Adaptive Manipulator Control: A Case Study

JEAN-JACQUES E. SLOTINE, MEMBER, IEEE, AND WEIPING LI

Abstract—**Adaptive control of linear time-invariant single-input single-output systems has been extensively studied, and a number of globally convergent controllers have been derived. Extensions of the results to nonlinear or multivariable systems have rarely been achieved. Yet, in the case of robot manipulators, which represent an important and unique class of nonlinear, time-varying, multiiput, multioutput dynamic systems, similar global convergence properties can indeed be obtained. Our earlier work [19] exploits the particular structure of manipulator dynamics to develop a simple, globally convergent adaptive controller for manipulator trajectory control problems. This paper, after summarizing the basic algorithm, demonstrates the approach on a high-speed two degree-of-freedom semi-direct-drive robot. It shows that the dynamic parameters of the manipulator, assumed to be initially unknown, can be estimated within the first half second of a typical run, and that accordingly, the manipulator trajectory can be precisely controlled. Furthermore, these experimental results demonstrate that the adaptive controller enjoys essentially the same level of robustness to unmodeled dynamics as a PD controller, yet achieves much better tracking accuracy than either PD or computed-torque schemes. Its superior performance for high-speed operations, in the presence of parametric and nonparametric uncertainties, and its relative computational simplicity, make it an attractive option both to address complex industrial tasks, and to simplify high-level programming of more standard operations.**

I. INTRODUCTION

THE development of effective adaptive controllers represents an important step towards versatile applications of high-speed and high-precision robots. Even in a well-structured industrial setting like the elusive, if proverbial, "factory of the future," robots still have to face uncertainty on the parameters describing the dynamic properties of the grasped load, such as moments of inertia or exact position of the center of mass in the end-effector. Since these parameters are difficult to compute or measure for geometrically complex objects, they limit the potential for robots to accurately manipulate objects of size and weights similar to their own, as the human arm routinely does. It is widely recognized that the accuracy of conventional approaches (such as the computed-torque method) in high-speed operations is greatly affected by the parameter uncertainties. This sensitivity is especially severe for direct-drive robots, for which no gear reduction is available. Two classes of approaches are being actively studied to maintain the performance of the manipulators in the presence of parameter uncertainties: robust control (e.g., [18], [6], [7], [23]) and adaptive control. An advantage of the adaptive approach is that the accuracy of a manipulator carrying unknown loads improves with time (because the adaptation mechanism keeps extracting parameter information from tracking errors), so that adaptive controllers potentially hold the promise of consistent performance in the face of very large load variations.

Adaptive control, however, as a branch of systems theory, is not yet quite mature, and the strong nonlinearity of robot dynamics makes them more complex to analyze than the linear dynamic systems on which most of the existing adaptive control theory has been traditionally focused. While many adaptive robot controllers have been proposed in the literature, most of them have to rely on assumptions or approximations such as local linearization, time-invariant, or decoupled-dynamics (see, e.g., [8], for a recent review) to guarantee their tracking convergence. Craig *et al.* [4] formulate a globally convergent adaptive controller which does not make these approximations, but their method requires numerical differentiation of the joint velocities to obtain estimates of joint accelerations (which tends to be quite noisy) and inversion of the estimated inertia matrix (which is computationally-intensive, and further requires ad-hoc procedures to ensure that the matrix remains invertible during the adaptation process). In [19] a new joint-space adaptive tracking control algorithm is presented, which consists of a PD feedback part and an adaptive full dynamics compensation part, with the unknown manipulator and payload parameters being estimated on-line. The algorithm guarantees global stability of the system and asymptotic convergence of the tracking errors, and further avoids both the aforementioned approximations and the need for acceleration measurements or inversion of the estimated inertia matrix. The algorithm is computationally simple, due to an effective exploitation of the structure of manipulator dynamics, and in particular of the natural relationship between the inertia matrix and the Coriolis and centripetal terms. Later work [20] extends the adaptive controller to Cartesian space motion, hybrid force/position control, and external control of unknown passive mechanisms. The adaptive controller in this paper is a direct adaptive controller, in the sense that parameter adaptation is driven by the motion tracking error. Our recent studies have indicated that an indirect approach [14], whose parameter estimation is driven by the prediction error of robot joint torque, and a composite approach [22], whose parameter adaptation is driven by both the motion tracking error and the torque prediction error, can also be used to guarantee globally convergent joint tracking. However, the indirect and composite approaches involve more computation than the direct approach.

Theoretical analysis and computer simulations of an adaptive controller are important but not sufficient, as reflected by the many discussions of the past few years concerning robustness issues (e.g., [17], [2]). Since inherent factors such as unmodeled high-frequency dynamics and measurement noise are generally neglected in the stability analysis, we believe that the ultimate justification for the value and applicability of an adaptive controller lies in its actual hardware implementation. Based on this perspective, this paper examines the joint-space adaptive tracking controller of [19] experimentally on a high-speed two degree-of-freedom semi-direct-drive robot. It shows that all the manipulator mass properties, assumed to be initially unknown, can be quickly estimated within the first half second of a typical run, and that the algorithm allows large loads of unknown mass properties to be precisely manipulated in the face of various nonparametric error sources.

Section II provides the theoretical background for the experiments, with Section II-B summarizing the formulation of the adaptive trajectory control approach of [19], and Section II-C discussing implementation aspects and the incorporation of sliding control terms into the basic algorithm for the purposes of

Manuscript received March 11, 1987; revised December 21, 1987. Paper recommended by Associate Editor, H. P. Geering. This work was supported in part by a grant from the Sloan Fund.

The authors are with the Nonlinear Systems Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139.

IEEE Log Number 8823320.

0018-9286/88/1100-0995$01.00

computational efficiency and enhanced robustness with respect to disturbances [22]. The manipulator and the experimental setup used in the tests are described in Section III. The experimental results are detailed in Section IV. Section V offers brief concluding remarks.

II. Manipulator Model

A. Manipulator Model

In the absence of friction and other disturbances, the dynamics of a rigid manipulator can be written as

$$H(q)\ddot{q}+C(q,\dot{q})\dot{q}+g(q)=\tau \tag{1}$$

where q is the $n \times 1$ vector of joint displacements, τ is the $n \times 1$ vector of applied joint torques (or forces), $H(q)$ is the $n \times n$ symmetric positive definite manipulator inertia matrix, $C(q, \dot{q})\dot{q}$ is the $n \times 1$ vector of centripetal and Coriolis torques, and $g(q)$ is the $n \times 1$ vector of gravitational torques.

Two simplifying properties should be noted about the above dynamic structure. First, as remarked by several authors (e.g., [1] and [11]), the two $n \times n$ matrices H and C are not independent. Specifically, given a proper definition of the matrix C (note that the centripetal and Coriolis torque vector $C\dot{q}$ is uniquely defined, but that the matrix C is not), the matrix $(\dot{H}-2C)$ is skew-symmetric, a property which can be easily derived from the Lagrangian formulation of the manipulator dynamics [22] and which reflects conservation of energy. This property can also be written

$$\dot{H}=C+C^T$$

since $\dot{H}$ is symmetric. The second important property is that the individual terms on the left-hand side of (1), and therefore the whole dynamics, are *linear* in terms of a *suitably selected set of equivalent manipulator and load parameters* [10], [3], as illustrated in Section III-B for a two-link manipulator.

B. Controller Design

The adaptive controller design problem is as follows: given the desired joint position $q_d(t)$, and with some or all the manipulator parameters unknown, derive a control law for the actuator torques, and an estimation law for the unknown parameters, such that the manipulator joint position $q(t)$ precisely tracks $q_d(t)$ after an initial adaptation process.

Let a be a constant m-dimensional vector containing the unknown elements in the suitably selected set of equivalent dynamic parameters, let $\hat{a}$ be its (time-varying) estimate, and let $\hat{H}$, $\hat{C}$, and $\hat{g}$ be the matrices obtained from the matrices H, C, and g by substituting the estimated $\hat{a}$ for the actual a. Then the linear parametrizability of the dynamics enables us to write

$$\hat{H}(q)\ddot{q}_r+\hat{C}(q,\dot{q})\dot{q}_r+\hat{g}(q)=Y(q,\dot{q},\dot{q}_r,\ddot{q}_r)\tilde{a} \tag{2}$$

where $\tilde{a}=\hat{a}-a$ is the parameter estimation error, Y is an $n \times m$ matrix independent of the dynamic parameters, and $\dot{q}_r$ is defined as

$$\dot{q}_r=\dot{q}_d-\Lambda\tilde{q}$$

with Λ being a positive definite matrix, and $\tilde{q}(t)=q(t)-q_d(t)$ denoting the position tracking error. The vector $\dot{q}_r$ formed by modifying the desired velocity $\dot{q}_d$ using the position error $\tilde{q}$, may be called "reference velocity," and is introduced to guarantee the convergence of the tracking error, as we shall see later. Intuitively, the reference velocity $\dot{q}_r$ increases if the actual trajectory q lags behind the desired q_d.

The following choice of control and adaptation laws were suggested in [19]:

$$\tau=\hat{H}\ddot{q}_r+\hat{C}(q,\dot{q})\dot{q}_r+\hat{g}(q)-K_Ds \tag{3}$$

$$\dot{\hat{a}}=-\Gamma Y^Ts \tag{4}$$

where Γ is a constant positive definite matrix, $K_D(t)$ is a uniformly positive definite matrix, and the vector s, which can be thought of as a measure of tracking accuracy, is defined as

$$s=\dot{q}-\dot{q}_r=\dot{\tilde{q}}+\Lambda\tilde{q}. \tag{5}$$

The above control and adaptation laws guarantee the global convergence of the positional and velocity tracking errors, as long as the desired trajectories q_d, $\dot{q}_d$, and $\ddot{q}_d$ are bounded. To prove this, we consider the Lyapunov function candidate

$$V(t)=\frac{1}{2}[s^THs+\tilde{a}^T\Gamma^{-1}\tilde{a}]. \tag{6}$$

The differentiation of $V(t)$ leads to

$$\begin{aligned}\dot{V}(t)&=s^T(H\ddot{q}-H\ddot{q}_r)+\tilde{a}^T\Gamma^{-1}\dot{\hat{a}}+(1/2)s^T\dot{H}s\\&=s^T(\tau-H\ddot{q}_r-C\dot{q}_r-g)+\tilde{a}^T\Gamma^{-1}\dot{\hat{a}}\end{aligned}$$

where the skew-symmetry of $(\dot{H}-2C)$ has been used to eliminate the term $(1/2)s^T\dot{H}s$ which reflects the time-varying nature of the inertia matrix. Substituting the control law (3) into the above expression, and using the linearity property (2), we obtain

$$\begin{aligned}\dot{V}(t)&=s^T(\tilde{H}\ddot{q}_r+\tilde{C}\dot{q}_r+\tilde{g}-K_Ds)+\tilde{a}^T\Gamma^{-1}\dot{\hat{a}}\\&=s^T[Y\tilde{a}-K_Ds]+\tilde{a}^T\Gamma^{-1}\dot{\hat{a}}=-s^TK_Ds+\tilde{a}^T[\Gamma^{-1}\dot{\hat{a}}+Y^Ts].\end{aligned}$$

The reason for the choice (4) of adaptation law is to remove the second term from the last expression. Indeed, substituting (4) leads to

$$\dot{V}(t)=-s^TK_Ds\leq 0. \tag{7}$$

Since $V(t)$ is lower bounded by zero and decreases for any nonzero s, as seen from (7), it seems plausible from the above equation that $\dot{V}(t)$, and therefore [from (7)] the tracking error measure s, must converge to zero. The strict mathematical proof of this result, based on showing the uniform continuity of $\dot{V}(t)$, is detailed in [22]. It can then be easily shown, from definition (5), that the convergence of s to zero in turn guarantees that $\tilde{q}$ and $\dot{\tilde{q}}$ also converge to zero. Intuitively, this corresponds to the fact that the output of a stable linear filter, whose input converges to zero, must also converge to zero. Therefore, both global stability of the system, and convergence of the tracking error, are guaranteed by the above adaptive controller.

The structure of the adaptive controller given by (3) and (4) is sketched in Fig. 1. The controller consists of two parts. The first part is a special form of full dynamics compensation, with three terms corresponding to inertial, centripetal and Coriolis, and gravitational torques. This part, based on the estimated parameters, attempts to provide the joint dynamic torques necessary to make the desired motions. The second part actually contains two terms representing PD feedback, since

$$-K_Ds=-K_D\dot{\tilde{q}}-K_D\Lambda\tilde{q}.$$

It intends to regulate the real trajectories about the desired trajectories. The required inputs to the controller are the desired joint position q_d, velocity $\dot{q}_d$, and acceleration $\ddot{q}_d$. The required measurements are the joint position q and velocity $\dot{q}$.

Remark 1: The closed-loop dynamics of the manipulator under adaptive control can be easily obtained by substituting the control law (3) into the system dynamics (1). Since (3) can be expressed

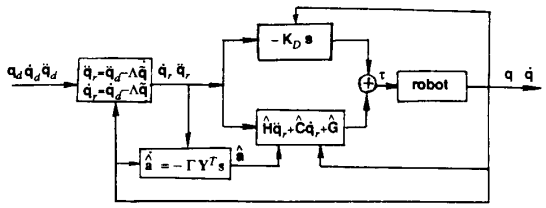

Fig. 1. The structure of the adaptive controller.

as

$$\tau = Y\hat{a} - K_D s$$

the closed-loop dynamics can be written in the simple form

$$H\dot{s} + (K_D + C)s = Y\tilde{a} \tag{8}$$

with $\tilde{a}$ determined by the adaptation law (4). Furthermore, note that the adaptation law can be expressed as

$$\dot{\hat{a}} = -\Gamma \frac{\partial \tau}{\partial \hat{a}} s.$$

This form is reminiscent of the intuitive "M.I.T. rule" in model-reference adaptive control, with the matrix Γ playing the role of adaptation gain. In the present adaptive controller, the magnitude of Γ does not affect the global stability of the system (as long as unmodeled dynamics are not excited), but it directly conditions the speed of adaptation, and therefore the system's performance.

Remark 2: The guaranteed convergence of the tracking errors to zero does not imply the convergence of the estimated parameters to the exact values. It is shown in [21] that the estimated parameters asymptotically converge to the true parameters if the matrix $Y(q_d, \dot{q}_d, \ddot{q}_d)$ is persistently exciting and uniformly continuous. By persistent excitation, we mean that there exist positive constants δ, α_1, and α_2 such that for all $t_1 \geq 0$

$$\alpha_1 I \leq \int_{t_1}^{t_1+\delta} Y_d^T Y_d \, dt \leq \alpha_2 I$$

where $Y_d = Y(q_d, \dot{q}_d, \ddot{q}_d, \dddot{q}_d)$, and I is the $m \times m$ identity matrix. It is interesting to note that this result, for the nonlinear robot dynamics, is similar to those for linear dynamics obtained by [16], although their derivation does not apply to the particular form (4), (8).

Remark 3: Besides constant positive-definite matrices, a natural choice for K_D is $K_D(t) = \lambda\hat{H}$, with λ a strictly positive constant. Indeed, with the associated choice $\Lambda = \lambda I$, control law (3) then becomes

$$\tau = \hat{H}[\ddot{q}_d - 2\lambda\tilde{\dot{q}} - \lambda^2\tilde{q}] + \hat{C}\dot{q}_r + \hat{g}$$

a form very close to a "computed-torque controller" with critically-damped error dynamics. However, since $\hat{H}$ is not guaranteed to remain uniformly positive-definite in the course of adaptation, one must then redefine Y in the adaptation law (4) to satisfy

$$\tilde{H}(q)(\ddot{q}_r - \lambda s) + \tilde{C}(q, \dot{q})\dot{q}_r + \tilde{g}(q) = Y(q, \dot{q}, \dot{q}_r, \ddot{q}_r)\tilde{a}$$

instead of (2). Global tracking convergence can be shown as before, since the same Lyapunov function (6) now yields

$$\dot{V}(t) = -s^T H s \leq 0$$

where the true inertial matrix H is uniformly positive definite.

Remark 4: Unlike the model-reference adaptive controllers in the robotics literature (see, e.g., [8] for a review), this adaptive controller does not require the mandatory use of a reference model. In tasks with only commanded joint position specified, a reference model can be used to supply the q_d, $\dot{q}_d$, and $\ddot{q}_d$ required by the tracking controller. But in many robotic tracking tasks, smooth Cartesian motions are planned beforehand and inverse kinematics are used to find the desired joint position, velocity, and acceleration. In these situations, the above adaptive controller is more advantageous than model-reference approaches since it follows the desired joint trajectories corresponding to the planned Cartesian motion, rather than trajectories altered by a reference model. Furthermore, our adaptive controller does not require assumptions or simplifications such as local linearization, time invariance, or decoupled-dynamics; nor does it need measurements of joint accelerations or inversion of the estimated inertia matrix.

C. Discussion

In this section, we discuss implementation aspects, computational efficiency, and combining adaptation of certain parameters with robustness to others and to disturbances.

1) Implementation Aspects:

i) Since the load is usually fixed with respect to the last link, it can be regarded as part of that link. Generally, the parameters of the manipulator itself need to be measured or estimated only once after its installation, possibly using the off-line estimation methods of [10] and [3], since these parameters do not change from task to task. In practice, the only parameters to be estimated on-line by the adaptive controller are, therefore, the equivalent dynamic parameters of the load grasped by the robot hand. Implementation of the adaptive controller on a six-degree-of-freedom manipulator carrying a rigid-body load requires the adaptation of 10 equivalent dynamic parameters corresponding to the mass of the load, its center of mass (three parameters), and its moments of inertia (six parameters, i.e., I_{xx}, I_{yy}, I_{zz}, I_{xy}, I_{xz}, and I_{yz}). If desired, continuous models of Coulomb and viscous friction may also be included in (1), and the corresponding coefficients can be estimated similarly.

ii) We can stop updating an unknown parameter when it reaches *a priori* known bounds, and resume updating as soon as the corresponding derivative changes sign. This intuitively motivated procedure can easily be shown to preserve tracking convergence.

iii) The algorithm can be implemented directly in Cartesian space simply by letting

$$\dot{q}_r = J^{-1}[\dot{x}_d + \Lambda(x_d - x)]$$

for a nonredundant manipulator [19], with J being the manipulator Jacobian matrix, and x and x_d denoting the actual and desired Cartesian positions.

iv) In the practical implementation of the above adaptive controller, the matrices $\hat{H}$, $\hat{C}$, and $\hat{g}$ may be updated at a slower rate than the rate used for the terms $\dot{q}_r$, $\ddot{q}_r$, and s, so as to reduce computations, since typically the tracking error terms vary much faster than the dynamic coefficient matrices (see, e.g., [9]). Due to the presence of $\dot{q}_r$ in the second term of control law (3), however, the controller cannot be implemented directly using fast recursive Newton–Euler algorithms [12]. The same is true of the matrix Y in the adaptation law (4). A simple modification of the adaptive algorithm, which allows recursive computations, is proposed in [19].

v) The previous analysis becomes invalid in the presence of actuator saturation, which occurs when one of the torques specified by the algorithm reaches the physical limit of the corresponding actuator. Many practical approaches can be used in order to deal with torque saturation. The speed of the desired trajectories may be reduced, thereby reducing the required magnitude of the actuator torques, since saturation typically occurs when the load is too heavy for the given speed and given torque capacity. The controller may also be switched temporarily into a conservative fixed-parameter mode (such as independent joint PD + fixed-parameter feedforward) when one of the specified torques exceeds the known physical limit of the

corresponding actuator; indeed, the adaptation mechanism (4), in extracting parameter information from the tracking errors, cannot be expected to distinguish whether those errors are due to parameter errors or to saturation effects. The on-line parameter estimators of [15] and [13], which are based solely on prediction errors and are therefore unaffected by torque saturation, may then be switched on to keep estimating the load parameters for later use. The full operation of the adaptive controller may then be resumed when the torques (3), computed based on the current parameter estimates, return within admissible values.

2) Combining Adaptation with Robust Control:

In practice, the computational efficiency of the algorithm may be enhanced by not estimating all unknown parameters. Some parameters, e.g., the frictions coefficients in direct-drive robot joints, or the cross moments of inertia I_{xy}, I_{xz}, and I_{yz} for loads with basically regular shape, may have relatively minor importance in the dynamics, in which case one may choose to make the controller robust to the uncertainty on these parameters, rather than estimating them on-line. Similarly, the center of mass of the load may have been estimated with reasonable precision through visual information or CAD data. Furthermore, the controller must be robust to residual time-varying disturbances, such as stiction or torque ripple.

Assume, without loss of generality, that only the first α unknown parameters are actually estimated on-line, while the rest are known or estimated off-line *a priori*, or simply taken to be zero. We can write, following [22]

$$\boldsymbol{a} = \{\boldsymbol{a}_E \; \boldsymbol{a}_R\}^T$$

with the row vectors $\boldsymbol{a}_E$ and $\boldsymbol{a}_R$ defined by

$$\boldsymbol{a}_E = \{a_j\}_{j=1,\cdots,\alpha} \quad \boldsymbol{a}_R = \{a_j\}_{j=\alpha+1,\cdots,m}$$

and let, correspondingly,

$$\boldsymbol{Y} = [\boldsymbol{Y}_E \; \boldsymbol{Y}_R].$$

Assume that the uncertainties on $\boldsymbol{a}_R$ as well as $\boldsymbol{d}(t)$, the disturbance torques reflected to the manipulator joints, are bounded by known positive constants or functions of time

$$|\tilde{a}_j| \le A_j \qquad j = \alpha+1, \cdots, m$$

$$|d_j(t)| \le D_j(t) \qquad j = 1, \cdots, n.$$

Let us add a sliding control term to torque input (3)

$$\tau = \hat{\boldsymbol{H}}\ddot{\boldsymbol{q}}_r + \hat{\boldsymbol{C}}\dot{\boldsymbol{q}}_r + \hat{\boldsymbol{g}} - \boldsymbol{K}_D\boldsymbol{s} - \boldsymbol{K}\,\mathrm{sgn}\,(\boldsymbol{s})$$

where the notation

$$\boldsymbol{K}\,\mathrm{sgn}\,(\boldsymbol{s}) = \{k_i\,\mathrm{sgn}\,(s_i)\}^T_{i=1,\cdots,n}$$

is used for simplicity, and the k_i's are yet to be chosen. With $\boldsymbol{a}_E$ and Γ_E in place of $\boldsymbol{a}$ and Γ in the Lyapunov function (6), we obtain

$$\dot{V}(t) = \boldsymbol{s}^T[-\boldsymbol{K}_D\boldsymbol{s} + \boldsymbol{Y}_R\tilde{\boldsymbol{a}}_R + \boldsymbol{d} - \boldsymbol{K}\,\mathrm{sgn}\,(\boldsymbol{s})] + \tilde{\boldsymbol{a}}_E^T[\Gamma_E^{-1}\dot{\tilde{\boldsymbol{a}}} + \boldsymbol{Y}_E^T\boldsymbol{s}]$$

since

$$\boldsymbol{K}\,\mathrm{sgn}\,(\boldsymbol{s}) - \boldsymbol{Y}_R\tilde{\boldsymbol{a}}_R - \boldsymbol{d} = \left\{k_i\,\mathrm{sgn}\,(s_i) - \sum_{j=\alpha+1}^{m} Y_{ij}\tilde{a}_j - d_i(t)\right\}^T_{i=1,\cdots,n}$$

Let us then take

$$k_i = \sum_{j=\alpha+1}^{m} |Y_{ij}|A_j + D_i(t) + \eta_i \qquad i = 1, \cdots, n$$

$$\dot{\hat{\boldsymbol{a}}}_E = -\Gamma_E\boldsymbol{Y}_E^T\boldsymbol{s}$$

where the η_i's are positive constants. This yields

$$\dot{V}(t) \le -\boldsymbol{s}^T[\boldsymbol{K}_D\boldsymbol{s} + \{\eta_i\,\mathrm{sgn}\,(s_i)\}^T_{i=1,\cdots,n}]$$

$$= -\sum_{i=1}^{n}\eta_i|s_i| - \boldsymbol{s}^T\boldsymbol{K}_D\boldsymbol{s} \le 0$$

so that the system trajectories are still guaranteed to reach the surface $\boldsymbol{s} = \boldsymbol{0}$, and therefore convergence of the tracking errors to zero is achieved.

Furthermore, in order to avoid undesirable control chattering, we can use the saturation function sat (s_i/ϕ_i) in place of the switching function sgn (s_i), with the ϕ_i's representing the thicknesses of the corresponding "boundary layers." Similar to [Slotine and Coetsee, 1986], parameter adaptation must then be stopped when the system trajectories are inside the boundary layers; indeed, by definition, disturbances and errors on $\boldsymbol{a}_R$ can drive the trajectories anywhere in the boundary layers without this providing any information about estimation error on $\boldsymbol{a}_E$. Then the components of $\boldsymbol{s}$ are guaranteed to remain in the boundary layers, with corresponding small tracking errors. This procedure also has the advantage of avoiding long-term drift of the estimated parameters.

Note that, if none of the unknown parameters is explicitly estimated ($\alpha = 0$), then a fixed-parameter sliding controller is obtained, which owes its simplicity to the exploitation of energy conservation, and of the linearity of the robot dynamics in terms of the unknown parameters.

III. Experimental Background

Before presenting the experimental results in Section IV we briefly discuss here the experimental equipment, the dynamic model of the robot, and the design of the adaptive controller.

A. Equipment

The equipment used [5] is a two degree-of-freedom semi-direct-drive robot arm developed at the Whitaker College of Health Sciences at M.I.T., with "semi" indicating that the second link is indirectly driven by a motor located at the base through a four-bar mechanism. The arm was designed to be used as an experimental apparatus for investigating human arm movements. The system consists of a two-link arm (Fig. 2), two dc servo motors with amplifiers, two optical encoders, two tachometers, and a microcomputer PDP 11/73. The arm lies in the horizontal plane, and therefore, the effects of gravity are absent (actually, the presence of gravity would further demonstrate the advantages of adaptive parameter estimation). The two links are made of aluminium, with lengths of 0.37 m and 0.34 m, and masses of 0.9 kg and 0.6 kg, respectively. Although it uses a four-bar linkage mechanism, the arm is not dynamically mass-balanced, and therefore presents full coupling effects.

The motors are driven by PMI SSA 40-10-20 pulse-width-modulated switching servo amplifiers. The inductance of the motors is low enough so that the amplifiers are considered as current sources to the motors. The two JR16M4CH motors, mounted on a rigid supporting frame, are rather large and heavy (16 kg each), but this does not represent a practical problem since they are both located at the manipulator base. A four-bar linkage mechanism is used to transmit the torque from the upper motor to the outer link. The motion of the relative angle between the inner and outer links ranges from 39° to 139°. This range is made possible by using an offset in the elbow of the manipulator. Due to the fluctuation and nonlinearity limitations of the amplifier capacity, the maximum torque each motor can generate is 9 N·m. In the controller implementation, the motors together with the amplifiers are regarded as having a constant of 1.12 N·m/V as their transfer function.

The joint positions are measured by incremental optical

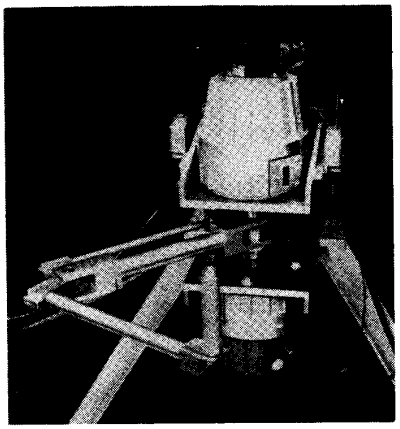

Fig. 2. The experimental two degree-of-freedom manipulator.

encoders attached to the output shaft of each torque motor, with a resolution of 12 bits/180°, i.e., 0.045°. The joint velocities are directly measured by tachometers. The tachometers, built in the motor housings, were originally designed for high-speed operations of motor shafts, and therefore have low output voltages for the comparatively low-speed rotations of the direct-drive arm shafts. Rather than modify the tachometers themselves, their output signals are amplified, which creates a considerable amount of noise. Another problem with the tachometers is their sensitivity to the vibrations of the supporting frame of the arm, which exhibits natural structural modes at about 400 Hz. Analysis in [5] suggests the use of a low-pass filter consisting of a cascade of four passive first-order filters, each with a nominal cutoff frequency of 50 Hz. While this filter is effective in eliminating the noises, it leads to significant phase lag (about 5 degrees at 1 Hz). This phase lag is believed to be one major source of residual tracking error for the adaptive control implementations.

Position and velocity measurements are sent to the PDP 11/73 for torque computation, with the control programs written in the C language. The resulting sampling frequency is 200 Hz.

B. Dynamic Model and Adaptive Controller Design

The dynamic model of the manipulator can be derived from Lagrange's equations to be

$$a_1\ddot{q}_1 + (a_3c_{21} + a_4s_{21})\ddot{q}_2 - a_3s_{21}\dot{q}_2^2 + a_4c_{21}\dot{q}_2^2 = \tau_1 \qquad (9\text{-a})$$

$$(a_3c_{21} + a_4s_{21})\ddot{q}_1 + a_2\ddot{q}_2 + a_3s_{21}\dot{q}_1^2 - a_4c_{21}\dot{q}_1^2 = \tau_2 \qquad (9\text{-b})$$

where $c_{21} = \cos(q_2 - q_1)$, $s_{21} = \sin(q_2 - q_1)$. It is clearly linear in terms of the four parameters a_1, a_2, a_3, a_4, which are related to the physical parameters of the links in Fig. 3 through

$$a_1 = J_1 + J_a + m_2l_1^2 + m_ah_a^2 \qquad a_2 = J_2 + J_b + m(2)h_2^2 + m_ar_a^2$$

$$a_3 = m_2h_2l_1\cos\delta - m_ah_ar_a \qquad a_4 = m_2h_2l_1\sin\delta \qquad (10)$$

with the load treated as part of the second link. Defining the components of the matrix $\boldsymbol{C}$ as

$$C(1,1) = C(2,2) = 0$$

$$C(1,2) = (a_4c_{21} - a_3s_{21})\dot{q}_2 \qquad C(2,1) = (a_3s_{21} - a_4c_{21})\dot{q}_1$$

the skew-symmetry of $\dot{\boldsymbol{H}} - 2\boldsymbol{C}$ can also be confirmed easily.

For simplicity, the feedback gain matrix $\boldsymbol{K}_D$ and the adaptation gain matrix Γ in the controller design are chosen to be diagonal

$$\boldsymbol{K}_D = \text{diag}(k_{d1}, k_{d2}) \qquad \Gamma = \text{diag}(\gamma_1, \gamma_2, \gamma_3, \gamma_4).$$

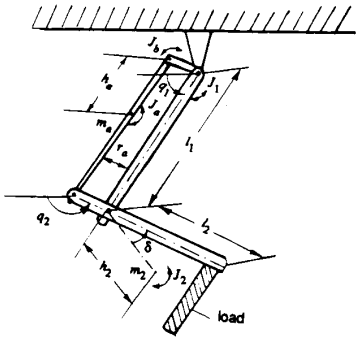

Fig. 3. The schematic structure of the manipulator.

The explicit form of the control law $\tau = \hat{\boldsymbol{H}}\ddot{\boldsymbol{q}}_r + \hat{\boldsymbol{C}}\dot{\boldsymbol{q}}_r - \boldsymbol{K}_D\boldsymbol{s}$ in terms of the parameter vector $\boldsymbol{a}$ is

$$\tau_1 = Y_{11}a_1 + Y_{13}a_3 + Y_{14}a_4 - k_{d1}s_1 \qquad (11\text{-a})$$

$$\tau_2 = Y_{22}a_2 + Y_{23}a_3 + Y_{24}a_4 - k_{d2}s_2 \qquad (11\text{-b})$$

where

$$Y_{11} = \ddot{q}_{r1} \qquad Y_{13} = c_{21}\ddot{q}_{r2} - s_{21}\dot{q}_2\dot{q}_{r2} \qquad Y_{14} = s_{21}\ddot{q}_{r2} + c_{21}\dot{q}_2\dot{q}_{r2}$$

$$Y_{22} = \ddot{q}_{r2} \qquad Y_{23} = c_{21}\ddot{q}_{r1} + s_{21}\dot{q}_{r1}\dot{q}_1 \qquad Y_{24} = s_{21}\ddot{q}_{r1} - c_{21}\dot{q}_{r1}\dot{q}_1.$$

The adaptation law can be explicitly written as

$$\dot{\hat{a}}_1 = -\gamma_1 Y_{11}s_1 \qquad (12\text{-a})$$

$$\dot{\hat{a}}_2 = -\gamma_2 Y_{22}s_2 \qquad (12\text{-b})$$

$$\dot{\hat{a}}_3 = -\gamma_3(Y_{13}s_1 + Y_{23}s_2) \qquad (12\text{-c})$$

$$\dot{\hat{a}}_4 = -\gamma_4(Y_{14}s_1 + Y_{24}s_2). \qquad (12\text{-d})$$

Note that the Coulomb and viscous frictions at the motor shafts and at the joints between links are neglected in the adaptive controller design and regarded as disturbances. The stiction effect at the joints is handled by an additional sliding control term, as explained later.

Since the adaptation law consists of a set of integrators, the high-frequency components of the measurement noises from the joint encoders and tachometers are mostly filtered out. Furthermore, small dead zones can be used in the computation of s, in order to account for the inaccuracies of the sensors. The adaptive controller is therefore expected to be reasonably robust to sensor noise.

C. Physical Interpretation of the Adaptation Mechanism

The above equations (12-a)–(12-d) have a strong intuitive appeal. The rate of adaptation is linearly proportional to the tracking error measure, s, with the γ_j's as adaptation gains. The estimate $\hat{a}_1$ is driven only by the tracking error of the first joint since a_1 appears only in the first joint motion equation (9-a); similarly for $\hat{a}_2$. On the other hand, $\hat{a}_3$ is driven by a combination of the first joint error and the second joint error, because a_3 appears in both (9-a) and (9-b); and similarly for $\hat{a}_4$.

Let us take a closer look at (12-a) and see how the adaptation mechanism extracts information about parameter uncertainty from the tracking error measure s. Consider, for instance, the case when Y_{11} is positive. For negative s_1, meaning a lag behind the desired motion in the first joint, adaptation law (12-a) increases the estimate $\hat{a}_1$. The reason for this increase is that the adaptation mechanism interprets the lag as indicating an insufficient joint torque τ_1, and since $(\partial\tau_1/\partial\hat{a}_1) = Y_{11} > 0$, this is in turn

interpreted as an insufficiency in $\hat{a}_1$. On the other hand, this increase in $\hat{a}_1$ leads to an increase in the torque τ_1, as seen from (11-a), and, accordingly, to a reduction in the lag behind the desired motion. The adaptation mechanism of $\hat{a}_2$ can be explained using the same argument. The variations of $\hat{a}_3$ and $\hat{a}_4$ can be explained similarly, with the difference that these parameter estimates are driven by weighted averages of s_1 and s_2. The weighting factors for $\hat{a}_3$, for instance, are Y_{13} and Y_{23}; their magnitudes reflect the extent of the contribution of a_3 to τ_1 and τ_2, and the corresponding lags in the joints.

Thus, the adaptation law generates the parameter estimates in an intuitively reasonable or "intelligent" manner.

IV. Experimental Results and Discussion

The purpose of the experiments is twofold, namely, to demonstrate the stability and the performance characteristics inferred from the theoretical development, and furthermore to compare the performance of the adaptive controller to those of the PD (local proportional-plus-derivative control) and computed torque controllers, two popular nonadaptive control methods. PD control, namely

$$\tau = -\boldsymbol{K}_D\dot{\tilde{q}} - \boldsymbol{K}_P\tilde{q}$$

is considered here instead of PID, since the lack of gravity, and therefore of corresponding steady-state errors, make an integral term superfluous. Three sets of experimental results are presented below, namely, comparison of PD and adaptive controllers, comparison of computed torque and adaptive controllers, and comparison of all controllers in the face of a large load. In the experiments, the design parameters of each controller are turned to their best values, in terms of the conflicting requirements of tracking accuracy in the joint motions and controller stability in the face of measurement noise, disturbances, and unmodeled high-frequency dynamics, so that the best performances of the three controllers can be compared.

The desired trajectories are the same for the three sets of experiments (Fig. 4). They last for 1 s with the first half-second for tracking motion between initial position $\boldsymbol{q}_s$ and end position $\boldsymbol{q}_e$, and the second half-second for regulation of the residual tracking errors to zero. The desired joint trajectories to be tracked are two fifth-order polynomials interpolated between $\boldsymbol{q}_s = [20.^\circ\ 100.^\circ]^T$ and $\boldsymbol{q}_e = [70.^\circ\ 100.^\circ]^T$, with zero desired velocities and accelerations at $t = 0$ and $t = 0.5$ s (note that, in terms of Remark 2 of Section II-B, this desired trajectory is rather "unexciting," and therefore represents a good challenge to an adaptive controller).

A. Comparison of PD and Adaptive Controllers

In this first case, no load is attached to the second link of the manipulator. For the adaptive controller, the initial values of the parameter estimation are taken to be zero, that is, the parameters of the arm are assumed to be totally unknown. The adaptive controller therefore starts as a PD controller and the feedforward part plays an increasingly effective role as the parameter adaptation is driven by the tracking errors. The adaptation gains γ_i are chosen to be all equal to 0.2.

For both controllers, proper choices of the feedback gains $\boldsymbol{K}_D$ and $\boldsymbol{K}_P$ (for the adaptive controller, $\boldsymbol{K}_P = \boldsymbol{K}_D\Lambda$) are important, since larger values lead to better tracking accuracy, but less robustness to measurement noise and unmodeled high-frequency dynamics. For simplicity, these matrices are all chosen to be diagonal in the experiments. By increasing $\boldsymbol{K}_D$ and $\boldsymbol{K}_P$, both controllers were found to become unstable at essentially the same values of these matrices. This suggests that the adaptive controller enjoys basically the same level of robustness to noises and high-frequency unmodeled dynamics as the PD controller. The diagonal values

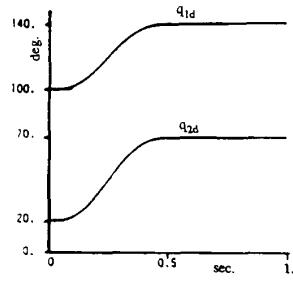

Fig. 4. The desired trajectories q_{1d} and q_{2d}.

$$k_{d1} = 2.0\ (\text{N} \cdot \text{m} \cdot \text{s/rad}) \quad k_{d2} = 2.0\ (\text{N} \cdot \text{m} \cdot \text{s/rad})$$

$$k_{p1} = 40.\ (\text{N} \cdot \text{m/rad}) \quad k_{p2} = 30.\ (\text{N} \cdot \text{m/rad})$$

or, equivalently for the adaptive controller, $\lambda_1 = 20$ and $\lambda_2 = 15$ were found to yield the best accuracy while avoiding noticeable excitation of the vibrational modes of the links.

A small dead zone is used in the computation of s, in order to handle the inaccuracies of the measurement signals and alleviate parameter drifts. The size of the dead zone is determined based on the resolutions of the encoders and the A/D converters for the tachometers.

The results of the PD controller are plotted in Fig. 5, while those of the adaptive controller are given in Fig. 6. The maximum joint errors for the PD controller are 6.57° and 4°, but those for the adaptive controller are only −2.12° and −2°. As expected, the errors and control torques of the two controllers are very close in the initial period, but the parameter estimates of the adaptive controller are quickly driven by the tracking errors. At roughly $t = 0.15$ s, the feedforward term in the adaptive control law, based on these estimates, is able to prevent the further growth of the tracking errors (which in the PD case, reach −5.° and −3.7°, respectively). The parameter estimates in Fig. 5(c) are seen to be very smooth, as expected from the integrator structure of the adaptation law. This smoothness is desirable because it avoids that the adaptation excite the vibrational modes of the links. At the end of the tracking operation in the first-half-second, usually an important instant in applications like pick-and-place tasks, the joint errors of the adaptive controller are only −0.7° in both angles, while the PD controller is suffering almost maximum joint errors. In the last-half-second, the joint errors are regulated by the controllers and the arm settles down to the end position $\boldsymbol{q}_e$.

Small steady-state joint errors of −0.09° (two encoder counts) and −0.135° (three encoder counts) are observed for the adaptive controller. These errors arise from the stiction effects of the static frictions at the motor shafts of this manipulator. The magnitudes of stiction were determined to be roughly 0.17 N·m at each joint. A small stiction-compensation term $[-0.15\ \text{sgn}\ (s_1),\ -0.10\ \text{sgn}\ (s_2)]^T$ was included in the control law of Fig. 6, otherwise the steady-state errors would have been about −0.3°. The sizes of these compensation terms are limited to small values by the condition that the corresponding accelerations should not make the system cross the dead zones in s in less than two sampling periods. The PD controller has the same level of steady-state errors, although they cannot be clearly seen in Fig. 4 due to the larger scale used.

B. Comparison of Computed Torque and Adaptive Controllers

In practice, the user always has some *a priori* knowledge of the robot parameters, possibly by computation based on design data. This information can be used to initialize the estimated parameters in the adaptive scheme, and to temporarily stop adaptation on a parameter if a known bound is reached. A popular alternative way

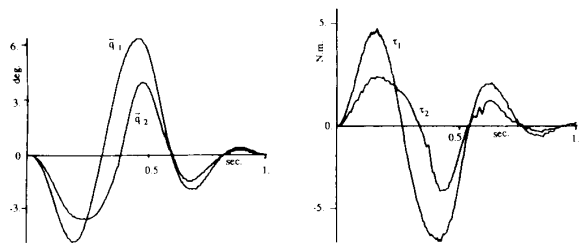

Fig. 5. PD control without load.

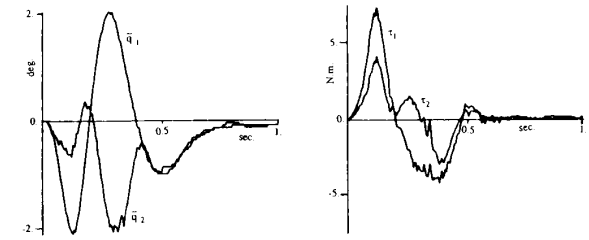

Fig. 6. Adaptive control starting from zero estimates.

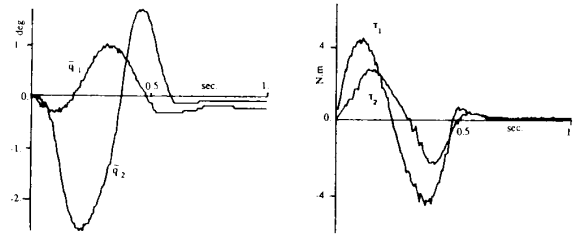

Fig. 7. Computed torque control.

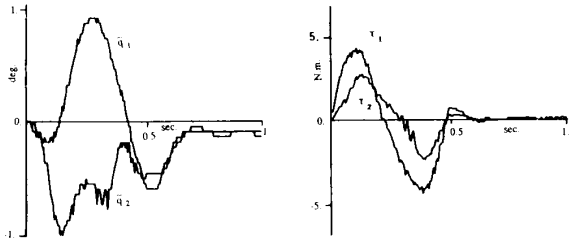

Fig. 8. Adaptive control starting with nonzero estimates.

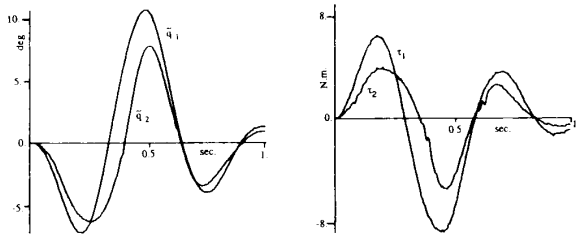

Fig. 9. PD control under load.

of using the *a priori* parameter estimates is to use fixed-parameter model-based controllers, such as the computed torque method. In this set of experiments, the performance of the computed torque and adaptive controllers are compared using *a priori* parameter estimates from [5] as the nominal parameters in the computed-torque and as the initial parameters in the adaptive controller.

The computed torque method is a fairly standard approach, whose formulation can be found in a number of papers (e.g., [12]). In the absence of the gravity, its input torque can be written as

$$\tau = \hat{H}(\ddot{q}_d - K_1\dot{\tilde{q}} - K_2\tilde{q}) + \hat{C}\dot{q}.$$

Let us take K_1 and K_2 as diagonal matrices

$$K_1 = \text{diag}\,(2\omega_1, 2\omega_2) \quad K_2 = \text{diag}\,(\omega_1^2, \omega_2^2)$$

where ω_1 and ω_2 are two positive constants. With this choice of K_1 and K_2, a critically damped error dynamics would be obtained if the exact parameters were used. Selecting K_1 and K_2 experimentally as before, the best values of ω_1 and ω_2 are determined to be $\omega_1 = 20$ rad·s^{-1} and $\omega_2 = 30$ rad·s^{-1}.

The design parameters of the adaptive controller are the same as before except that Γ is increased by a factor of two, since reasonable initial parameters are already available. The parameter values used for the computed torque method and as initial values of the adaptive control are

$$a_1 = 0.11 \text{ kg·m}^2 \quad a_2 = 0.0285 \text{ kg·m}^2 \quad a_3 = 0.033 \text{ kg·m}^2 \quad a_4 = 0.$$

They are computed from the engineering drawings of the arm links. In addition to the discrepancies between the real quantities and those on the drawings, the mass of the force sensor attached to the endpoint also causes some inaccuracy in the above values.

The results are shown in Fig. 7 for the computed torque method and Fig. 8 for the adaptive controller. The maximum joint tracking errors for the computed torque are 1° and −2.5°, respectively, while those for the adaptive controller are 0.95° and −0.96°. The tracking error of the first joint is smaller because the parameter uncertainty is larger in the four-bar mechanism associated with the second link.

C. Comparisons in the Presence of a Large Load

Since the underlying purpose of the adaptive controller is to maintain tracking accuracy in the face of significant uncertainty in the load parameters, a large load is attached to the end of the second link of the arm to demonstrate the performance of the adaptive controller. The load is a clutch which has roughly half the size and weight of the second link. The same trajectory is successively controlled by the PD, the computed torque, and the adaptive controllers. The PD controller is identical to the one in Section IV-A and the computed torque and adaptive controllers are the same as the ones in Section IV-B, both in design parameters and initial parameters. The results are plotted in Figs. 9–11.

For the PD controller, the maximum tracking errors increase to 10.2° and 7.4°. Even at the end of the last-half-second regulation, the tracking errors are still far from settled down (1.3° and 0.9°, respectively). The maximum tracking errors of the computed torque controller are 1.9° and −4.7°, representing increases of

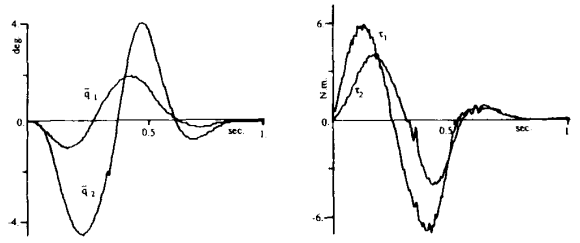

Fig. 10. Computed torque control under large load.

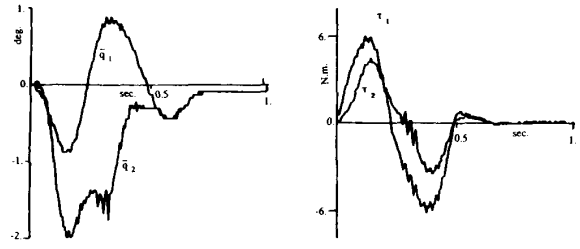

Fig. 11. Adaptive control under load.

0.9° and 2.2°. The much larger increase in the second joint is due to the fact that the attached clutch causes a larger increase in the uncertainty of the second joint dynamics than in that of the first joint. The maximum tracking errors for the adaptive controller are now 0.9° and −2.0°, with the second joint error increased by 1.04°, but no increase in the first joint error.

When the parameter estimates obtained at the end of this run are used for a second adaptive run, the maximum errors of both joints are found to stay within 1°. The PD and computed-torque controllers would, of course, essentially repeat their errors. In further adaptive experiments along varied and longer trajectories, long-term parameter drift was found to be very slight, due to the stopping of adaptation in the dead zone. Parameter drift does not seem to be a severe problem in high-speed robotic manipulation tasks because the tasks last for at most a few seconds.

D. Error Sources

The theory in Section II predicts that the tracking errors of the adaptive controller globally converge to zero, while actually the tracking errors in the second and later runs along the same trajectory merely stay within the one degree range, without exact convergence. The discrepancy is due to various noises, disturbances, and unmodeled dynamics, which are inherent in the experiments but ignored in the theoretical analysis. The tracking errors observed arise from many hardware or software sources. These include, in particular the following.

1) Arm Modeling Errors: The substantial Coulomb frictions at the motor shafts are not modeled. The Coulomb and viscous frictions at the linkage joints are not compensated at all in the experiments. The neglected frictions are believed to contribute a significant portion of the tracking errors. Furthermore, the vibrational dynamics of the links may also have contributed a certain amount of error.

2) Actuation Errors: For simplicity, the amplifiers and the motors have been modeled as constant gains, while they actually have dynamics of their own, which may not be negligible for this half-second fast operation. Furthermore, preliminary testing indicates an error of about 3 percent in the amplifier gain specified by the manufacturer and used in this experiment, in addition to a small torque ripple. In the absence of joint torque sensors, no attempt has been made to compensate for the torque inaccuracies.

3) Measurement Errors: Joint velocity measurements contain a considerable amount of error. The tachometer signals are small and are therefore sensitive to noise. The signals after amplification and filtering contain quite severe phase lag. In addition, the A/D converters for the velocity signals create a certain amount of error. For this high accuracy controller, the relatively low resolution of the optical encoders (0.045°, i.e., 12 bits/180°) does not allow numerical differentiation to advantageously replace the tachometer signals.

4) Real-Time Computing Limitations: Roundoff errors and, especially, sampling limitations, are further sources of error at these high speeds.

V. Concluding Remarks

These experimental results demonstrate that, despite all the nonparametric error sources, the adaptive controller enjoys essentially the same level of robustness to measurement noise and unmodeled dynamics as the PD controller, yet achieves much better tracking accuracy than either PD or computed-torque schemes. Its superior performance for high-speed operations, in the presence of parametric and nonparametric uncertainties, and its relative computational simplicity, make it an attractive option both to address complex industrial tasks, and to simplify high-level programming of more standard operations.

The presented comparisons, however, should not be interpreted as suggesting the replacement of PD or computed-torque controllers by adaptive control in all robotic applications. An advanced robotic control system should switch between different control modes depending on the speed, precision, and parameter adaptation requirements of the tasks.

Acknowledgment

The authors are grateful to Prof. N. Hogan for letting them borrow the fine manipulator and experimental setup described above, and to I. Faye for his help in the implementation.

References

[1] S. Arimoto and F. Miyazaki, "On the stability of P.I.D. feedback with sensory information, " in *Proc. Int. Symp. Robotics Res., Bretton Woods,* Cambridge, MA: M.I.T. Press, 1984.

[2] K. J. Astrom, "Interaction between excitation and unmodelled dynamics in adaptive control," presented at the IEEE Conf. Decision Contr., Las Vegas, NV, 1984.

[3] C. G. Atkeson, C. H. An, and J. M. Hollerbach, "Estimation of inertial parameters of rigid body links of manipulators," presented at the IEEE Conf. Decision Contr., Fort Lauderdale, FL, 1985.

[4] J. J. Craig, P. Hsu, and S. Sastry, "Adaptive control of mechanical manipulators," presented at the IEEE Int. Conf. Robotics Automat., San Francisco, CA, 1986.

[5] I. C. Faye, "An impedance controlled manipulandum for human movement studies," M.S. thesis, Dep. Mechanical Eng., Mass. Inst. Technol., Cambridge, MA, 1986.

[6] L. Guzzella and H. P. Geering, "Model following variable structure control for a class of uncertain mechanical systems," presented at the IEEE Conf. Decision Contr., Athens, Greece, 1986.

[7] I. J. Ha and E. G. Gilbert, "Robust tracking in nonlinear systems," *IEEE Trans. Automat. Contr.,* vol. AC-32, pp. 763-771, Sept. 1987.

[8] T. C. Hsia, "Adaptive control of robot manipulators—A review," presented at the IEEE Int. Conf. Robotics Automat., San Francisco, CA, 1986.

[9] O. Khatib, presented at the US-Japan Symp. Flexible Automation, Osaka, Japan, 1986.

[10] P. Khosla and T. Kanade, "Parameter identification of robot dynamics," presented at the IEEE Conf. Decision Contr., Fort Lauderdale, FL, 1985.

[11] D. Koditschek, "Natural motion of robot arms," presented at the IEEE Conf. Decision Contr., Las Vegas, NV, 1984.

[12] J. Luh, M. Walker, and R. Paul, "On-line computational schemes for mechanical manipulators," *ASME J. Dynam. Syst. Meas. Contr.,* vol. 102, 1980.

[13] W. Li and J. J. E. Slotine, "Parameter estimation strategies for robotic applications," presented at the ASME Winter Annual Meet., Boston, MA, 1987.

[14] ——, "Indirect adaptive robot control," presented at the IEEE Int. Conf. Robotics Automat., Philadelphia, PA, 1988.

[15] R. H. Middleton and G. C. Goodwin, "Adaptive computed torque control for rigid link manipulators," presented at the IEEE Conf. Decision Contr., Athens, Greece, 1986.

[16] A. P. Morgan and K. S. Narendra, "On the uniform asymptotic stability of certain linear nonautonomous differential equations," *SIAM J. Contr. Optimiz.,* vol. 15, 1977.
[17] C. E. Rohrs, L. S. Valavani, M. Athans, and G. Stein, "Robustness of continuous-time adaptive control algorithms in the presence of unmodelled dynamics," *IEEE Trans. Automat. Contr.,* vol. AC-30, Sept. 1982.
[18] J. J. E. Slotine, "The robust control of robot manipulators," *Int. J. Robotics Res.,* vol. 4, no. 2, 1985.
[19] J. J. E. Slotine and W. Li, "On the adaptive control of robot manipulators," presented at the ASME Winter Annual Meet., Anaheim, CA, 1986.
[20] ——, "Adaptive strategies in constrained manipulation," presented at the IEEE Int. Conf. Robotics Automat., Raleigh, NC, 1987.
[21] ——, "Theoretical issues in adaptive manipulator control," in *Proc. 5th Yale Workshop on Apl. Adaptive Systems Theory,* 1987.
[22] ——, "On the adaptive control of robot manipulators," *Int. J. Robotics Res.,* vol. 6, no. 3, 1987.
[23] M. W. Spong and M. Vidyasagar, "Robust linear compensator design for nonlinear robotic control," *IEEE J. Robotics Automat.,* vol. RA-3, no. 4, 1987.

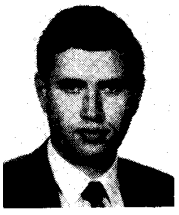

Jean-Jacques E. Slotine (S'82-M'83) was born in Paris, France, on October 18, 1959. He received the Ingenieur degree from the French Ecole Nationale Superieure de l' Aeronautique et de l' Espace in 1981, and the Ph.D. degree from the Massachusetts Institute of Technology, Cambridge, in 1983.

He is currently Assistant Professor of Mechanical Engineering, Doherty Professor in Ocean Utilization, and Director of the Nonlinear Systems Laboratory at M.I.T. His research interests include applied nonlinear control, robotics, and learning systems. He is the coauthor of the textbook *Robot Analysis and Control* (New York: Wiley, 1986), and is a frequent consultant to industry and to Woods Hole Oceanographic Institution.

Dr. Slotine is a Technical Editor of the IEEE JOURNAL OF ROBOTICS AND AUTOMATION, an Associate Editor of *The Robotics Review,* and an Associate Editor of the *International Journal of Robotics and Automation.*

Weiping Li was born in Hunan, China, on July 11, 1962. He received the B.E. degree in aircraft design from the Beijing Institute of Aeronautics and Astronautics, Beijing, China, in 1982, and the M.S. degree in aeronautics from the Massachusetts Institute of Technology, Cambridge, in 1986.

From 1982 to 1984, he worked in the Vibration Research Center of the Nanjing Aeronautical Institute, China, where he performed vibration and flutter control experiments on various helicopters and jet aircraft. He is currently completing work towards the Ph.D. degree at the Massachusetts Institute of Technology. His research interests include robotics, adaptive control theory, neural networks, and computer vision.