose stability could be proven rigorously with respect to the highly nonlinear rigid-body dynamical model. While many of these authors empirically demonstrated significant performance gains over traditional PD controllers, no systematic empirical comparisons *between* the provably correct rigid-body model-based schemes, as applied to various robot plants, seem to have been attempted. Moreover, while these algorithms were typically implemented on high-performance laboratory arms, prior [30] and subsequent [17] authors have argued that infidelities present in industrial robots must vitiate any real benefits from model-based controllers that rely upon the rigid-body assumptions.

The purpose of the present paper is twofold. First, we offer the first reported empirical comparison within this family of closely related but conceptually and algorithmically distinct adaptive controllers. Second, we present a new rigid-body model-based adaptive controller that achieves a slight but po-

tentially significant theoretical advance over past contributions. The experiments were performed both on an industrial SCARA arm and on our new “Yale Bühgler” three-degree-of-freedom (3-DOF) direct-drive juggling robot. Our data corroborate, in part, claims made both by the proponents and the detractors of model-reference adaptive control for robot arms. In particular, they suggest the following:

- 1) The tracking performance of rigid-body model-based controllers is generally superior to conventional PD algorithms.
- 2) Adaptive model-based control algorithms consistently outperform their nonadaptive counterparts. There is only a marginal performance distinction between the various adaptive controllers.
- 3) Model-based algorithms that feedforward reference trajectory information rather than actual state information yield significant performance benefits *when the controller model is valid*; they fail dramatically (in relative terms) when the actuator model is violated (such as actuator torque saturation).
- 4) As has been independently verified, e.g., [1], the degree of performance improvement afforded by all model-based algorithms is strictly limited by the accuracy of the plant model employed.

Of course, as in any other academic paper, there is a methodological “subtext” lurking in this presentation worth articulating directly. We have taken the modest step of presenting an objective measure of tracking performance—root-mean-squared error—in addition to displaying plots of individual joint tracking error over time. We presume that the robotics research community will benefit by adopting some uniform convention for objective and statistically meaningful presentation of tracking data, and we call explicit attention to our own choice in this paper as a means of inviting further public debate concerning the proper standard.

II. THEORY REVIEW

The equations of motion of a mechanical system in local coordinates resulting from application of the Euler–Lagrange operator, to the kinetic energy $\kappa = (1/2)\dot{q}^T M(q)\dot{q}$ of a kinematic chain in the presence of external forces arising from the Earth’s gravitational potential g , and independently controlled torque actuators τ , take the form

$$\begin{aligned}\dot{q}_1 &= q_2 \\ \dot{q}_2 &= -M^{-1}(q_1)[C(q_1, q_2)q_2 + g(q_1) - \tau].\end{aligned}\quad (1)$$

Manuscript received February 21, 1991; revised September 13, 1991. This work was supported in part by SGS Thomson-INMOS Corporation, The Superior Electric Corporation, Fanuc Robotics Corporation, and the National Science Foundation under a Presidential Young Investigator Award held by D. E. Koditschek. Portions of this paper were published in the *Proceedings of the 1991 IEEE International Conference on Robotics and Automation*.

L. L. Whitcomb was with the Department of Electrical Engineering, Yale University, New Haven, CT 06520. He is now with the Department of Mathematical Engineering and Information Physics, University of Tokyo, Bunkyo-ku, Tokyo 113, Japan, and with the Department of Applied Ocean Physics and Engineering, Woods Hole Oceanographic Institution, Woods Hole, MA 02543.

A. A. Rizzi and D. E. Koditschek were with the Department of Electrical Engineering, Yale University, New Haven, CT 06520. They are now with the Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109.

IEEE Log Number 9205083.

Supposing that some reference trajectory $r(t)$ has been specified in advance, we have in our possession as well the additional signals $\dot{r}(t), \ddot{r}(t)$. We assume that all of these are bounded, $\sup_{t \in \mathbb{R}} \|d^{(i)}r/dt^i\| = \rho_i < \infty, i = 0, 1, 2$.

The problem addressed in this paper is the construction of a control law $\tau(t)$ that causes the robot's position to track r asymptotically exactly, that is, $q_1 \rightarrow r$. This section reviews a family of provably correct solutions to the stated problem that have been proposed in the literature and presents a new addition to this family.

A. PD: Fixed Proportional and Derivative Feedback Control

Adopt the standard choice of error coordinates, $e = [e_1^T, e_2^T]^T = [(r - q)^T, (\dot{r} - \dot{q})^T]^T$. The control algorithm against which all controllers are measured is the proportional-derivative linear controller, labeled *PD* in the figures of Section IV, given by $\tau_{pd} = Ke$; $K = [K_1, K_2]$. We assume throughout this paper that $K_i = K_i^T > 0$. With the addition of an (often gain-scheduled) integral term, it is used in nearly every industrial robot available today. While this simple algorithm does not provide asymptotically exact tracking, it does guarantee a bounded error $\|e(t)\|$ when the gain matrices are positive definite and symmetric. Moreover, the steady-state magnitude of $\|e\|$ may be reduced by selecting higher feedback gains [15], [16], [33].

B. ID: Fixed Inverse Dynamics

As of this writing, the most widely familiar algorithm that achieves robot tracking is the "computed torque controller"

$$\tau_{ct} = g(q_1) + C(q_1, q_2)q_2 + M(q_1)[\ddot{r} + \tau_{pd}] \quad (2)$$

resulting in asymptotically stable linear time-invariant error dynamics, and thus asymptotically exact tracking [12], [20]. We shall use, instead of (2), a less well known variation, (3), labeled *ID* in the figures of Section IV, that provides for asymptotically exact tracking *without* exact linearization. We choose this approach because it admits of adaptive extensions that are globally convergent in both state and parameter error, unlike adaptive versions of (2), which have been shown to be globally convergent in plant state error and only locally stable in parameter error and which may require instrumentation of joint acceleration [11]. Consider the control law

$$\tau_{id} = g(q_1) + C(q_1, q_2)\dot{r} + M(q_1)\ddot{r} + \tau_{pd}. \quad (3)$$

With an eye toward subsequent adaptive versions of this algorithm, note that it may be equivalently written in the form

$$\tau_{id} = W(q_1, q_2, \dot{r}, \ddot{r})\theta^* + \tau_{pd} \quad (4)$$

where θ^* is a vector of dynamical parameters (which we still presume to be known at this point in the discussion), and W is an array of known nonlinear functions that depend upon the robot's known kinematics [1], [36]. This control law results in the nonautonomous nonlinear error dynamics

$$\begin{aligned} \dot{e}_1 &= e_2 \\ \dot{e}_2 &= -M^{-1}(q_1)[C(q_1, q_2)e_2 + Ke]. \end{aligned} \quad (5)$$

As in the case of the computed torque algorithm (2), this error system is globally asymptotically stable when $K_i = K_i^T > 0$. However, the demonstration is no longer as straightforward. The (time-varying analogue of) "total energy," $\eta = (1/2)e_1^T K_1 e_1 + (1/2)e_2^T M(q_1)e_2$, has a time derivative along the motions of this system that is negative semidefinite. Unfortunately, this is of no use in consideration of asymptotic stability, since LaSalle's invariance principle does not apply to nonautonomous systems. A complete stability argument for this error system is given in Subsection B of the Appendix.

C. IDA: A New Adaptive Controller

The stability proof of the fixed model-based controller (3) in turn affords for the first time a proof for the following adaptive version:

$$\tau_{ida} = W(q, \dot{q}, \dot{r}, \ddot{r})\hat{\theta} + \tau_{pd} \quad (6)$$

$$\dot{\hat{\theta}} = K_g W^T [e_2 + \epsilon(e_1)e_1]; \quad \epsilon(e_1) = \epsilon_0(1 + \|e_1\|)^{-1} \quad (7)$$

(derived in Subsection C of the Appendix), which is both globally asymptotically stable in plant tracking error and globally stable in controller parameter error.

D. IDC and IDCA: "Critically Damped" Inverse Dynamics

Several years ago it was observed that the solutions to the closed-loop system arising from the controller

$$\tau_{idc} = W(q_1, q_2, \dot{r}', \ddot{r}')\theta^* + \tau_{pd}, \quad \dot{r}' = \dot{r} + \Lambda e_1 \quad (8)$$

labeled *IDC* in the following sections, and the plant (1) converge, in an \mathcal{L}^2 sense, to the stable first-order subspace $e_2 = -\Lambda e_1$. Since an exponentially stable system forced by an input that decays to zero has an output that decays to zero, it follows that $e \rightarrow 0$. This useful observation was reported independently, first by Slotine and Li [28] and subsequently by Horowitz and Sadegh [26], leading to the first provably correct differentiator free adaptive controller for a robot arm, (9) and (10), labeled *IDCA* in the figures to follow.

$$\tau_{idca} = W(q, \dot{q}, \dot{r}', \ddot{r}')\hat{\theta} + \tau_{pd} \quad (9)$$

$$\dot{\hat{\theta}} = K_g W(q, \dot{q}, \dot{r}', \ddot{r}')^T K_2^{-1} K e. \quad (10)$$

E. IDR and IDRA: Pure Feedforward Versions

Several researchers [27], [33] have observed that a substitution of reference information for state information in the feedforward portion of (3)

$$\tau_{idr} = W(r, \dot{r}, \ddot{r})\theta^* + \tau_{pd} + \sigma_n \|e\|^2 K_2^{-1} K e \quad (11)$$

can still provide asymptotically exact tracking. This algorithm will be denoted *IDR* in the figures below. Moreover, a pure feedforward version of the adaptive algorithm (9), denoted *IDRA* in the figures below, has been presented in [27].

$$\tau_{idra} = W(r, \dot{r}, \ddot{r})\hat{\theta} + \tau_{pd} + \sigma_n \|e_1\|^2 K_2^{-1} K e \quad (12)$$

$$\dot{\hat{\theta}} = K_g W(r, \dot{r}, \ddot{r})^T K_2^{-1} K e \quad (13)$$




Fig. 1. (a) The GMFanuc model A-500 and (b) the Yale Bühgler.

III. EXPERIMENTAL SETUP

A. Hardware

We will study the performance of the controllers implemented on two different robots. The first, the Yale-GMFanuc A-500 Industrial Arm, is a classical industrial manipulator. The second, the Yale Bühgler, was designed to support our research program in robot juggling [24].

The GMFanuc Robotics Model A-500, a 4-DOF SCARA-type arm, is shown in Fig. 1. The manufacturer's control system was replaced with our own custom system. The first two revolute degrees of freedom were used in these experiments. The "elbow joint" (joint 1) is driven by a GMFanuc Model 1-0 three-phase dc brushless motor and a 47:1 spiroidal gear box. It is representative of the highly geared drive systems used in most conventional industrial robots. The "shoulder joint" (joint 0), is direct-drive driven by a 14-in NSK direct-drive variable-reluctance (VR) motor capable of delivering a static torque of 250 N·m.

The Yale Bühgler Arm, a direct-drive 3-DOF *essentially spherical* robot arm, is depicted in Fig. 1. Each joint is driven directly by a VR motor manufactured by the Superior Electric Corporation. These are high-performance units whose low-velocity peak-torque and relatively high torque-to-mass ratio makes them particularly suited to direct-drive design.

The computational hardware for these implementations is the Yale XP/DCS [19], a distributed real-time controller based upon the SGS-Thomson INMOS Transputer floating-point microprocessor [24].

B. Software

It has long been known [5] that the mathematical representation of the controllers explored here hides to a considerable degree their rather startling computational complexity. Indeed, many experimenters use approximations of reduced complexity either by making simplifying assumptions about the structure of link inertia tensors [7] or by using (theoretically unjustified) approximations to the exact dynamics expression [18]. It is well known, e.g., [6], that the computational complexity of evaluating the rigid-body robot dynamical model varies linearly with n , the degrees of freedom. For example, [10] reports a complexity of $126n - 99$ multiplications and $106n - 92$ additions.

We employ the *exact* Lagrangian dynamical equations for fully general link inertial tensors (including the off-diagonal terms), without omission or approximation of a single term. The equations were generated by a program¹ written for the symbolic mathematics environment Mathematica, though any home-made or commercially available general derivation program would suffice. In either case, with a derivation utility in hand, its application to any particular robot is a straightforward and nowadays a commonplace exercise. The input to these derivation programs is a file containing the kinematic (four per link) and dynamic (ten per link) parameters of a robot.² The symbolic derivations for the A-500 and Bühgler each take less than 25 s of CPU time on a Sparc-1 workstation. The output is an explicit closed-form expression for the equations listed above, from which computer source code is automatically generated.

The A-500 control laws are all completely evaluated at 1 kHz. The Bühgler control laws are all completely evaluated at two time scales—the feedback terms at 1 kHz and the model-based terms at 400 Hz. Sampling issues in the discrete time distributed control of robots are explored in [24] and [35].

C. Reference Trajectories, Feedback Gains, and Magic Parameters

It is always possible to demonstrate the "superiority" of a favorite control algorithm by contriving an appropriately clever example. The practicing control engineer, however, is justifiably skeptical of anecdotal special-case examples that may not accurately represent typical overall performance. We have endeavored to demonstrate typical overall performance, rather than special case examples, by adopting the following rules of procedure.

The various reference trajectories were selected to exhibit both "slow" friction-dominated behavior and "fast" rigid-body dynamics dominated behavior, and still lie within the actuators' torque saturation limits. To this end we employed sinusoidal joint-space reference trajectories of the form $r_i(t) = \text{offset} + \text{magnitude} \times \sin(\text{phase} + \omega \times t)$. A great many (offset, magnitude, phase, ω) combinations were used in the experiments to explore the different regions of behavior.

The feedback gain matrices used in all controllers were identical: they were selected empirically to give an approximately critically damped response to the individual joints when in independent motion. Note that the feedback gain matrices employed (identical between controllers) were considerably lower than the limit dictated by the usual experimental technique of increasing gains as high as possible (to the verge of instability). Unfortunately, since robot manipulators constitute highly nonlinear plants, gain settings optimized for one reference trajectory may well result in unstable performance for another. We were interested in *comparing the relative* performance of the different controllers in an unbiased fashion over a wide range of performance regimes, and accordingly did *not* push gains to the verge of instability to obtain the

¹ Available from the authors.

² We use, without loss of correctness, the commonly accepted technique of employing a smaller dimensional set of base parameters, e.g., [21].

smallest tracking error magnitude for each reference trajectory. Higher feedback gains were observed (of course) to provide uniformly smaller steady-state tracking errors but identical relative performance between the various controllers. In short, the gain margin is very much a function of the reference trajectory in such nonlinear systems, and we chose low enough values to permit experimentation with identical gains over the entire performance regime examined.

The adaptive gain matrix K_g can, according to theory, be any symmetric positive definite matrix. In practice, the numerical integration of the adaptive law destabilizes the entire system for sufficiently high adaptation gains. In these experiments, we set the adaptation gains to be as large a multiple of I as possible while preserving stability. It would be useful to have a more complete theoretical understanding and an automatic procedure to accomplish the manual fine-tuning of adaptation gains reported in previous implementations [13].

The designer also must choose initial values for the adaptive controller model parameters for each run. Except where noted, we have in all cases initialized the adaptive model parameters to those values used in the fixed (nonadaptive) controllers. It is worth noting in passing that fixed parameter values were, in turn, obtained by running an adaptive algorithm for a short period of time from wildly inaccurate (zero) initial parameter estimates as discussed below.

Included (except where noted) in the model-based controller implementations, though omitted for clarity from the equations of Section II, is a Coulomb and viscous friction compensation term [32] for each joint.

IV. EXPERIMENTAL RESULTS

The overall conclusion to be derived from these experiments is best summarized by Fig. 2. This plot depicts the mean and variance of root-mean-square position error norms (for all joints) achieved by each of the seven controllers described in Section II. The ensemble of runs over which these descriptive statistics are gathered comprise ten very different reference trajectories—differing not only in frequency content but in the region and volume of joint space they encompass. The results are normalized for convenience with respect to the simple PD controller since all physical significance of the joint angle errors is compromised by the diversity of trajectories being compared. Thus, the plot displays in succinct form the experimental data that justify the broadly stated conclusions in the introduction to this paper.

It is clear from Fig. 2 as well as the other plots given below that the model-based controllers offer far better tracking performance than the PD in almost every case and that the adaptive model-based controllers (as a group) outperform the nonadaptive controllers. In general, the performance of each model-based controller is improved roughly 50% by its adaptive counterpart as shown in Fig. 2. Furthermore, there is a nearly identical performance ranking within both the nonadaptive and adaptive controller groups. We will examine those conclusions in greater detail and with reference to more selective experiments in this section.

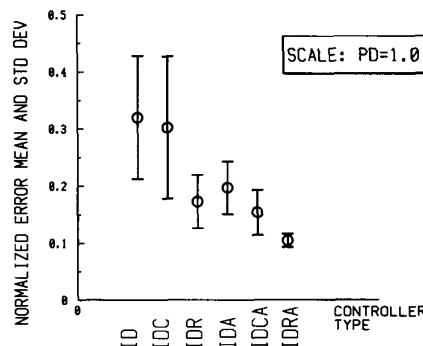


Fig. 2. Bühgler normalized joints 0, 1, and 2 \mathcal{L}^2 position error norm ensemble mean and standard deviation over (a sample size of) ten different reference trajectories, repeated for each of seven controllers.

A. Data Presentation

Since there is, at present, no general nonlinear counterpart to classical linear systems performance summaries (such as Bode plots), it has become accepted practice in the robotics community to compare controller performance by the visual examination of tracking error curves as a function of time for a “representative” or “standard” reference trajectory. Fig. 5 (given later), for example, shows the position tracking error of three controllers, as a function of time for joints 0 and 1 of the A-500 robot when both joints were tracking a sinusoidal reference trajectory. We wish to compare controller performance for no less than seven controllers over a variety of reference trajectories. While the curves in Fig. 5 provide a palpable representation of tracking performance, the visual comparison of a succession of such graphs quickly becomes an act of aesthetic judgment rather than empirical analysis.

As an alternative, we have employed the scalar-valued \mathcal{L}^2 norm as an objective numerical measure of tracking performance for an entire error curve. The \mathcal{L}^2 norm is given by $\mathcal{L}^2[e(t)] = ((1/t) \int_0^t \|e(t)\|^2 dt)^{1/2}$ where $e(t)$ is a selected scalar (or vector) valued tracking error. The norm measures the root-mean-square “average” of the tracking error. A smaller \mathcal{L}^2 norm represents smaller tracking error—and thus better performance.

Assertions based on experimental results are meaningful only to the extent that the results may be reproduced reliably. To test this, we ran and computed the error norm average and standard-deviation over ten runs with identical controllers, plants, and reference trajectories for each of the listed controllers. The standard-deviation of the error norm over the ten runs was observed to be typically less than 1% of their mean [36].

B. Performance Benefits Due to the Adaptive Algorithms

Fig. 3 shows PD position tracking error norm ($\mathcal{L}^2[e(t)]$; $e(t) = [e_1(t), e_2(t), e_3(t)]^T$) of between 13° and 25° , IDR errors of 3° to 6° , and IDRA errors of roughly 2° over a range of reference trajectories. The figure shows the \mathcal{L}^2 norm of the Bühgler joint position tracking error³ vector at steady

³ $1^\circ = 0.0174$ rad.




Fig. 3. Bühgler joints 0, 1, and 2 \mathcal{L}^2 position error norm (radians) versus reference trajectory nominal frequency.

state for each controller. The reference trajectory for each joint was a sinusoid differing slightly from a nominal frequency. The error norms are plotted at three different nominal reference trajectory nominal frequencies. The frequency range was chosen to include slow friction-dominated operation at one end to dynamics-dominated operation at the other. In this plot the slowest frequency corresponds to peak gripper velocity of 0.5 m/s, and the highest frequency corresponds to peak gripper speeds of 3.0 m/s. These data represent a typical instance of the statistical trends reflected in Fig. 2. The model-based controllers provide tracking performance superior (smaller error norm) to conventional PD control at equal feedback gains; their adaptive counterparts perform still better.

Among the model-based controllers (both fixed and adaptive), the IDR and IDRA (see Section II-E) controllers that utilize reference trajectory values in their plant model are uniformly outperforming the controllers that use sensor values in their plant model. In Section IV-D-2 we will demonstrate that the advantages thus gained are not without peril.

The IDC and IDCA controllers were observed to marginally outperform the ID and IDA controllers, respectively. This consistent difference is discussed in Section IV-D-1. Finally, adaptive controllers were observed to be less robust than the nonadaptive controllers in the presence of certain unmodeled effects such as link vibration modes, actuator saturation, numerical integration, and the like, which may occur when the reference trajectories exceed the system's design capability.

C. The Effect of Parameter Values

It is commonly agreed that effective nonadaptive model-based control relies on the availability of "correct" model parameter values. A common misconception, however, is that "any model is better than none"—that an "approximately correct" parameter set will result in better tracking than that obtained by PD control alone.

1) *Incorrect Parameters:* Fig. 4 shows the \mathcal{L}^2 error norm for all joints of the Bühgler obtained with the same reference trajectories, controllers, and initial parameters as Fig. 3. However, in this run the absence of the former "gripper payload," an L-shaped steel bar now removed from the distal link, significantly degraded the performance of both the PD and (via parameter mismatch) nonadaptive model-based controllers.




Fig. 4. Incorrect parameters: Bühgler joints 0, 1, and 2 \mathcal{L}^2 position error norm (radians) versus reference trajectory nominal frequency.

The adaptive controllers, in contrast, compensate automatically for the change. Not surprisingly, instances were also observed where incorrect parameter values resulted in *poorer* fixed model-based controller performance than that of PD.

The sensitivity of nonadaptive model-based algorithms to incorrect parameter values is doubly problematic because these parameters, representing the link inertia tensors, are difficult to measure by hand. Off-diagonal inertia tensor elements (products of inertia) are sufficiently difficult to measure that even the most capable investigators typically do not attempt it [3]. In contrast, the adaptive controllers estimate all terms of the parameterization—including those arising from the off-diagonal inertia tensor elements—with indifference⁴.

2) *Zero Initial Parameter Values:* An important special case of "incorrect parameter values" arises when the designer has no knowledge of the plant parameters and must rely on adaptation for a "cold start." Fig. 5 shows the tracking performance of the PD, IDR (using "correct" fixed parameters from a previously converged adaptive run), and IDRA using zero initial parameter values. While the IDRA is seen to have a large initial transient, it recovers almost immediately to outperform PD. Within 10 s it is already performing nearly as well as the "correct parameter" fixed IDR controller.

3) *Obtaining "Correct" Fixed Parameters:* A fair comparison between the fixed and adaptive model-based algorithms is complicated by the issue of where to obtain the necessary parameter estimates for the former class. We observed that parameter sets produced by the adaptive controllers, when used in the fixed controllers, provided performance superior to the dismal performance obtained using hand-measured physical parameters.

On the other hand, when the fixed model controllers are given parameters resulting from their adaptive counterparts' convergence over a long run, and their performance is compared with respect to exactly the same reference trajectory, then it is not surprising that the adaptive controllers perform little better if at all. Yet, since parameters "optimally tuned" for one reference trajectory are in general "suboptimal" with respect to any other, the fixed controllers always perform less

⁴These direct adaptive controllers provide for asymptotically exact reference tracking and stability of parameters but do not promise (without additional conditions) convergence of controller parameters to the "true" plant values.

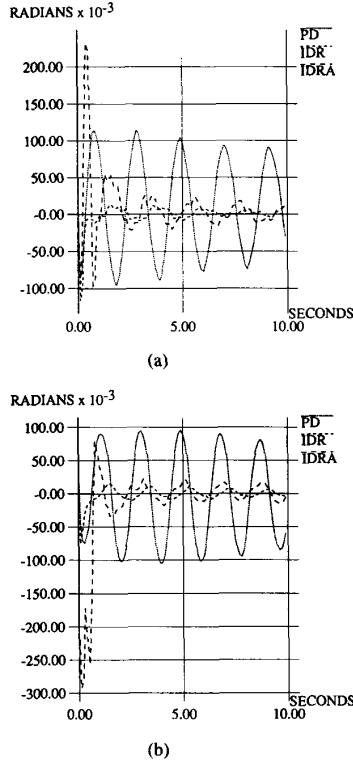


Fig. 5. A-500 tracking error versus time: (a) joint 0 and (b) joint 1.

capably than their adaptive counterparts in any other context. This is reflected not merely in the lower means of Fig. 2 but in the comparatively smaller variance of the adaptive algorithms' performance relative to their fixed parameter counterparts.

The (theoretically illicit) sensitivity of parameters to reference trajectory type as well as the capacity of the adaptive algorithms to adjust from type to type with little transient error will be examined in Section IV-E. In general, we have tried to show fixed algorithms at their best (smallest tracking error) by choosing parameters resulting from previous adaptive runs on similar (but not identical) trajectories.

D. Ranking the Model-Based Controllers

The marginal ranking between the three fixed model-based controllers (ID, IDC, and IDR) and the nearly identical ranking between their adaptive counterparts (IDA, IDCA, and IDRA) was observed to hold under a variety of reference trajectories (Fig. 2) and for both plants (Section IV-F). Section IV-D-1 shows the tracking performance of ID and IDC and their adaptive versions to be nearly identical—the apparent performance distinction between them a consequence of their differing feedback structure. Section IV-D-2 suggests reasons for the reference controllers' (IDR, IDRA) superior performance and examines a context in which their performance is compromised.

1) *Performance Differences: ID and IDC:* In the statistical plot, Fig. 2, it is clear that the IDC has a slightly lower mean error than the ID algorithm and that the adaptive versions

of these algorithms display the same pattern. The apparent distinction between ID and IDC performance turns out to be an unintended consequence of exactly the phenomenon we had hoped to avoid—higher *effective* feedback gains resulting from the variant structure of the IDC and IDCA controllers—by retaining the same set of uniform PD gain matrices ($K_2, K_1 = K_2\Lambda$) over each comparative run.

To see this, recall that the ID controller (3) is written

$$\tau_{id} = W(q_1, q_2, \dot{r}, \ddot{r})\theta^* + \underbrace{K_2}_{\text{derivative gain}} e_2 + \underbrace{K_1}_{\text{proportional gain}} e_1$$

and the IDC controller, (8) is written

$$\begin{aligned} \tau_{idc} &= W(q_1, q_2, \dot{r}', \ddot{r}')\theta^* + \tau_{pd}, \quad \dot{r}' = \dot{r} + \Lambda e_1 \\ &= W(q_1, q_2, \dot{r}, \ddot{r})\theta^* + K_2 K e + M \Lambda e_2 + C \Lambda e_1 \\ &= W(q_1, q_2, \dot{r}, \ddot{r})\theta^* + \underbrace{[K_2 + M \Lambda]}_{\text{effective derivative gain}} e_2 + \underbrace{[(K_2 + C) \Lambda]}_{\text{effective proportional gain}} e_1. \end{aligned}$$

When, as in all the Bühler experiments reported above, Λ is a multiple of the identity matrix, then $M\Lambda$ is positive definite and the effective derivative gain for the IDC and IDCA algorithms $[K_2 + M\Lambda]$ exceeds that of ID and IDA $[K_2]$. Moreover, while C and, hence, $C\Lambda$, have an indefinite symmetric part, the values attained by C in the course of these runs are often small; hence, the effective proportional gain for the IDC and IDCA algorithms $[(K_2 + C)\Lambda]$ is only insignificantly less than that of the ID and IDA ($[K_1]$ where $K_1 = K_2\Lambda$). In all the data discussed above, we set $K_1 = \text{diag}\{-50, -25, -5\}$, $K_2 = \text{diag}\{-10, -5, -1\}$, and $\Lambda = \text{diag}\{5, 5, 5\}$.

For example, consider the point in state space $q_1 = [0, 0, 0]^T$ and $q_2 = [1, 1, 1]^T$. Here, the IDC effective derivative gain matrix $[K_2 + M\Lambda]$ has eigenvalues 150% greater than ID's derivative gain K_2 . The IDC effective proportional gain matrix $[(K_2 + C)\Lambda]$ is essentially identical to ID's $K_1 = K_2\Lambda$. Increasing the ID derivative gain by 150% to make it more equivalent to IDC's "effective" gain was observed to close the performance gap [36]. Conversely, "detuned" values of Λ (e.g., unequal elements on the diagonal) that can result in the symmetric part of $M\Lambda$ being indefinite (and thus lowering the effective derivative gain of IDC) result in poorer relative performance of IDC in comparison to ID [36]. Unfortunately, because ID's feedback structure is fixed and IDC's is (in part) time varying, there is no constant set of feedback gain matrices that will allow them to be compared exactly "evenly" out of context.

2) Reference Trajectory Feedforward Algorithm Performance:

It has been noted that the model-based controllers IDR (11) and IDRA (12), using reference trajectory signals in place of actual ("exact") sensed signals in model computation [27], [33], provided generally superior tracking performance relative to their exact counterparts. In addition to enabling (in some applications) a reduction of on-line computation, the trajectory algorithms substitute a "clean" reference velocity for the inherently noisy sensor-derived velocity in the feedforward

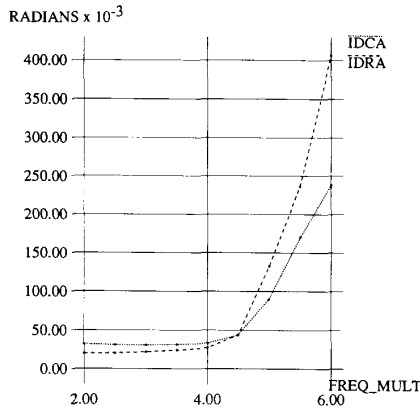


Fig. 6. The effect of actuator saturation on IDCA and IDRA: Bühler joints 0, 1, and 2 \mathcal{L}^2 position error norm (radians) versus reference trajectory nominal frequency.

and adaptation laws. The latter property might account for their superior tracking performance.

To test the second assertion, we smoothed the (numerically differentiated) velocity signal with a first-order linear filter. We observed a slight performance improvement for filter time-constants up to 3 ms and performance degradation for higher values. On the whole, however, our data were inconclusive on this point. This suggests that elucidating the reference algorithms' advantage requires a more careful and systematic investigation. Ongoing work on an observer theory for Lagrangian mechanical systems [8] may yield insight to this important point.

All model-based controllers may fail in varying degrees due to mismatch between the dynamic model and plant caused by unmodeled dynamics, torque saturation, actuator dynamics, friction, and the like. A classic example where such a mismatch is likely to occur is the potentially unbounded joint torque commands resulting from a (seemingly reasonable) workspace trajectory near a kinematic singularity. It is therefore essential to investigate the behavior of the various control algorithms under the pervasive condition of torque saturation.

Fig. 6 shows the position tracking error norms for the IDCA and IDRA controllers as a function of the nominal frequency of the sinusoidal reference trajectory. The actuator torque limits were reduced to 20% of their previous values. At lower frequencies (with correspondingly low velocities and torques), where saturation is absent, the IDRA outperforms the IDCA controller. At the highest frequencies, where saturation is almost continual, we see the following remarkable differences: 1) The absolute magnitude of the tracking error is greater for both controllers, as is expected in the presence of actuator saturation. 2) The IDRA (12) controller does not perform as well as the exact model-based controller. This may be attributed to the significant modeling error introduced in the IDRA algorithms by the large tracking error. 3) At slightly higher reference trajectory frequencies than those plotted, the IDRA algorithm becomes unstable while the IDCA algorithm remains stable. Indeed, the stability proofs for *all* of the adaptive controllers depend intimately on instantaneous and unbounded torque

capabilities—we have seen all of the adaptive controllers become unstable when sufficiently compromised.

E. "Long-Term Memory" Effects

Recently, a fundamentally different set of "learning" techniques—neural networks [13], memory-based learning [4], and repetitive learning [2], [22] methods—have challenged the hegemony of Lagrangian model-based methods in robust controller design. The principal advantage of the "learning" control algorithms would be the promise of accurately controlling enormously complicated plants without explicitly modeling the plant's underlying dynamics. Their disadvantage is the need to repetitively learn each unique task—they are unable to apply the knowledge of "learned" parameters to any but the original task. While much of the early work in this area was heuristic, recent results directly address stability and robustness [2], [22], thus establishing some of these techniques as theoretically sound alternatives to model-based adaptive robot control.

In contrast, the adaptive robot controllers' "learning" processes (parameter convergence) occur simultaneously with task execution, obviating the need for a separate "learning" phase. Moreover, after achieving parameter convergence, they can (in theory) apply this knowledge to track *any* smooth reference trajectory. In practice, however, these advantages are compromised by the following effects. First, adaptive parameter convergence relies on richness properties of signals within the system [23] that commonplace workplace tasks may fail to produce. Second, we observe in practice that adaptive parameters converge to slightly different "optimal" values for differing reference trajectories rather than converging to a single value for all rich trajectories, and they exhibit (theoretically disallowed) transients when transitioning from one reference to another.

We tried a variety of reference trajectory combinations to find examples evoking a transient response in which the IDR controller outperformed its adaptive counterpart. The response to one such trajectory, Fig. 7, shows the instantaneous position error norm for the smooth transition from extremely slow reference sinusoids (friction dominates), on the left, to fast sinusoids, on the right (dynamics dominates). The IDRA adaptive parameter drift was sufficient to produce a larger transient excursion than the IDR controller. Within a few seconds after the transition, the IDRA controller recovers to equal the performance of the IDR controller, and it is superior at steady state.

Do model-based controllers offer "generically" superior performance in practice? Fig. 8 shows the tracking error norms and standard deviations for the Bühler arm, computed on 10-s intervals, for 10 different reference trajectory combinations. In each run a reference trajectory "switch" occurred at $t = 60$ s, as may be observed from the rise in the mean and variance spread in the plot. These data confirm the validity of the anecdotal observations above. The model-based controllers outperform PD, and the adaptive controllers outperform the fixed controllers. The standard-deviation separation demonstrates that, over a variety of reference trajectories, the ordering

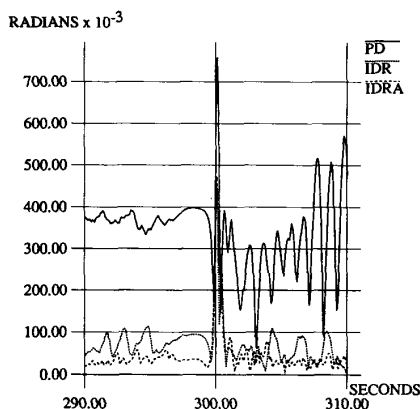


Fig. 7. Long-term memory effects: Bühler joints 0, 1, and 2 \mathcal{L}^2 instantaneous position error norm (radians) versus time (seconds). Smooth reference "switch" at $t = 300$.




Fig. 8. Bühler joints 0, 1, and 2 \mathcal{L}^2 position error norm ensemble mean and (+/-) one standard deviation of ten different reference trajectories computed on 10-s intervals versus time for three controllers. Smooth reference "switch" at $t = 60$.

of their tracking error means is both *average* as well as (in a statistical sense) *typical*.

We conclude that the performance of model-based, and in particular adaptive model-based, controllers is not seriously compromised by the imperfect parameter convergence mentioned above. The adaptive controller usually outperforms both its nonadaptive counterpart as well as PD control; the worst case "defects" of the adaptive algorithms (brief transients) are relatively innocuous in comparison to its demonstrated advantages over all of the fixed controllers.

F. Contrast between Direct-Drive and Geared Joint Performance

1) *Geared Joints Require Friction Models:* Fig. 9 shows the \mathcal{L}^2 norm of the (gear-driven) joint 1 position tracking error at steady state for the A-500 robot under seven controllers at each of three different reference trajectory frequencies both with (Fig. 9(a)) and without (Fig. 9(b)) friction compensation. The dramatically superior performance of the friction-compensated controllers over both the PD and (nonfriction-compensated) model-based controllers corroborates previously




Fig. 9. Geared joint: A-500 joint 1 \mathcal{L}^2 position tracking error (radians) versus reference trajectory nominal frequency: (a) with and (b) without friction compensation.

reported results, e.g., [9], [32]. Thus, the dynamics of the geared joint (joint 1) appears dominated by frictional forces instead of link inertial forces. Indeed, this agrees with the well known fact that a link's contribution to the total inertia seen by a joint motor of gear ratio n is scaled by a ratio of $1/n^2$.

The spiroidal gear box used in joint 1 of the A-500 is, in fact, a dynamically simple mechanism with only two moving parts and providing high stiffness. Given the significant performance degradation that arises from even this relatively simple (in comparison to, for example, a harmonic drive) geared actuator, it is not surprising that other researchers have observed more curious performance defects in robots with more complicated actuator systems, e.g., [17].

In contrast, Fig. 10 shows the corresponding error norms for the (direct-drive) joint 0. This plot demonstrates only a marginal performance improvement between the friction-compensated and nonfriction-compensated controllers at higher velocities and (not surprisingly) shows significant improvement only at low velocities. Supported by similar results for the direct-drive Bühler Arm, we conclude that the rigid-body dynamics appears to govern the direct-drive joint (joint 0). Friction compensation appears to offer improved performance for these types of direct-drive joints only at low velocities.

2) *Geared Joints Do Not Always Require Lagrangian Models:* The ubiquity of industrial gear-driven manipulators suggests the possible utility of adding only friction-compensation terms to PD controllers and ignoring the Lagrangian dynamical model. Indeed, several motor and robot manufacturers claim to have implemented similar friction-compensation features, which we shall term PDF. Such a




Fig. 10. Direct-drive joint: A-500 joint 0 L^2 position tracking error (radians) versus reference trajectory nominal frequency: (a) with and (b) without friction compensation.

controller might provide superior tracking performance for geared actuators at relatively modest computational cost. Fig. 11 shows the tracking error norms of the A-500 for the PDF controller in addition to the original seven controllers.

Fig. 11(a) shows that tracking error norms for (direct-drive) joint 0 for each of the eight controllers. Here, the PDF controller provides improved performance (over PD) only at low velocities and no significant improvement at high reference velocities. Fig. 11(b) shows the PDF tracking error for (gear-drive) joint 1 to be essentially on par with the full model-based controllers (which incorporate both friction and rigid-body models). This and the lackluster geared-joint performance of model-based controllers without friction compensation (Fig. 9(b)) confirm that friction dynamics dominates the rigid-body dynamics of this geared joint.

V. CONCLUSION

This paper has reviewed the stability literature for a class of model reference parameter adaptive controllers for robot-arm manipulators based upon the ID (3) variant of the popular computed torque algorithm (2). It provides for the first time a rigorous and global stability proof for IDA (6), a member of this class that has heretofore eluded a complete analysis. Comparative experiments of all these variants have been performed on a standard industrial SCARA manipulator and a fast direct-drive robot arm developed at the Yale Robotics Laboratory. The highlights of the observations of Section IV having been previewed in the introduction of the paper, we will only briefly summarize and amplify here.




Fig. 11. A-500 joint L^2 position tracking error norm (radians) versus reference trajectory nominal frequency: (a) direct-drive joint 0 and (b) geared joint 1.

A. Summary of Results to Date

Fixed model-based controllers dramatically outperform the PD controller and their adaptive counterparts perform still better. Thus, if a designer is committed to a computed torque-like controller, since the computationally intensive W (4) terms must be computed anyway (the parameter adaptation integrals represent very minor additional computational burden), the adaptive variant should be preferred. Most importantly, adaptation is the *easiest, fastest, and most accurate* method for obtaining parameter values for use in fixed-parameter controllers.

The degree of performance improvement afforded by model-based algorithms is strictly limited by the accuracy of the plant model employed. We conclude that an appropriate dynamical model is a more important antecedent to superior tracking than the particular model-based control law employed. Unfortunately, given the huge variety of available actuators and torque-amplifiers, there is no single "right" model for model-based control. The practicing engineer must carefully match controller and actuator capabilities in the context of a desired application—when better models are incorporated, performance benefits are immediate.

We (cautiously) conclude that the worries of learning theorists [4], [25] about the "long-term memory" effects of model-based adaptive methods appear to be unjustified. Namely, after parameters have adapted to yield good tracking in the face of a particular class of reference trajectories, new trajectories may incur substantial transient error—but the transients are typically no worse than those incurred by fixed controllers, and the steady-state performance is superior. The matter deserves continued careful examination.

B. Future Experiments

As usual, these preliminary experiments have quickly led to further questions that suggest some interesting near-term extensions to the empirical work presented here. The extent to which we can improve the performance of conventional industrial robot manipulators seems limited by the validity of the controller's actuator models. Thus, the development of such models (e.g., [29]) is of considerable practical importance. The ubiquity of digital joint position encoders in actual machines suggests the utility of the recent work on observer theory for Lagrangian dynamical systems [8]. We presume the (theoretically arbitrary) adaptive gain matrix (see (7), (10), and (13)) to be of untapped practical usefulness. Finally, to the best of our knowledge, a systematic experimental investigation of the various provably correct "learning" algorithms in comparison to model-based adaptive algorithms has never been undertaken.

APPENDIX A NEW ADAPTIVE CONTROLLER

Two years ago, Koditschek [16] presented a new strict global Lyapunov function for general mechanical systems that includes in its general purview the ID error dynamics (5). The specialization of the general idea to the present case amounts to nothing more than a modification of the bilinear cross term of the strict (but local) Lyapunov function [15], [31], [33], as will be seen below. Using this new Lyapunov function, we derive in this appendix an adaptive law to accompany the IDA controller (6) that yields global asymptotic stability in the state and parameter errors. A more leisurely derivation of these results is given in [36], to which the reader is referred for further details.

A. Notation and Terminology

We will require a notation for induced operator norms of constant-, linear-, and bilinear-operator-valued functions on a vector space. For example, let C be a map from some space \mathcal{J} into the set of bilinear operators on the vector space product $\mathcal{X} \times \mathcal{Y}$. Then the upper norm of C is defined to be

$$\|C\| \triangleq \sup_{q \in \mathcal{J}} \sup_{\substack{\|x\|=1 \\ x \in \mathcal{X}}} \sup_{\substack{\|y\|=1 \\ y \in \mathcal{Y}}} \|C(q, x)y\|$$

and the lower norm $\|C\|$ is defined analogously. A slight abuse of notation permits the same notation regardless of whether the domain of C is \mathcal{Y} , $\mathcal{X} \times \mathcal{Y}$, or $\mathcal{J} \times \mathcal{X} \times \mathcal{Y}$.

B. A New Lyapunov Function

Now consider the modified Lyapunov candidate

$$\vartheta \triangleq \eta + \epsilon(e_1)e_1^T M(q_1)e_2 = \frac{1}{2}e^T \begin{bmatrix} K_1 & \epsilon M \\ \epsilon M & M \end{bmatrix} e$$

where

$$\epsilon(e_1) \triangleq \epsilon_0(1 + \|e_1\|)^{-1}.$$

Note that a sufficiently small choice of ϵ_0 , for example,

$$\epsilon_0 \leq (\|K_1\| \|M\|)^{\frac{1}{2}} / \|M\|$$

guarantees that this is a positive definite function with respect to the error coordinates e_1, e_2 for any time varying trajectory $q(t)$.

Taking the derivative of ϑ along a motion of the error system (5), we have

$$\dot{\vartheta} = \dot{\eta} + \epsilon e_2^T M e_2 - \epsilon e_1^T [C e_2 + K_2 e_2 + K_1 e_1] + \epsilon e_1^T \dot{M} e_2 + \dot{\epsilon} e_1^T M e_2.$$

Defining, for any vector v , $\dot{M}_q(v) \triangleq (v^T \otimes I) D_q M^s(q)$, we have

$$\begin{aligned} \epsilon e_1^T (\dot{M} - C) e_2 &= \frac{1}{2} \epsilon (e_1) e_1^T \left[\dot{M}(q_2) e_2 - \dot{M}(e_2) q_2 \right. \\ &\quad \left. - \dot{M}(q_2)^T e_2 \right] \\ \left| \epsilon e_1^T (\dot{M} - C) \right| &\leq (3\epsilon(e_1)/2) \|\dot{M}\| \|e_1\| \cdot \|e_2\| \cdot \|q_2\| \\ &\leq (3\epsilon(e_1)/2) \|\dot{M}\| \|e_1\| \cdot \|e_2\| \cdot (\|e_2\| + \rho_1) \\ &\leq (3\epsilon_0/2) \|\dot{M}\| \|e_2\|^2 + \\ &\quad (3\rho_1\epsilon/2) \|\dot{M}\| \|e_1\| \cdot \|e_2\| \end{aligned}$$

where ρ_1 is defined in Section II. Moreover, we have

$$\begin{aligned} \dot{\epsilon} e_1^T M e_2 &= -\frac{\epsilon_0}{\|e_1\|(1 + \|e_1\|)^2} \cdot e_1^T e_2 \cdot e_1^T M e_2 \\ |\dot{\epsilon} e_1^T M e_2| &\leq \frac{\epsilon_0 \|\dot{M}\| \|e_1\|^2 \cdot \|e_2\|^2}{\|e_1\|(1 + \|e_1\|)^2} \leq \epsilon \|\dot{M}\| \|e_2\|^2. \end{aligned}$$

It now follows that

$$\begin{aligned} \dot{\vartheta} &\leq -\frac{1}{2} e_2^T (K_2 - \epsilon M) e_2 + \frac{1}{2} (\epsilon \|M\| + 3\epsilon_0 \|\dot{M}\|) \|e_2\|^2 \\ &\quad - \frac{1}{2} e_2^T K_2 e_2 - \epsilon e_1^T K_1 e_1 + (3\rho_1\epsilon/2) \|\dot{M}\| \|e_1\| \cdot \|e_2\| \\ &\quad + \epsilon e_1^T K_2 e_2 \\ &\leq -\frac{1}{2} [\|K_2\| - \epsilon_0 (2\|M\| + \frac{3}{2}\|\dot{M}\|)] \|e_2\|^2 \\ &\quad - e^T \begin{bmatrix} \epsilon \|K_1\| & \frac{1}{2} \epsilon (\|K_2\| + \frac{3}{2}\rho_1 \|\dot{M}\|) \\ \frac{1}{2} \epsilon (\|K_2\| + \frac{3}{2}\rho_1 \|\dot{M}\|) & \frac{1}{2} \|K_2\| \end{bmatrix} e \end{aligned}$$

is negative definite as long as

$$\epsilon_0 < \min \left[\frac{\|K_2\|}{4\|M\| + 3\|\dot{M}\|}, \frac{4\|K_1\| \|K_2\|}{(\|K_2\| + 3\rho_1 \|\dot{M}\|)^2} \right].$$

In this case we have

$$\begin{aligned} \dot{\vartheta} &\leq -\epsilon \|Q\| \|e\|^2 \\ Q &\triangleq \begin{bmatrix} \|K_1\| & \frac{1}{2} (\|K_2\| + \frac{3}{2}\rho_1 \|\dot{M}\|) \\ \frac{1}{2} (\|K_2\| + \frac{3}{2}\rho_1 \|\dot{M}\|) & \|K_2\|/2\epsilon_0 \end{bmatrix}. \end{aligned} \quad (A1)$$

C. Stability of the New Algorithm

The IDA controller (6) results in the error system

$$\begin{aligned} \dot{e}_1 &= e_2 \\ \dot{e}_2 &= -M^{-1}(q_1)[C e_2 + K e + W \hat{\theta}] \end{aligned} \quad (A2)$$

where $\hat{\theta} \triangleq \hat{\theta}^* - \hat{\theta}$ denotes the parameter error vector whose adjustment over time must now be established in such a fashion that $e \rightarrow 0$ as if $\hat{\theta}^*$ were known. Define the adaptive

law to be as in (7). The scalar valued function $v \triangleq \dot{\theta} + \frac{1}{2}\tilde{\theta}^T K_g^{-1} \tilde{\theta}$ has a derivative along the motion of the full adaptive system

$$\begin{aligned} \dot{v} = & -e_2^T (K_2 e_2 + W\tilde{\theta}) + \epsilon e_2^T M e_2 - \epsilon e_1^T [C e_2 + K e + W\tilde{\theta}] \\ & + \epsilon e_1^T \dot{M} e_2 + \epsilon e_1^T M e_2 + \tilde{\theta}^T K_g^{-1} \dot{\tilde{\theta}} \\ \leq & -\epsilon \|Q\| \|e\|^2 - [e_2 + \epsilon e_1] W \tilde{\theta} + \tilde{\theta}^T K_g^{-1} \dot{\tilde{\theta}} \leq -\epsilon \|Q\| \|e\|^2 \end{aligned}$$

that is nonpositive. It follows that v is bounded, hence, $\sqrt{\epsilon} \|e\|$ is an \mathcal{L}^2 function [23]. But an \mathcal{L}^2 function whose derivative is bounded must tend to zero [23]; hence, $\|e\| \rightarrow 0$ as desired.

ACKNOWLEDGMENT

Prof. M. Bühler of McGill University was responsible in large measure for the original conception and design of the Yale "Bühler" Direct Drive Robot. M. Aiken of the Superior Electric Corporation and Prof. P. Kindlmann of the Yale Department of Electrical Engineering provided crucial support in the design and integration the Bühler Variable-Reluctance actuators and control electronics. An anonymous referee provided an unusually cogent and helpful review. We thank them for their invaluable advice and assistance.


REFERENCES

- [1] C. H. An, C. G. Atkeson, and J. M. Hollerbach, *Model-Based Control of a Robot Manipulator*. Cambridge, MA: MIT Press, 1988.
- [2] S. Arimoto, "Robustness of learning control for robot manipulators," in *Proc. IEEE Int. Conf. Robotics Automat.* (Cincinnati, OH), May 1990, pp. 1528-1533.
- [3] B. Armstrong, O. Khatib, and J. Burdick, "The explicit dynamic model and inertial parameters of the puma 560 arm," in *Proc. IEEE Int. Conf. Robotics Automat.* (San Francisco), Apr. 1986, pp. 510-518.
- [4] C. G. Atkeson, "Memory-based approaches to approximating continuous functions," in *Proc. 6th Yale Workshop Adaptive Learning Syst.*, (New Haven, CT), 1990, pp. 212-217.
- [5] A. K. Bejczy, "Robot arm dynamics and control," Jet Propulsion Lab., Pasadena, CA, Tech. Rep. 33-699, 1974.
- [6] M. Brady, J. Hollerbach, T. Johnson, T. Lozano-Perez, and M. Mason, Eds., *Robot Motion: Planning and Control*. Cambridge, MA: MIT Press, 1982.
- [7] J. Burdick, "An algorithm for generation of efficient manipulator dynamic equations," in *Proc. IEEE Int. Conf. Robotics Automat.* (San Francisco), Apr. 1986, pp. 212-218.
- [8] C. Canudas de Wit and N. Fixot, "Adaptive control of robot manipulators via velocity estimated feedback," in *Proc. IEEE Int. Conf. Robotics Automat.* (Sacramento, CA), 1991, pp. 16-21.
- [9] C. Canudas de Wit, P. Noël, A. Aubin, and B. Brogliato, "Adaptive friction compensation for robot manipulators: Low velocities," *Int. J. Robotics Res.*, vol. 10, no. 3, pp. 189-199, June 1991.
- [10] J. J. Craig, *Introduction to Robotics Mechanics and Control*. Reading, MA: Addison-Wesley, 1986.
- [11] J. J. Craig, P. Hsu, and S. Sastry, "Adaptive control of mechanical manipulators," Electronics Res. Lab., Univ. of California, Berkeley, CA, Memo. M86/3, Jan. 1986.
- [12] E. Freund, "Fast nonlinear control with arbitrary pole placement for industrial robots and manipulators," *Int. J. Robotics Res.*, vol. 1, no. 1, pp. 65-78, 1983.
- [13] M. A. Johnson and M. B. Leahy Jr., "Adaptive model-based neural network control," in *Proc. IEEE Int. Conf. Robotics Automat.* (Cincinnati, OH), May 1990, pp. 1704-1709.
- [14] D. E. Koditschek, "Adaptive techniques for mechanical systems," in *5th Yale Workshop Applicat. Adaptive Syst. Theory*, (New Haven, CT), May 1987, pp. 259-265.
- [15] —, "High gain feedback and telerobotic tracking," in *Workshop Space Telerobotics* (Jet Propulsion Lab., Pasadena, CA), Jan. 1987, pp. 355-363.
- [16] —, "Strict global Lyapunov functions for mechanical systems," in *Proc. Amer. Contr. Conf.* (Atlanta, GA), June 1988, pp. 1770-1775.
- [17] M. B. Leahy Jr., D. E. Bossert, and P. V. Whalen, "Robust model-based control: An experimental case study," in *Proc. IEEE Int. Conf. Robotics Automat.* (Cincinnati, OH), May 1990, pp. 1982-1987.
- [18] M. B. Leahy, Jr., L. M. Nugent, K. P. Valavanis, and G. N. Saridis, "Efficient dynamics for a Puma-600," in *Proc. IEEE Int. Conf. Robotics Automat.* (San Francisco), 1986, pp. 519-524.
- [19] F. Levin, M. Bühler, L. Whitcomb, and D. E. Koditschek, "Transputer computer juggles real-time robotics," *Electron. Syst. Design*, vol. 19, no. 2, pp. 77-82, Feb. 1989.
- [20] J. Y. S. Luh, M. W. Walker, and R. P. Paul, "Resolved acceleration control of mechanical manipulators," *IEEE Trans. Automat. Contr.*, vol. AC-25, pp. 468-474, 1980.
- [21] H. Mayeda, K. Yoshida, and K. Oshun, "Base parameters of manipulator dynamic models," *IEEE Trans. Robotics Automat.*, vol. 6, no. 3, pp. 312-321, June 1990.
- [22] W. Messner, W. Kao, R. Horowitz, and M. Boals, "A new adaptive learning rule," in *Proc. IEEE Int. Conf. Robotics Automat.* (Cincinnati, OH), May 1990, pp. 1522-1527.
- [23] K. S. Narendra and A. Annaswamy, *Stable Adaptive Systems*. Englewood Cliffs, NJ: Prentice-Hall, 1988.
- [24] A. A. Rizzi, L. L. Whitcomb, and D. E. Koditschek, "Distributed real-time control of a spatial robot juggler," *Computer Mag.*, vol. 25, no. 5, May 1992.
- [25] M. W. Roth, "Neurobiological learning, Markov models, and the long term memory problem for artificial neural networks," in *Proc. 6th Yale Workshop Adaptive Learning Syst.* (New Haven, CT), 1990, pp. 110-119.
- [26] N. Sadegh and R. Horowitz, "Stability analysis of an adaptive controller for robotic manipulators," in *Proc. IEEE Int. Conf. Robotics Automat.* (Raleigh, NC), Apr. 1987.
- [27] —, "Stability and robustness analysis of a class of adaptive controllers for robotic manipulators," *Int. J. Robotics Res.*, vol. 9, no. 3, pp. 74-92, June 1990.
- [28] J.-J. E. Slotine and W. Li, "On the adaptive control of robot manipulators," in *Proc. ASME Winter Annu. Meeting* (Anaheim, CA), Dec. 1986.
- [29] M. W. Spong, "Adaptive control of flexible joint manipulators," in *Proc. IEEE Int. Conf. Robotics Automat.* (Scottsdale, AZ), May 1989, pp. 1188-1193.
- [30] L. M. Sweet and M. C. Good, "Redefinition of the robot motion-control problem," *IEEE Contr. Syst. Mag.*, vol. 5, no. 3, pp. 18-25, Aug. 1985.
- [31] M. Takegaki and S. Arimoto, "Stability and robustness of pid feedback control for robot manipulators of sensory capability," in *Robotics Res., Proc. 1st Int. Symp.*. Cambridge, MA: MIT Press, May 1984.
- [32] T. Tarn, A. Bejczy, X. Yun, and Z. Li, "Effect of motor dynamics on nonlinear feedback robot arm control," *IEEE Trans. Robotics Automat.*, vol. 7, no. 1, pp. 114-122, Feb. 1991.
- [33] J. T. Wen and D. S. Bayard, "Robust control for robotic manipulators, Part I: Non-adaptive case," Jet Propulsion Lab., Pasadena, CA, Tech. Rep. 347-87-203, 1987.
- [34] —, "Robust control for robotic manipulators, Part II: Adaptive case," Jet Propulsion Lab., Pasadena, CA, Tech. Rep. 347-87-204, 1987.
- [35] L. L. Whitcomb and D. E. Koditschek, "Robot control in a message passing environment," in *Proc. IEEE Int. Conf. Robotics Automat.* (Cincinnati, OH), May 1990, pp. 1198-1203.
- [36] L. L. Whitcomb, A. A. Rizzi, and D. E. Koditschek, "Comparative experiments with a new adaptive controller for robot arms," Yale Univ., Center for Syst. Sci., New Haven, CT, Tech. Rep. 9101, Jan. 1991.



Louis L. Whitcomb received the B.S. degree in mechanical engineering in 1984 and the Ph.D. degree in electrical engineering in 1992, both from Yale University, New Haven, CT. His dissertation work, under Professor Daniel Koditschek, comprised the theory and implementation of distributed computational architectures for real-time robot control; the theory, implementation, and extensive experimentation with adaptive control of robot manipulators; and automated assembly via artificial potential functions.


From 1984 to 1986, he was a member of the research and development staff at the GMFanuc Robotics Corporation. He is presently completing a post-doctoral fellowship at the Woods Hole Oceanographic Institution, Woods Hole, MA, as well as an NSF/JSPS post-doctoral fellowship at the University of Tokyo, Tokyo, Japan.



Alfred A. Rizzi (M'89) received the Sc.B. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1986 and the M.Sc. degree in electrical engineering from Yale University, New Haven, CT. He is currently working toward the Ph.D. degree at Yale.

From 1986 to 1988, he worked as a designer in the aerospace industry. In addition to his doctoral studies, he is a Research Associate at Yale. His current interests include the study of robot systems capable of dynamical dexterous manipulation, and

the theory and implementation of distributed real-time robot controllers.



Daniel E. Koditschek (S'80-M'83) received the B.S. degree in 1977 in engineering and applied science and the Ph.D. degree in 1983 in electrical engineering from Yale University, New Haven, CT.

He was appointed Assistant Professor of Electrical Engineering at Yale in 1984 and is presently an Associate Professor in that department. His interests include the application of dynamical systems theory to autonomous machine design, nonlinear control theory, and the application of computational theory

and hardware to feedback control of physical processes.