

May 1988

Report No. STAN-CS-88-1205

Thesis

2

DTIC FILE COPY

Dynamics for Robot Control: Friction Modeling and Ensuring Excitation During Parameter Identification


by

Brian Stewart Randall Armstrong



Department of Computer Science

Stanford University
Stanford, California 94305



AD-A198 732

ADA198732

REPORT DOCUMENTATION PAGE				Form Approved GSA No. 0704-0188 Exp. Date: Jun 30, 1986	
1a. REPORT SECURITY CLASSIFICATION unclassified			1b. RESTRICTIVE MARKINGS		
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release: Distribution Unlimited		
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE					
4. PERFORMING ORGANIZATION REPORT NUMBER(S) STAN-CS-88-1205			5. MONITORING ORGANIZATION REPORT NUMBER(S)		
6a. NAME OF PERFORMING ORGANIZATION Computer Science Department		6b. OFFICE SYMBOL (if applicable)	7a. NAME OF MONITORING ORGANIZATION		
6c. ADDRESS (City, State, and ZIP Code) Stanford University Stanford, CA 94305			7b. ADDRESS (City, State, and ZIP Code)		
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DARPA		8b. OFFICE SYMBOL (if applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER		
8c. ADDRESS (City, State, and ZIP Code) 1400 Wilson Blvd. Arlington, VA 22209			10. SOURCE OF FUNDING NUMBERS		
			PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.
11. TITLE (Include Security Classification)					
12. PERSONAL AUTHOR(S) Brian S. R. Armstrong					
13a. TYPE OF REPORT dissertation		13b. TIME COVERED FROM _____ TO _____		14. DATE OF REPORT (Year, Month, Day) May 1988	
15. PAGE COUNT 198					
16. SUPPLEMENTARY NOTATION					
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)		
FIELD	GROUP	SUB-GROUP			
19. ABSTRACT (Continue on reverse if necessary and identify by block number)					
To accurately control any mechanism it is necessary to know the relationship between applied forces and the resultant motion. These forces may be simple to compute, as is the case for many single degree of freedom machines; or they may be quite complex. Two steps toward the accurate prediction of motion forces are presented in this thesis: an experimental investigation of friction, and a study of the sensitivity of robot inertial parameter identification methods to noise.					
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS				21. ABSTRACT SECURITY CLASSIFICATION	
22a. NAME OF RESPONSIBLE INDIVIDUAL			22b. TELEPHONE (Include Area Code)		22c. OFFICE SYMBOL

DYNAMICS FOR ROBOT CONTROL: FRICTION MODELING
AND ENSURING EXCITATION DURING PARAMETER IDENTIFICATION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF
ELECTRICAL ENGINEERING
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY



By
Brian Stewart Randall Armstrong

May 1988

Accession For	
NTIS CRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

© Copyright by Brian Armstrong 1988

All Rights Reserved

Chas. F. Franklin

Thomas O. Binford

David Wein

iii

ABSTRACT

→ To accurately control any mechanism it is necessary to know the relationship between applied forces and the resultant motion. These forces may be simple to compute, as is the case for many single degree of freedom machines; or they may be quite complex. Two steps toward the accurate prediction of motion forces are presented in this thesis: an experimental investigation of friction, and a study of the sensitivity of robot inertial parameter identification methods to noise.

The friction study begins with an experimental investigation of the most basic properties required for predictive modeling: repeatability and structure. Friction is found to be surprisingly repeatable; position dependence is found, and a destabilizing effect - the Stribeck effect - is observed at low velocity. The experimental work is specific to a particular mechanism: the PUMA 560 arm; but many of the observations, particularly the study of the Stribeck effect, will extend to a broad class of machines. Using the friction model developed and an inertial model reported elsewhere, *open-loop control* of the PUMA robot is carried out, demonstrating the accuracy of the friction model.

When designing an identification experiment for a system described by nonlinear functions, such as those of manipulator dynamics, it is necessary to consider whether the excitation is sufficient to provide an accurate estimate of the parameters in the presence of experimental noise. It is shown that the convergence rate and noise immunity of a parameter identification experiment depend directly upon the condition number of the input correlation matrix, a measure of excitation. The

sensitivity of an identification experiment to unmodeled dynamics is also studied; a dimensionless measure of this sensitivity - Bias Susceptibility - is proposed and related to excitation. The issue how exciting a trajectory may be is addressed, and a method is presented to maximize the excitation. Two identification experiments reported in the literature are studied; analysis of these experiments shows that intuitively selected trajectories may provide poor excitation and that considerable improvement results from employing the optimization to maximize excitation.

TABLE OF CONTENTS

SECTION I: Friction Modeling

Chapter	9.1	Introduction	1
Chapter	2	Experiment Design	11
	2.1	Acceleration Measurement	
	2.2	Velocity Estimation	
	2.3	Contact Force Sensing	
	2.4	Torque Control	
	2.5	Break Away Torque	
	2.6	Linear Parameter Estimation	
	2.7	Non-Linear Parameter Estimation	
	2.8	Low Pass Filtering	
Chapter	3	Repeatability	16
Chapter	4	Break Away Experiments	24
	4.1	Experimental Issues in Measuring Break Away Torque	
	4.2	Building the Correction Table	
Chapter	5	Friction as a Function of Velocity	34
	5.1	Analysis of Variance on the Motion Friction Data	
	5.2	Motion Friction at Low Velocities	
	5.3	Motion Friction During Compliant Motions	
	5.4	The Dahl Effect	
	5.5	The Stribeck Effect	
Chapter	6	Demonstrations of Friction Compensation	61
	6.1	Open Loop Motion of One Joint	
	6.2	Open Loop Motion of Three Joints	
	6.3	Friction Compensated Force Control	

Chapter	7	Recommendations to the Engineer _____	80
	7.1	Recommendations on Experimental Technique	
	7.2	Recommendations on Control	
	7.3	Recommendations on Design	
Chapter	8	Conclusions _____	86
Appendix	A	Small Studies _____	88
	A.1	Friction as a Function of Motor Angle	
	A.2	Joint 2 Motor Alone and Joint 2 Link Alone	
	A.3	Trials with Dither	
	A.4	Friction as a Function of Load	
	A.5	Creep	
	A.6	Effects that were not Observed	
Annotated Bibliography		_____	103
		Additional References	

SECTION II: Excitation

Chapter	9	Introduction	114
	9.1	Introduction	
	9.2	Prior Work	
Chapter	10	The System and Noise Models	124
	10.1	A Model of Manipulator Rigid Body Dynamics	
	10.2	The Noise Model	
	10.3	A Note on Notation	
Chapter	11	Noise Analysis of the Sequential Algorithm	130
	11.1	The Basic LMS Parameter Identification Algorithm	
	11.2	The Consideration of Noise and Evaluation of Bias	
	11.3	Bias Due to Random Sensor Noise	
	11.4	Bias Due to Unmodeled Dynamics and Bias Susceptibility	
	11.5	The Consideration of Noise and Evaluation of Variance and Convergence Rate	
	11.6	Making the Choice of Γ	
	11.7	Noise Analysis of the Batch Least Squares Identification Algorithm	
	11.8	Summary of Noise Results	
Chapter	12	Generation of Optimal Identification Trajectories	151
	12.1	Quadratic Cost Applied to Acceleration and Jerk	
	12.2	Implementation Issues	
	12.3	Local Minima, the Global Minima and Path Constraints	
Chapter	13	Results	160
	13.1	Comparison of Original and Very Exciting Trajectories	
	13.2	Systematic Error: A Case Study	

Chapter 14	Discussion	177
14.1	Why Seemingly Active Trajectories May Provide Poor Excitation	
14.2	Trajectory Tracking versus Accurate Identification	
14.3	Adaptive Control and Reduce Model Order	
14.4	Conclusions	
Bibliography		181

ACKNOWLEDGEMENTS

First and foremost I would like to thank and acknowledge my Parents, Norman and Phyllis Armstrong; in the sense that a Saturn V had a lot to do with getting to the moon, they have had a lot to do with my arrival here. I would like to thank my sisters, Jennifer and Sara, for their loving support, and my brother-in-laws, Andy and Nick, for being dudes. I had a lot of momentum when I left home and public school; some of the teachers who made a difference were Ron Crampton, Al Gloor, Charles Lang, Phil Pfenniger and John Truell.

It was my great fortune to connect with my advisors, Professors Thomas Binford and Gene Franklin. The connection was made part of chance, part of mutual interest and part of affinity, and has been the greatest contributor to my success here at Stanford. Professor Binford long ago convinced me that the control of manipulators is a tough and interesting problem. While at Stanford, my greatest intellectual growth has been in mathematics; professor Franklin, by silent, possibly maddening insistence drove me to it; and I am especially grateful: it will stand me in good stead.

It was not purely by luck that I spent my time at Stanford working with such cool guys and robo jocks as Joel Burdick, Ron Fearing, Oussama Khatib, Shashank Shekhar, Bradley Chen, David Kriegman and Chung-sheng Cai - I thought about it. First I met Joel's feet; they were sticking out from under a pdp-11 computer, as if it had just been deposited by a tornado. Later I met Oussama, probably over a brew, and Ron, Shashank, David and Bradley. Joel and Oussama and I have pulled more

than a few all-nighters. Oussama, a research scientist, taught me what robots are about and, with my advisors, how to do research.

A university can be a self serving thing; Stanford University's fascination with its own reputation often seems complete. Students are tolerated at Stanford, as long as they don't get in the way of eminence.

None-the-less, there are rays of light that shine out from within the University: one that stands out in my experience is the libraries. From the lofty garret of the university archives to the deep basement of the medical library, I have used the libraries. And all but universally the aid of the librarians and staff has been freely given and well informed; and I appreciate it. I would especially like to mention Lois Sher, of the Engineering School Library, who has seen as much of me as my house-mate and who, on one occasion, allowed me to borrow two eighty year old German engineering journals, that were otherwise non-circulating, so that I could take them to my Austrian, illegal-alien, computer science, post-doctoral, carpenter friend for help with translation. It proved to be a key.

And thanks, of course, to my carpenter friend, Anu Ragi, for his help with translation from German. And to my good friend, the man of accidental adventures, Phil Kesten, for help with translation to Latin.

For three of my years at Stanford I was supported by the Hewlett Packard Corporation as part of their faculty development fellowship program. I appreciate their support and hope their program be successful. For two summers and an academic year I was supported by Professor Binford and the Stanford Institute for Manufacturing Automation, and for one quarter by the Dean of the School of

Engineering. This support was vital to carrying this work beyond the bare minimum and has carried me a great distance. I would also like to thank Professor Raju of the University of Ohio and Shawn Buckley of Cochlea Corporation for opportunities to work part time.

The Systron-Donner Corporation lent this project a rotational accelerometer of extraordinary performance, this piece of equipment facilitated the study of friction with sensitivity far beyond what would otherwise have been possible.

While at Stanford I have spent considerable effort resisting the offer or avoiding the use of military money; but I have failed. In no small part, the experimental and analytic portions of this research were carried out on apparatus and computing machinery provided by the Department of Defense.

Chapter 1

Introduction

Friction is universally present in the motion of bodies in contact. The modern science of Tribology seeks to explain the atomic details of friction. But the universality of friction may also be understood from a different perspective. Leonardo da Vinci, among his many investigations, studied the relationship between friction and the music of the heavens. He knew the music to be produced by the bumping and rubbing of the heavenly spheres, and he was concerned with the possibility of friction between these heavenly bodies:

“Had however this friction really existed, in the many centuries that these heavens have revolved they would have been consumed by their own immense speed of every day ... we arrive therefore at the conclusion that the friction would have rubbed away the boundaries of each heaven, and in proportion as its movement is swifter towards the centre than toward the poles it would be more consumed in the centre than at the poles, and then there would not be friction any more, and the sound would cease, and the dancers would stop ...”

Leonardo da Vinci (1452-1519),
The Notebooks, F 56 V

Leonardo understood that friction is absent from the state of grace. Thus friction is confined to this mortal world: friction is a consequence of original sin.

Friction plays a role in the simplest actions of living, such as walking, grasping and stacking. In many cases of importance the forces of friction are not small. But for all of this, in the discussion of dynamics for control of mechanical systems friction is but little studied and often completely omitted. When friction is addressed, the models are often those of Leonardo da Vinci or Charles de Coulomb. Experimental evidence pertinent to the situation under scrutiny is rarely sought out or presented. This arises, perhaps, out of the worth ascribed to problem independent solutions, the impetus toward results that may be applied *sine mutationibus mutandis* to every situation.

This is a study of friction in a particular machine: a dc electric robot with spur gears and ball bearings. We shall find along the way a number of effects that extend beyond the confines of this particular mechanism. But more importantly, an example is presented of what may be achieved by mechanism *specific* study. The potential for capturing friction forces in a predictive model is explored; and such fundamental, but long neglected issues as repeatability and structure are addressed directly. The issues addressed in this work - the fraction of friction forces that can be predicted, the model components that are dominant, the characterization of friction at low velocities, a demonstration of experimental procedure for determining basic friction properties - will extend *cum mutationibus mutandis* to a class of mechanisms far broader than that directly studied. After all, Leonardo's conclusions, derived from the study of the motion of bricks on a board, continue after five hundred years to dominate the design of control when friction is considered. A little experimental work can go a long way.

The control of contact forces between environment and machine, that is force control, presents a special challenge. This is due, in part, to the unforgiving coupling

between actuation and applied force. When the environment and mechanism are very stiff, there is no leeway between the control command and the contact force: a change in control effort is reflected immediately by a change in the contact force. Friction plays a dominant role in limiting the quality of force control. We are accustomed to specifying the upper performance limits of continuously controlled mechanisms: the maximum speed, the greatest motion. The simple theory, even when extended by Leonardo's model of friction, fails to predict that there are lower limits. But as a consequence of non-linear low-velocity friction, there are lower limits to motion, a *minimum* speed, a *minimum* distance. These lower limits to motion translate to limits on the fidelity of force control, substantial limits in the case of many practical machines.

A sense of the challenge of force control may be had by considering, for a particular manipulator, the motion corresponding to a small change of force. In a typical configuration the PUMA arm exhibits an end effector stiffness of 20,000 Newton-meters (N-m) per radian. If the desired force resolution is one tenth of a Newton (about one third ounce), at a typical radius of half a meter, it must be possible to govern motions of the mechanism as small as $1/400,000$ of a radian. Comparing this to the workspace of joint 1, which is roughly 5 radians, gives a ratio of desired resolution to total motion of 1 to 2,000,000. If this same requirement were applied to other control systems, it would lead to an elevator for the world's tallest building that could position its car with an error less than the thickness of this page, or a disk head the position of which may be commanded to within a tenth of a wave length of light. It is interesting to note the field outside of robotics where active compensation of friction has been undertaken: telescope pointing. *All* implementations of active friction compensation reported prior to interest in robotic

force control involve tracking by telescope [Walrath 84; Gilbert and Winston 74]. The attainable pointing resolution is in the neighborhood of $1/4$ arc second, which gives a ratio of resolution to total motion of $1 : 5,000,000$; not far from our goal with the PUMA arm.

More compliant mechanisms will reduce the bandwidth and resolution requirements of force control. Work has progressed toward governing the oscillatory behavior of compliant mechanisms [Cannon and Schmitz 84; Maples 85; Wang and Vidyasagar 87] and, in work specifically focused on the need for compliance in force control, towards achieving good positioning accuracy with compliant mechanisms [Roberts 84]. One impact of increased compliance is improvement in the ability of feedback control to compensate for friction induced errors. Attention is also being given to lower friction actuators and transmissions for robots [Asada and Youcef-Toumi 84; Pfeffer, Khatib and Hake 86]. Brushless and direct drive designs have shown tremendous reduction in friction [Asada and Youcef-Toumi 84] But, given the fact of original sin, friction will continue to be with us. Friction modeling will be important for compliant mechanisms, where active damping depends upon accurate system modeling. And in the investigation of low friction designs, detailed understanding of friction will be key to understanding the trade-offs among friction, performance, weight, payload and cost. And, especially because of the destabilizing character of bearing friction near zero velocity, friction will continue to be a performance limiting factor in force control.

The study of non-linear friction in servo-mechanisms has a long history. Tustin examined the effect of backlash and non-linear friction in a feedback loop using a vector graphic technique that anticipates the modern describing function analysis [Tustin 47]. Tou and Schultheiss thoroughly study the impact on control of

static and kinetic friction using describing function analysis [Tou 1953; Tou and Schultheiss 1953]. More recently Townsend and Salisbury [87] have employed describing function techniques to investigate a system similar to that studied by Tou and Schultheiss. The analysis of Townsend and Salisbury extends that of Tou and Schultheiss by paying careful attention to the regions of stability; they identify a region of input-dependent stability that is likely to occur in experimentally tuned controllers.

One motivation shared by the above authors has been to explain observed limit cycle behavior that a purely linear system cannot exhibit. Tou and Schultheiss [53] show the possibility of limit cycling during slow motions when the controller includes integral terms. They propose the use of a velocity feedback term to eliminate the oscillation and demonstrate the method using an analog simulation. Townsend and Salisbury [87] also show the possibility of limit cycling when integral control is used. Tustin [47], and in a qualitative way Gogoussis and Donath [87], chose a different structure for his model of friction at low velocity and was able to show the possibility of limit cycling when only proportional and derivative feedback are used.

Before proceeding, several definitions are provided - words describing friction are sometimes used imprecisely:

Tribology:

Literally, the study of rubbing; the name given to the modern study of friction and wear of rubbing surfaces.

Static Friction (Sticktion):

The torque (force) necessary to initiate motion from rest. It is often greater than the kinetic friction.

Kinetic Friction (Coulomb friction, Dynamic friction):

A friction component that is independent of the magnitude of the velocity.

Viscous Friction:

A friction component that is proportional to velocity and, in particular, goes to zero at zero velocity.

Negative Viscous Friction:

A phenomenon sometimes reported or assumed in which the friction at low velocities decreases with increasing velocity, a source of instability.

Break-Away:

The transition from rest (static friction) to motion (kinetic friction).

Break-Away Torque (Force):

The amount of torque (force) required to overcome static friction.

Break-Away Distance:

The distance traveled during break-away; that is, the distance over which static friction operates, a consequence of the materials used and forces applied.

Dahl Friction or the Dahl Effect:

A friction phenomenon that arises from the elastic deformation of bonding sites between two surfaces which are locked in static friction. The Dahl effect causes a sliding junction to behave as a linear spring for small displacements.

Stribeck Friction or the Stribeck Effect:

A friction phenomenon that arises from the use of fluid lubrication and gives rise to decreasing friction with increasing velocity at low velocity. See figure 1.3.

Deviation:

Many measurements are reported in this thesis that exhibit trial to trial variation; deviation is a measure of the trial to trial variation: it is the square root of the sum of the squared variations. Deviation is used as a measure of variation rather than variance because the magnitude of the deviation can be compared directly to the magnitude of the signal itself.

There is no agreement on the structure appropriate for a model of servo-mechanism friction. Mechanists and tribologists who have studied physical mechanisms of friction find numerous phenomena, varying with the materials and surface preparation used. Among this data there is evidence for decreasing friction with increasing velocity [Bowden and Tabor 73], increasing friction with increasing velocity [Bowden and Tabor 73], very short break-away distances [Rabinowicz 51], much longer break-away distances [Dahl 77], the presence of a static friction component [Bowden and Tabor 73; Rabinowicz 51], and the absence of a static friction component [Dahl 77]. Among control experimentalists who attempt to model friction, the choice is often a kinetic plus viscous friction model, as shown in figure 1.1 below [Gilbert and Winston 74; Walrath 84; Canudas, Astrom and Braun 86; Craig 86]. Control theorists, with a few exceptions, employ a static plus kinetic plus viscous model, as shown in figure 1.2 [Tou and Schultheiss 53; Satyendra 56; Bohacek and Tuteur 61; Shen 62; Shen and Wang 64; Townsend and Salisbury 87; an exception is Tustin 47]. Tustin is important because his model incorporates negative viscous friction. He alone is able to explain limit cycling in systems with no integral control, a phenomenon observed in practical systems. Experimentalists resist employing the static friction component because of its extreme dependence upon velocity, whereas the theorists require it to explain limit cycling.

Feed-forward compensation of friction has been undertaken using (non-linear) analog control [Gilbart and Winston 74]; and digital control [Walrath 84; Craig 86; Kubo, Anwar and Tomizuka 86]. Gilbart and Winston achieved a 6 to 1 reduction in the RMS pointing error of an optical tracking telescope by feed-forward friction compensation. They adaptively tuned a kinetic friction parameter. Walrath reports a feed-forward compensation scheme to point optical tracking instruments




Figure 1.1 Shape of a kinetic and Viscous Friction Model.




Figure 1.2 Shape of a Sticktion, kinetic and Viscous Friction Model.




Figure 1.3. Shape of a Sticktion, Kinetic and Viscous Friction Model with a Negative Velocity Dependence at low Velocities.

on a moving platform. He does extensive experimentation to develop and justify his friction model, which is a modified Dahl model [Dahl 77]. Walrath's system adaptively identifies one parameter and shows 10 to 1 reduction in the RMS pointing error by the use of friction compensation. Craig [86], and Kubo, Anwar and Tomizuka [86] have worked with friction compensation in robotics. Craig incorporates friction parameters in the model for his work in adaptive control, but does little to show that the friction parameters identify true physical processes or improve control. Kubo, Anwar and Tomizuka are specifically concerned with friction; they justify the use of kinetic friction compensation theoretically and demonstrate it practically on a DC motor driven robot.

This thesis addresses a lack of basic experimental data concerning friction in servo-mechanisms. Engineers building feed-forward controllers with friction

compensation employ experimental data to adjust model parameters, but no careful investigation of the properties of friction typical of servo-mechanisms has been undertaken. The theoretical work cited above - Tustin; Tou and Schultheiss; Satyendra; Bohacek and Tuteur; Shen; Shen and Wang; Townsend and Salisbury 87 - contains no reference whatever to experimental work involving servo-mechanisms. In chapter 2 the experimental procedures used are described. In chapter 3 the repeatability of friction phenomena is examined: any attempt at modeling and compensation is predicated upon repeatability. The motion data show a correlation between measured friction force and position. In chapter 4 this correlation is examined and an experiment that has been highly successful in identifying spatial dependencies is discussed. Chapter 5 addresses the relationship between friction and velocity; an experiment that provided sensitive measurements of friction at exceptionally low velocities is described. In chapter 6 friction-compensated free motion is demonstrated: the friction model is adequate to precompute torques and to move the arm under feed-forward, open-loop control. Cumulative position error during open-loop motions was less than 10%. In the last section of chapter 6 a demonstration of force control is presented. And in chapter 7 the engineering implications of this work are discussed and recommendations for the design of control for friction-afflicted mechanisms are presented.

Chapter 2

Experiment Design

The experiments reported here have been carried out with a PUMA 560 arm controlled with the NYMPH computer [Chen *et.al.* 86]. Two procedures for measuring friction have been used:

1. Measuring acceleration and subtracting computed inertial torques from known motor torques;
2. Measuring the minimum torque necessary to initiate motion.

2.1 Acceleration Measurement

Measuring acceleration is a substantial challenge. Using a two-sided Kalman smoother, acceleration can be estimated from recorded position data; however if noise rejection is turned up to an adequate level, bandwidth is quite poor. In trials with PUMA position data, acceptable noise rejection was determined to be 18 dB reduction of white quantization noise, which gave an RMS acceleration deviation of 0.3 rad/sec^2 . This level of noise rejection results in an effective bandwidth of roughly 2 hertz, a severe experimental limitation. The Systron-Donner corporation generously lent a model 8160 rotational accelerometer to this project. This extraordinarily sensitive instrument is specified to have a threshold and resolution of 0.005% of full scale; in the case of this instrument, that is 1 milli-radian per

second squared. The response of the instrument is flat to a second order roll-off at 40 hertz. An accurate acceleration measurement has been key to the measurement of friction during motion.

2.2 Velocity Estimation

In most of the experiments presented here, velocity was estimated off line using a two pass Kalman smoother running on position and torque data. This procedure gives a reliable, unbiased estimate of velocity, and has a bandwidth of roughly 10 hertz. For the low speed trials reported in chapter 5, a real time velocity estimate was required: a digital cross-over filter was implemented that had the high frequency response of the integrated acceleration signal and the low frequency response of the differentiated position measurement. The cross over frequency was 10 radians per second.

2.3 Contact Force Sensing

Measurements of contact force were made with the Stanford force sensing fingers (see [Khatib and Burdick 86]). The force finger sensor consists of a three axis load cell with a stiffness of 60,000 Newtons per meter. The sensitivity is roughly 0.01 Newtons, or 0.04 ounces.

2.4 Torque Control

The Unimate controller employs high gain current amplifiers to drive the motors. During operational tests the current was observed to slew to a new command value in less than 500 μ sec and hold that value to within 0.5%. The motor model used here, the standard dc motor model, provides that torque is proportional to current

and independent of velocity; a multiplicative torque ripple of a few percent may be included. The PUMA motors in particular, which are Magnetic Technologies model 3069-381-016, are specified to have a torque constant of .261 Newton-meters (N-m)per Amp, and an average to peak ripple of 4% at 25 cycles per revolution. The maximum available motor torque is 1.70 N-m; reflected through the transmission, this gives a maximum torque of 100 N-m at joint 1.

2.5 Break Away Torque

The break-away torque was measured by stopping the arm, setting the torque to zero (or to the gravity compensation value, if any) and ramping up the torque in 0.2 Newton-meter steps at a rate of 40 steps per second. Break-away was established when motion (one or more shaft encoder pulses) was observed during two consecutive steps. The torque applied during the first interval with motion was taken to be the break-away torque. Highly variable "wind-up" was observed in the time prior to steady motion, and labeling the first shaft encoder pulse to arrive as the break-away event led to non-repeatable results. The step size and rate of this experiment were chosen to maximize repeatability.

2.6 Linear Parameter Estimation

When parameters to be identified appear linearly in a model, as will be the case in the main of these experiments, solving the normal equations produces an estimate of the parameters which minimizes the squared error. The estimation is constructed in the following way:

If the model may be written

$$\tau(k) = \phi'(k) \theta^* \quad (2.1)$$

where $\tau(k)$ is the torque applied at time step k ;
 ϕ is the regressor vector, it contains functions of the manipulator state;
 θ is the parameter vector, it contains fixed parameters;
 θ^* is the vector of true parameter values;
 $\hat{\theta}$ is the estimated parameter vector;
 k is the time step in the process;

then the normal equations may be constructed by writing

$$\mathbf{T} = \begin{bmatrix} \tau(0) \\ \vdots \\ \tau(K) \end{bmatrix} \quad \Phi = \begin{bmatrix} \phi'(0) \\ \vdots \\ \phi'(K) \end{bmatrix}$$

and estimating the parameters according to (see section II, chapter 11.7):

$$\hat{\theta} = [\Phi' \Phi]^{-1} \Phi' \mathbf{T} \quad (2.2)$$

When the normal equations are used to estimate friction parameters, ϕ will commonly comprise the joint acceleration, velocity and sign(velocity):

$$\phi = \begin{bmatrix} \ddot{q} \\ \dot{q} \\ \text{sign}(\dot{q}) \end{bmatrix} \quad \theta = \begin{bmatrix} \text{mass} \\ \text{viscous friction} \\ \text{kinetic friction} \end{bmatrix}$$

So that

$$\begin{aligned} \tau(k) = & \text{mass} * \ddot{q}(k) + (\text{viscous friction}) * \dot{q}(k) \\ & + (\text{kinetic friction}) * \text{sign}(\dot{q}(k)). \end{aligned}$$

Once the parameters are computed, it is possible to compute the residual error according to

$$\tilde{T} = T - \Phi \hat{\theta} \quad (2.3)$$

where $\tilde{T}$ is torque not predicted by the model.

2.7 Non-Linear Estimation

In chapter 5 velocity parameters which appear in an exponential friction model are estimated. This was done by a simple, manual, gradient search method. The cost function minimized was squared model error; the gradient search continued until no further improvement was possible.

2.8 Low Pass Filtering

In some cases below, the data presented have been low pass filtered. To avoid the phase lag associated with causal filter implementations, this has been done off line using convolution with a gaussian curve. Where a cut-off frequency is cited, this is the width to half maximum of the filter frequency response.

Chapter 3

Repeatability

Perhaps the most fundamental issue in an effort to model any process is repeatability. Whatever hope exists of capturing the process in a predictive model is predicated upon repeatability. Tribologists have studied the variance of friction forces in carefully controlled situations [Rabinowicz, *et. al.* 55; Rabinowicz 56]. But it is not straightforward to extrapolate from these data to complex mechanisms. Repeatability may be inferred from the successful compensation of friction demonstrated in [Gilbert and Winston 74; and Walrath 84]. Walrath in particular presents correlation data that suggest a high degree of repeatability. But neither of these investigators explicitly addresses repeatability.

Here I undertake to measure repeatability in the simplest possible way: by playing out a pre-determined sequence of torques and observing the motion. A stiff PID controller, the standard Unimate controller, is used to pre-position the arm. The industrial controller is able to attain a desired position to within ± 0.001 radians. Using precomputed torques, as opposed to closed-loop control, ensures that the applied torques are the same from trial to trial. A typical result is presented in figure 3.1 where the velocity profiles of three successive motions are shown.

Prior to collecting these data the arm was 'warmed up' by a minute of motion throughout the workspace. The effects and importance of warming up the




Figure 3.1 Velocity resulting from the application of a constant torque to joint one of a PUMA 560 robot.

mechanism are discussed in chapter 4.1. The torque sequence used is shown in figure 3.2. The three velocity profiles of figure 3.1 show considerable structure: the arm repeatedly accelerates and decelerates in a roughly sinusoidal pattern. If the velocities are plotted as a function of position, the correspondence becomes more dramatic, as shown in figure 3.3.

Using accelerometer data it is possible to determine the mass, viscous friction and kinetic friction observed during the motion by solving the normal equation, (2.2). Equation (2.3) then gives the residual torques, the torques which are unexplained by the model. The residual torques for the motions of figure 3.3 are plotted in figure 3.4, here they are shown as a function of position.

By sampling laterally across several data sets, it is possible to estimate the variance in friction torque. To produce figure 3.5 the residual friction torques




Figure 3.2 Torque Applied during each of the three motions depicted in figures 3.1.




Figure 3.3 Velocity Profiles of Figure 3.1 plotted against position.




Figure 3.4 Torques that are Unexplained by a Mass plus Viscous plus Kinetic Friction Model, plotted against position.
(Low Pass Filtered to 20 Hz.)

were determined as they were for figure 3.4. At each time step the residuals of five motions (including the three of figure 3.4) were averaged and the variance determined. Figure 3.5 is a plot of the average across these five moves; note that large residual torques occur at the beginning and end of the trajectory. Figure 3.6 is more interesting; it is a plot of the standard deviation in each of the bins made by the averaging process which gave figure 3.5. The standard deviation is remarkably flat.

Figure 3.4 shows a high degree of position dependence in the torque error. This dependence is confirmed in figure 3.7, which is a plot of average and deviation in the variance of the five trials, as in figures 3.5 and 3.6, but here they are averaged and plotted as a function of position. During the cruise portion of the trajectory, the variance of figure 3.7 is quite small. In table 3.1 the magnitude of the friction torque,




Figure 3.5.a Average of the Unexplained Torque During Five Repetitions of the motion of Figure 3.1. The Data have been binned by Time. Scale chosen to show peaks.




Figure 3.5.b Average of the Unexplained Torque During Five Repetitions of the motion of Figure 3.1. The Data have been binned by Time. Scale chosen to show detail and match figure 3.6.




Figure 3.6 Standard Deviation of the Unexplained Torque During Five Repetitions of the motion of Figure 3.1. The Data have been binned by Time.

the residuals and residual deviations are presented. The deviation of the residuals is of greatest interest: the deviation is a measure of the non-repeatable component of the friction process. The degree of repeatability seen here is remarkably great, the deviation in friction between identical trails is only 1% as great as the magnitude of the friction itself: 0.107 (N-m) of Deviation versus 10.407 (N-m) of Friction.

The reduction deviation arising from correlating the friction data with position, from $\sigma = 0.284$ to $\sigma = 0.107$ corresponds to a Fisher statistic of $F_{(1000,500)} = 12.1$. Use of the Fisher statistic is described in greater detail in chapter 5.1 below. Let it suffice here to say that the F statistic can be used to test whether an improvement in model accuracy is achieved by chance or by actually explaining an underlying systematic process. In this case an $F_{(1000,500)}$ of 12.1 yields a confidence of 99.99999999% that a systematic correspondence between position and friction exists.

In this chapter a fundamental aspect of friction has been addressed: repeatability. Surprising repeatability is observed. Study of the friction force that is

Table 3.1 Magnitude of Friction Torque Captured by a Kinetic plus Viscous Model and of the Residuals and the Deviations of the Residuals.

Friction Signal	Magnitude (N-m)
Modeled Friction	10.407
RMS Residual, Sampled According to Time Along the Trajectory	0.282
Mean Standard Deviation of the RMS Residual, Time Sampled	0.284
RMS Residual, Sampled According to Position Along the Trajectory	0.379
Mean Standard Deviation of the RMS Residual, Position Sampled	0.107




Figure 3.7 Average and Deviation of the Unexplained Torque During Five Repetitions of the motion of Figure 3.1. The Data have been binned by Position.

not explained by a kinetic plus viscous friction model shows a substantial component which, with very high confidence, is systematically correlated with position. Measuring this position-dependent friction will be the subject of the next chapter.

Chapter 4

Break-Away Experiments

The break-away experiments were conducted as described in section 2.5. The beauty of the break-away experiment is its ability to sample friction at extremely high spatial frequency. The break-away experiment does not suffer from problems of mechanical compliance, phase lag and instrument bandwidth that are associated with motion experiments. The data are useful because they align so well with the position-dependent disturbance observed in the motion data. An example of the alignment between break-away data and the unexplained friction in the motion data is shown in figure 4.1, where the solid trace is the error torque experienced during a motion, similar to that presented in figure 3.4. The dashed trace is taken from break-away data, suitably scaled and offset to align with the motion data. Both traces have been low pass filtered to 20 cycles per radian. As an aside, the break-away data were collected on June 6, 1987, and the motion data on January 12, 1988, six months and a thousand machine operating hours later.

4.1 Experimental Issues in Measuring Break Away Torque

Four aspects of this experiment were tuned to achieve maximum repeatability: a warm up exercise was developed, fine sampling was used, the torque rate was tuned, and the definition of break away selected to maximize repeatability. Early in




Figure 4.1. The alignment of break-away friction data and Motion Friction Data.

this experimental work it was observed that after an extended period of inactivity (over night) the friction would start high and diminish quickly to a steady value. A series of empirical trials showed the motion friction to come to its steady value after a minute of moderate velocity activity spanning the work space. The motion friction would remain at its steady value through tens of minutes of inactivity. The break away friction was observed to increase with over night inactivity, as does the motion friction. A faster process, affecting the friction after minutes of inactivity, was also observed in the break-away experiment. The underlying physics of neither process is known to the author; as a hypothesis, I suspect that the longer time process is due to oxidation at the surface of the lubricant, and that the shorter time process is due to the physical displacement of lubricants as part of the Stribeck effect (see chapter 5.5). The longer time process was addressed with a vigorous motion at the beginning of each session. While sampling the break away friction, a

shorter back and forth motion, spanning the space to be sampled during the coming interval, was carried out about once per minute.

The break away data were observed to have features at extremely high spatial frequencies. The initial sampling density was inadequate and problems arose from aliasing and bias introduced by the sample location selection. To collect break away data at a spatial resolution higher than that achievable by the standard industrial controller, an open type sample location selection was used. The Unimate controller can achieve commanded positions to within ± 0.001 radians, whereas the shaft encoder has a resolution ten times finer. Rather than attempting to control the motion to a desired sample point, a sample was simply collected at the location of the arm. Each sample moves the arm forward a distance of several bins. After many (possibly ten) sweeps over each region nearly all of the locations will have been sampled. Experience showed that some locations were unreachable, either by feedback control or by chance !! I believe that the steep friction gradient surrounding these points made them unreachable. Quadratic interpolation was later used to estimate the friction at these locations. For the table building experiments, 2,000 bins per radian were used. The highest prominent spatial frequency observed is 355 cycles per radian; thus 2,000 samples per radian provides nearly a factor of three head room above the Nyquist limit.

The break away experiments required considerable time (60,000 data points $\times$ 2 directions $\times$ 3 joints $\times$ 2 seconds per point = 200 hours). It was therefore important to maximize the rate of the experiment. As roughly 80% of the experimental time was spent ramping up the torque, the size of torque step and step rate were critical experimental issues. The results of several trials are presented in table 4.1; 0.2 Newtons per step and 40 steps per second was chosen as a good compromise

between repeatability and speed. The result of this effort toward achieving good repeatability in the break away measurements is shown in figure 4.2. The three lines of the figure are the mean measured break away force in each position bin, and plus and minus one standard deviation. The mean break away force of figure 4.2 is -8.43 N-m, the mean standard deviation within the bins is .237 N-m, or 3% of the mean friction.

Table 4.1 Measurement of Deviation of Break Away Friction Measurement as a Function of the Experiment Parameters.

Sample Rate	Torque Step	Mean Break Away	Standard Deviation
(Hz)	Size (N-m)	Torque (N-M)	(N-m)
40	.2	8.51	0.165
40	.1	8.63	0.151
40	.4	9.00	0.330
80	.2	9.10	0.220
80	.1	8.86	0.124

Initially the break away torque was taken to be that applied when the first shaft encoder pulse was observed. This led to highly non-repeatable results. Trials were made requiring two, three and four shaft encoder pulses to certify motion and determine the break away torque, but with little improvement in repeatability. Imposing what is in essence a velocity requirement, that shaft encoder pulses come in two consecutive intervals, markedly improved the repeatability and was adopted as the condition to certify break away. The combination of 0.1 N-m steps at 80 Hz, which is suggested by table 4.1, was rejected because of the velocity and thus acceleration required to achieve two shaft encoder pulses in consecutive sample intervals. The lack of repeatability in the observation of the initial pulse has been connected to the Dahl effect (see chapter 5.4).




Figure 4.2 A Portion of the Break Away Data on a Greatly Expanded Position Scale. The Mean and Plus and Minus One Standard Deviation Curves are shown.

4.2 Building the Correction Table

The break-away experiment provides an excellent means to fill out a friction correction lookup table. To capture phenomena at spatial frequencies as high as 355 cycles per radian a sampling frequency of 2,000 bins per radian was selected. To prepare a lookup table for 5 radians of joint motion, 10,000 bins were needed. Collecting 60,000 data points allowed averaging and improved repeatability.

Examination of the curves of figure 4.1 suggests that there should be a few dominant spatial frequencies. The major periodicity makes 3 cycles across the figure, and a less prominent periodicity appears at about 5 cycles per cycle of the major periodicity. A spatial FFT of the break-away data fails, however, to reveal these periodicities, as shown in figure 4.3. Figure 4.3 is a spatial FFT of the joint 1 break-away data set, sampled at 2,000 samples per radian, low pass filtered to 100

cycles per radian and pre-processed with a Blackman-Harris window. The result is a flat transform, not showing the dominant periodicities that the eye sees in figure 4.1. The major periodicity of figure 4.1 makes 2.8 cycles per radian or 15 cycles per revolution of the joint. This is the rotation frequency of the intermediate gear of joint one, suggesting that periodic friction phenomenon is related to a non-uniformity of the intermediate gear. The higher frequency apparent in figure 4.1 occurs once per revolution of the motor, suggesting a non-uniformity in the motor or motor gear.




Figure 4.3 Spatial FFT of the Break Away data for Joint 1 of a PUMA 560 Arm. Note the lack of Periodicities Apparent in Figure 4.1 or Figure 4.7.

To correct for motion error a lookup table was constructed using data from the break-away experiment. The velocity profile of a position friction compensated open-loop motion is shown in figure 4.4. Applying the position-dependent compensation

from the lookup table reduced the variance of the velocity during the motion. It was observed that the table correction was most effective at lower spatial frequencies, the variance in velocity and acceleration was measured with correction tables applied that were low pass filtered to dc (0 Hz), 0.8, 4.0, 20, 100, and 500 cycles per radian (data collected at 2,000 samples per radian). The deviation (square root variance) in velocity measured along the motion of figure 4.4 is shown as a function of the table bandwidth in figure 4.5. In figure 4.6 the deviation in acceleration is shown.




Figure 4.4 Velocity Profile Recorded During an *Open-Loop* Motion of the Arm with Position-Dependent Correction Applied.

The best reduction in velocity deviation is roughly a factor of 2. There is no apparent reduction in acceleration deviation. The table lookup scheme predicted the position one and a half sample times ahead, to try to select the position corresponding to the middle of the sample interval during which the torque would




Figure 4.5 RMS Deviation in Velocity as a function of the Highest Spatial Frequencies of the Position-Dependent Friction Correction Used.




Figure 4.6 RMS Deviation in Acceleration as a function of the Highest Spatial Frequencies of the Position-Dependent Friction Correction Used.

be applied. The break away data were low pass filtered to 20 cycles per radian and then compressed into a look table with a sampling frequency of 200 bins per radian. At the moderate arm velocity of 1 radian per second, the arm would cross only one bin during one 5 milli-second cycle of the control system. No attempt was made to average the correction from several bins according to the velocity of the arm.

A low pass filter of 20 cycles per radian was chosen for the correction table. An offset was added to make the lookup table zero-mean: kinetic friction correction is provided as a separate term. The resulting correction torque table is shown graphically in figure 4.7, this correction was used to generate the motion of figure 4.4, the velocity profile of which should be contrasted with those of figure 3.1, for which no correction table was used.




Figure 4.7 Lookup Table Correction to Torques Applied at Joint One of the PUMA 560.

In this chapter an experiment that can accurately measure friction as a function of position is presented. Again the repeatability is striking, this time the repeatability of the torque required to achieve break-away. The importance of modeling the position-dependence in the friction for control will depend upon the application. But the importance of modeling this effect for other experimental work is tremendous; the model of position dependence will be used many times in what is yet to come.

Chapter 5

Friction as a Function of Velocity (Negative Viscous Friction Revealed)

The control engineer often assumes a viscous friction model. The tribologist rarely includes any velocity dependence at all, and when he does it is as likely to be negative as positive. Based on the experimental work presented here, the kinetic plus viscous friction model seems very accurate at velocities above a minimum critical velocity. Below the critical velocity, decreasing friction with increasing velocity is observed. The Stribeck effect is strongly indicated as the cause of this unusual friction behavior.


The relationship between friction and velocity has been measured in three ways:

1. Open-loop, steady torque;
2. Closed, stiff velocity loop, constant velocity;
3. In contact, an open-loop torque ramp against a compliant surface.

Each of these experiments offers an advantage: the open-loop gliding motions provide a test bed for measuring the repeatability of the friction forces - a basic objective of this research - and to test for the dependence of motion friction upon position and load. The closed loop control provides a means to penetrate the

unstable regime with a realistic control experiment. And the compliant motions - the most successful experiment of this research - provide a means of mapping the entire regime of unstable velocities.

Figure 5.1 is a plot of friction as a function of velocity, measured on joint one of the PUMA arm. These data were collected by applying a steady torque and observing the average resulting velocity. The position-dependent friction was compensated by a zero-mean table lookup torque.



Negative Direction
Viscous 3.45 N-m/rad/sec
Kinetic -8.26 N-m

Positive Direction
Viscous 4.94 N-m/rad/sec
Kinetic 8.43 N-m

Figure 5.1 Friction Torque as a Function of Velocity. (Note that the kinetic and viscous parameters are different for positive and negative direction rotations.)

5.1 Analysis of Variance in the Motion Friction Data

Figure 5.1 is based on five measurements of velocity, at each of four positions, at each of four steady state torques, in the positive and negative directions: a total of 160 measurements. This data set is sufficient to allow the use of standard analysis of variance techniques to test whether separate positive and negative direction parameters should be used, and to test whether the motion friction is dependent upon position. The analysis of variance (ANOVA) is a technique that allows the variance explained by the extension of a model to be compared with the residual variance. When linear estimation is used, adding parameters to a model will nearly always result in an improvement in the fit of the model to the data. Using ANOVA and the F statistic, a standard test of significance, it is possible to determine whether the improvement given by adding model parameters is significantly better than chance, and thus justified. The test boils down to comparing the improvement achieved per degree of freedom added to the model, the MST, with the residual variance per degree freedom, MSE. If the extra parameters have nothing whatever to do with the source of variance (which may simply be random), one would expect the F statistic, $F = MST/MSE$, to be nearly one. A large F statistic indicates that a lot of variance is explained by the extension of the model and that the increased model size is justified. The threshold F statistic is a function of the confidence level required and the number of degrees of freedom involved.

The hypotheses to be tested are:

Hypothesis 1: The model should include separate kinetic and viscous friction parameters for the positive and negative rotation directions.

Hypothesis 2: There is position dependence in the kinetic and viscous parameters, and separate parameters should be used for each of the four starting points used. This would be a position dependence arising during motion, separate from that corrected for by the table lookup compensation.

The variance measures are defined to be:

MST: Mean squared error explained by the use of a larger model.

MSE: Mean squared error still unexplained, even with the use of a larger model.

Table 5.1 ANOVA Testing of Hypothesis 1 and 2.
(Squared errors are in units of torque (N-m) squared.)

	MST	MSE	F	Threshold F, 95% Confidence	Accept Hypothesis?
Hypothesis 1	13.209	0.00605	2182.0	2.15 (F 2,28)	Yes
Hypothesis 2					
Positive Direction	0.00475	0.00295	1.61	9.03 (F 6,8)	No
Negative Direction	0.00828	0.00847	0.98	9.03 (F 6,8)	No

The use of positive and negative direction kinetic and viscous parameters is strongly indicated (Hypothesis 1). The position dependence of the kinetic and viscous parameters (Hypothesis 2) is not supported, the F value of 1.61 is little better than chance and the F value of 0.98 is worse than chance.

5.2 Friction at Low Velocities

Stick-slip oscillation during low speed motion has been observed by many investigators; and a sizable literature exists studying the contribution of static friction to this oscillation: [Tou and Schultheiss 53; Satyendra 56; Bohacek and Tuteur 61; Shen 62; Shen and Wang 64; Kubo, Anwar and Tomizuka 86; Townsend and Salisbury 87]. In all of the above reports, describing function analysis is applied to a simple sticktion model, such as that of figure 1.2; the outcome is oscillations when integral control is used. None of the work based on a simple sticktion model can explain oscillation without an integral control term. It is odd that none of the authors recommends control without an integral term.

During unrelated experimental work at Stanford [Khatib and Burdick 86] unstable motions were observed even when no integral control term was used. This apparent conflict between theory and practice remained a quandary until a much-neglected paper by Tustin [47] came to light. Tustin studies this effect under the assumption that friction follows an exponential curve from its static value to its kinetic value. Tustin's model is:

$$F(v) = F_0 - F_c(1 - e^{(-v/V_c)}) + F_v v \quad (5.1)$$

where $F(v)$ is the friction as a function of velocity;

F_0 is the static friction;

F_c is the difference between static friction and kinetic friction;

F_v is the viscous friction parameter;

v is the motion velocity;

V_c is a constant with units of velocity giving the characteristic velocity at which the system transits to kinetic friction.

Tustin predicts oscillations at low speed, even in the absence of integral control. Tustin provides no experimental evidence, either direct or cited, to support his choice of model structure. None-the-less, Tustin's contribution stands out: it explains an observable phenomenon neglected by a host of authors.

Tustin's hypothesis suggested an important path to pursue. To measure the motion friction at low velocities, a stiff velocity control loop was implemented and the average torque required to sustain steady motion was measured. With a sample rate of 200 Hertz and using the first difference of position as an estimate of the velocity, it was not possible to implement a velocity error gain greater than 50 Newton-meters (N-m) per rad/sec. This gain permitted a minimum velocity of 0.15 rad/sec before the onset of stick-slip. By integrating the accelerometer signal, the velocity error gain could be increased to 90 N-m per rad/sec, which allowed testing motions as slow as 0.015 rad/sec. A velocity of 0.012 rad/sec was achieved with direct accelerometer feedback. Note that using only the position measurement and a simple control structure (the standard industrial implementation), the apparent minimum velocity of 0.15 rad/sec is 7% of the maximum velocity achievable at the joint.

The measurements made with a stiff velocity loop are shown in figure 5.2, a downward bend in the low velocities is clearly evident. Five samples were taken at each velocity, from this the variance was estimated giving the 90% confidence intervals shown.

Fitting Tustin's model to the variance weighted data gives $V_c = 0.019$ rad/sec. The friction is given by:

$$F(v) = 9.56 - 1.13(1 - e^{(-v/0.019)}) + 4.94 v \quad (\text{N} - \text{m}).$$

(Viscous friction term taken from the fit to high velocity data).




Figure 5.2 Friction as a Function of Velocity at Low Velocities, fit with Tustin's Exponential Model.

5.3 Friction During Compliant Motions

The data collected by closed loop motions showed clear evidence of an upward turn in the friction curve at low but significant velocities. However the model and experiment were unsatisfactory in two respects: the model did not fit the data well, the lowest velocity data point showing a 2 standard deviation excursion; and no measurements could be made at extremely low velocities, a velocity regime important to force control. Data collected in a preliminary force control experiment

showed both a structural flaw in the model and a means to collect data in the extremely slow regime. Figure 5.4 shows the result of an open-loop force move. In this trial, the arm, with a force sensor, was pressed against a hard surface, and the torque is ramped up to a value well above static friction, and down again. The configuration of the arm, actuation and sensor are shown in figure 5.3. A number of features are evident in this motion, but one stands out: *there is no stick slip behavior in the resultant force trajectory*. Tustin's model predicts stick slip down to zero velocity. The shaft encoder data from the trial of figure 5.4 is shown in figure 5.5; the motion is extremely slow, about 330 micro-radians (μR) per second.




Figure 5.3 Apparatus of the Compliant Motion and Force Control Experiments.




Figure 5.4 Applied Torque and Measured Force (scaled by radius) for a Motion in Contact with a Hard Surface.

Figures 5.4 and 5.5 show that by applying varying force between a manipulator and a compliant surface very slow motion can be achieved. The motion velocity is given by

$$\hat{q} = (1/k) \times \dot{\tau}$$

where $\hat{q}$ is the estimated velocity;

k is the stiffness of the combined manipulator and environment;

$\dot{\tau}$ is the torque rate.

Data were acquired by this means using a number of coil springs. The springs yielded effective stiffness from 450 to 12,000 N-m per radian. Applying a torque




Figure 5.5 Position Reading During the Motion of Figure 5.4.

ramp with a rate of 10 N-m per second yielded velocities ranging from 0.0008 to 0.022 radians per second. The study of low speed friction required exceptionally sensitive measurement of velocity and careful accounting of forces. The data analysis led to the observation of the Dahl effect [Dahl 68]. The Dahl effect, which has substantial importance for force control, had to be compensated to achieve an accurate measurement of low velocity friction.

The first experimental task was to develop the ability to measure velocities of hundreds of micro-radians per second. Careful study of the noise processes of the rotational accelerometer and the force sensing fingers was required. Systron Donner specifies a threshold and resolution for the loaned accelerometer of 0.005% of full range, substantially below the quantization of the twelve bit A/D converter. The range of the A/D converter and full scale of the instrument are well matched; but when the output signal of the instrument is small, it is possible to make better use

of the resolution of the A/D converter by increasing the gain of the analog amplifiers buffering the accelerometer signal. A computer controlled variable gain amplifier permitted an 8x multiplication of small accelerometer signals thus reducing the effect of quantization. Simpson's rule integration replaced Euler integration in off-line processing to improve the estimate of velocity; and a sensitive bias null that used an average of many samples was implemented. The combined effect was an order of magnitude reduction in the error in position estimated from the second integral of acceleration, from 1 part per thousand to 1 part per ten thousand.

During compliant motion with constant stiffness, velocity can also be estimated from the derivative of force:

$$\hat{q} = (1/k) \times \dot{f}$$

where $\dot{f}$ is the derivative of measured force.

Note: The friction, applied torque and torsional stiffness are measured in rotational units. Here the Contact forces and stiffness are measured in linear units. Conversion is made by scaling by the appropriate radius, 0.83 meters for most of this work.

As the stiffness, k , is large, reasonable force rates correspond to very small velocities. The force fingers proved to be substantially more sensitive than was previously thought. Adjusting the signal conditioning amplifier gain to give a good match to A/D converter range resulted in an RMS noise power of 0.01 Newtons, mostly at 60 Hz. The fingers however, suffered from considerable slow drift. These effects gave changes in bias in the range of 0.1 Newtons; careful attention to maintaining the null reading of the sensor was required.

Figure 5.6 shows both the estimate of velocity derived from the acceleration signal and that derived from the force rate. The correspondence is fantastic, even

at extremely low velocities. This gives good confidence that both methods of measurement are quite accurate. The integrated accelerometer signal, merged with the derivative of the position signal in a cross-over filter, was used as the velocity estimate.




Figure 5.6 Velocity Measured during slow motion by Integration of Acceleration and by Differentiation of Contact Force.

In addition to accurate force and velocity sensing, the measurement of friction at extremely low velocities depended upon accurate accounting of other forces affecting the motion. In the simple model of compliant motions, friction is given by:

$$\begin{aligned} \text{Friction Torque} = & \text{Applied Torque} - \text{Measured Contact Torque} \\ & - \text{Motion Torque.} \end{aligned}$$

The applied torque is known, the contact torque may be determined from the force sensor measurements by scaling by the moment arm; and the motion torque is computable from the measured accelerations and the known inertia. The resulting estimate of friction force for an early trial is shown in figure 5.7. By combining the velocity and friction data, a plot can be made of friction as a function of velocity, as shown in figure 5.8. At first it was believed that the open curves of figure 5.8 were the result of a hysteresis like process in the friction, that the friction level was somehow lower during deceleration than acceleration; simple lack of repeatability was also suspected as the source of the spread. However, study of several different trials frustrated explanation in terms of a sensible, velocity-related process; and examination of several identical trials showed the repeatability to be quite good. Insight came from noting that the spread was greater for motions with softer springs, i.e., those that covered greater distance. Comparison of the spread with the position-dependent correction of chapter 4 showed a high correlation. Figure 5.9 is the result of compensating for the position-dependent friction with:

$$\begin{aligned} \text{Friction Torque} = & \text{Applied Torque} - \text{Measured Contact Torque} \\ & - \text{Motion Torque} - \text{Position-Dependent Correction.} \end{aligned} \tag{5.2}$$

Combining the friction measurements of a number of trials with different spring stiffnesses gives figure 5.10, not yet a repeatable measurement of friction as a function of velocity. Again the study of motions carried out with different spring stiffnesses provided the clue to sorting out simultaneous processes. The several




Figure 5.7 Compliant Motion Friction Force Plotted Against Time.

curves of figure 5.10 show different angles of rise from zero friction at zero velocity to a relatively steady friction level at higher velocities. The length of this rise was seen to be variable in terms of time, and variable in terms of velocity, but roughly constant in terms of position change. This observation led to the hypothesis that the Dahl effect was behind the scatter of friction measurements.




Figure 5.8 Compliant Motion Friction Force Plotted Against Velocity.




Figure 5.9 Data of Figure 5.8 Corrected According to Equation 5.2.




Figure 5.10 Friction data from several Trials, Corrected According to Equation 5.2, but not Corrected for the Dahl Effect.

5.4 The Dahl Effect

Dahl's work with friction analysis and modeling is the work on friction best known to the controls community [Dahl 68, 76, 77]. After studying experimental work with the friction of ball bearings done by H. Shibata at the Aerospace Corporation, Dahl proposed that a process of plastic deformation and failure occurs during the transition from static to kinetic friction. These processes had been previously described and ingeniously demonstrated by radiological techniques [Bowden, Moore and Tabor 43; Rabinowicz and Tabor 51]; Dahl's development is apparently independent. Intuitively the process is this: the asperities of one surface weld to those of the other when the two are in contact for an adequate time. When motion begins the separation is not immediate, but rather the asperities deform in a fashion comparable to that of the bulk material. This process entails both elastic

and plastic deformation, and leads to hysteresis if the motion is reversed. Figure 5.11 is taken from [Dahl 77], friction is plotted as a function of displacement - *not* velocity - for a cyclic motion of the bearing. Over the linear region of the curve, that near zero displacement, the process is conservative. A ball bearing, for example, can be observed to rock back and forth on a flat hard surface under the influence of this effect [Dahl 77].




Figure 11. New Bearing Hysteresis Loop.

Figure 5.11 Friction Plotted as a Function of Position for a Ball Bearing Apparatus (taken from [Dahl 77], reprinted courtesy of author). Note the linear relationship between force and displacement for small displacements.

Rabinowicz [51] did a brilliantly simple experiment to measure the transition distance from static to kinetic friction. With an apparatus not unlike that of Leonardo da Vinci, Rabinowicz studied the sliding motion of blocks on a flat inclined plane. He would initiate motion by rolling a ball a prescribed distance down the plane to strike the block, thus delivering a known impulse. The relationship between impulse, sliding distance, plate slope and lubricant yielded information about the transition from static to kinetic friction, among other things. For lubricated smooth metal surfaces, the transition distances lie in a range from 1 to 7 microns (μm), roughly the size of the surface asperities (the micro-bumps). Shibata's data show a typical transition distance that is much greater, roughly 200 microns. This inconsistency and Rabinowicz's evidence for transition on a scale too small to be observed discouraged investigation of the Dahl model in this work. Only when examination of the compliant motion data strongly suggested the presence of the Dahl effect on an observable scale was the possibility studied.

The inconsistency of distance scales is resolved by considering that Rabinowicz was studying the motion of blocks on flat surfaces while Shibata was studying the motion of ball bearings. When a ball is rolling, the relative translation across the interface between the ball and contacting surface (due to part deformation) is much smaller than the motion of the center of the ball. Dahl's friction model is not predictive. Neither Dahl nor investigators who have employed his model [Walrath 84] have attempted to relate the transition distance of the Dahl model to the micro-morphology of the surfaces in contact. However, such a relationship may be derivable, as suggested by table 5.2. The transition distance for Dahl [77] was computed from the geometry of the apparatus used: 6 mm balls rolling in a 40 mm diameter race. The ball translation distance of 200 microns corresponds to

a separation of 6.7 microns between opposing faces of the friction interface. For the simple, inelastic model the translation is mostly strain rather than shear. A detailed study of the elastic contact (Hertzian contact) should be conducted to relate ball translation to shear between the boundaries and the distance of the Dahl effect to the size of the asperities. The transition distance for joint 1 of the PUMA arm was computed from the 0.19 meter diameter of the bull gear, assuming a 10% slip coefficient for the gear cut. The externally observable transition distance was 0.0003 radians. These three measurements of transition distance all lie near the typical asperity dimension of finished hard metals.

Table 5.2 Transition Distances from Static to Kinetic Friction.

Investigator	Transition Distance	Interface Characteristics
E. Rabinowicz [51]	1 to 7 μm	Various Metals Sliding on Steel
P.R. Dahl [77]	6.7 μm	Industrial Ball Bearing, 6mm Balls
B. Armstrong	1.9 μm	Bull Gear of the PUMA Joint 1

5.5 The Stribeck Effect

The Stribeck effect appears at low velocity and gives rise to friction which decreases with increasing velocity. To observe the Stribeck effect alone, it was necessary to account for the influence of the Dahl effect in the data. Rather than attempting to model the Dahl effect accurately, it was removed from the data by throwing out all data points taken near to a zero crossing in velocity. Welding occurs at zero velocity and affects the mechanism behavior for some translation away from the point of welding. When a sufficient distance is traversed in steady motion, the Dahl effect is no longer a substantial influence. A number of distances were tried, five shaft encoder ticks (0.0005 radians) was found to be effective. This

filtering step, removing all data points collected within five shaft encoder ticks of a velocity zero crossing, I call "de-Dahl'ing".

In the final step of processing, the friction measurements were placed into bins according to velocity. Only those bins containing 30 or more points were retained, a requirement of sufficient sample size to achieve a good estimate of the mean. The result of the "de-Dahl'ing" and sample size filters, applied to the data of figure 5.10, is shown in figure 5.12. The separate lines indicate data collected by different combinations of spring stiffness and torque rate. Binning the data of figure 5.12 according to velocity and computing the mean and deviation give figure 5.13. The negative dependence of friction upon velocity is evident at a level many times the uncertainty in the measurements. The 90% confidence interval shown is taken from the distribution of friction measurements within each bin. The total data set giving figure 5.13 is 4,000 data points, taken in 9 separate motions involving 5 different springs, ranging from 770 to 3400 Newtons per meter stiffness, and 3 different torque rates: 5, 10 and 15 Newton-meters per second. The curves of figure 5.13 are shown in expanded view in figure 5.14.

The model curve shown in figure 5.13 was created by fitting an exponential function of velocity squared to the data. The model function is:

$$F(v) = F_0 - F_{c1}(1 - e^{(v/V_{c1})^2}) - F_{c2}(1 - e^{(v/V_{c2})^2}), \quad (5.3)$$

minimizing the squared error yields:

$$F(v) = 10.11 + 4.97 v - 1.02(1 - e^{(v/0.0061)^2}) - 0.71(1 - e^{(v/0.048)^2}). \quad (5.4)$$




Figure 5.12 Data from several Trials Corrected for the Dahl Effect.

I believe that the presence of two break points reflects linear sliding occurring at two different velocities in the drive train.

This fit was done to a data set that merged data from the open-loop, constant-torque motions, closed loop, constant-velocity motions and the compliant motions. The values of kinetic friction demonstrated by these experiments were observed to be different and an adjusting factor of 0.49 N-m was added to the constant velocity data and 1.94 N-m was added to the compliant motion data. This variability in the kinetic friction may be due to the separation in time between the experiments. Drift in the kinetic friction parameter was observed and is discussed in relation to adaptive control, though it was generally less than 10% or 0.8 N-m. An experimental bias affecting the compliant motion data is indicated by the large shift in kinetic friction relative to the constant torque and constant velocity motion data. The




Figure 5.13 Friction as a Function of Velocity During Motions Against a Compliant Surface.




Figure 5.14 Friction as a Function of Velocity During Motions Against a Compliant Surface (Expanded Scale).

merged friction data set and the model curve are shown in figure 5.15, which spans the velocity range of joint one.




Figure 5.15 Friction as a Function of Velocity; Compliant, Constant Velocity and Constant Torque Motion data Merged to Span Full Motion Range.

The exponential model of equation (5.3) is empirical. A model with an additional parameter in each exponential break was fit to the data, the additional parameter allowed the break point to be shifted in velocity:

$$F(v) = F_0 - F_{c1}(1 - \xi^{((v-V_{o1})/V_{c1})^2}) - F_{c2}(1 - \xi^{((v-V_{o2})/V_{c2})^2}) \quad (5.5)$$

$$\text{where } \xi(x) = \begin{cases} 1, & \text{if } x \leq 0; \\ e^{(x)}, & \text{if } x > 0. \end{cases}$$

But this increase in model complexity did not produce a substantial improvement in the fit. The Fisher statistic for the increased model complexity is $F(30, 2) = .63$; a Fisher statistic that is between zero and one indicates an improvement in the model fit, but an improvement that is less than chance.

The curve of figure 5.13 is an example of a generalized Stribeck curve [Stribeck 02; Rabinowicz 65; Czichos 78]. The effect has a rather straightforward tribological explanation: in a lubricated system, when the surfaces are not in motion the lubricant may be largely expelled from between the contact surfaces and metal to metal contact results; this is called boundary lubrication. When the surfaces are in motion at high velocity, they are fully separated by the lubricant, which remains between the surfaces because of its viscosity. In this regime, termed full fluid lubrication, viscous friction is in evidence, as shown in the high velocity portion of figure 5.15. At middle velocities, the sliding surface is only partially lifted by the lubricant, and metal to metal contact bears a portion of the load which diminishes with increasing velocity. This velocity regime is termed mixed lubrication and is marked by decreasing friction owing to the decreasing load borne by metal to metal contact. The three stages of the Stribeck effect are shown in figure 5.16.

Break-away and the transition to full fluid lubrication are not simple, single parameter processes; in addition to material properties and load, the transitions must depend upon residence time and acceleration. The fact that the data from 9 trials group so tightly in figure 5.12 gives confidence that the measurements accurately reflect the underlying process over a range of physical configurations. However, if the experiment were constructed in a different way the measurements might differ; though the underlying mechanism of the Stribeck effect and structure of the Stribeck curve would remain unchanged.




Figure 5.16 The Generalized Stribeck Curve. (From [Czichos 78], reprinted with permission of the publisher).

No predictive model of the Stribeck effect is available; a complicated interdependence arises between asperity size, lubricant viscosity and contact geometry. The viscosity of lubricants under pressure may provide the greatest barrier to analysis. It is now known that some lubricants, particularly metallic soaps (grease), become very viscous under extreme pressure; so much so that the lubricant itself may contribute to plastic deformation of the bearing metal surface, confounding analysis of the interaction. Thus, the functional form of the model, as well as the parameter values, are based on empirical observations. A. Sommerfeld, while he held the Chair of Technical Mechanics in the Technische Hochschule in Aachen, studied the viscous drag on a journal bearing, solving the Reynolds equation with the assumption of incompressible flow [Sommerfeld 04]. Sommerfeld's solution was later generalized by introduction of the Sommerfeld transformation. Sommerfeld's solution predicts a

curve with the same shape as that of the Stribeck curve, and probably provided the prototype, as Stribeck's own data mapped only the mixed and full fluid lubrication regimes. It is an interesting twist in the history of science that the "Stribeck curve" was never drawn by Stribeck, nor was the "Sommerfeld transformation" ever employed by Sommerfeld; both came along later as the original contributions were generalized [Dowson 79].

The Stribeck effect has profound implications for the stability of controlled mechanisms. An instrumentality capable of providing stable control in the region of mixed lubrication would be quite incredible. The derivative of the friction model of equation (5.3) is shown in figure 5.17. The derivative of friction with respect to velocity is proportional to the natural damping of the system: a negative derivative indicates an unstable system. The low velocity limit of the three control structures tested - shaft encoder feedback, integrated accelerometer feedback and direct accelerometer feedback - are indicated in figure 5.17. The degree of instability that must be compensated for at the lowest point of figure 5.17 is more than ten times as great as that achieved by direct accelerometer feedback. To allow a mechanism to operate stably throughout the velocity spectrum would require effective elimination of the Stribeck effect. But the effect is general to fluid lubricated interfaces, and may extend to dry lubricated interfaces where the dry lubricant deforms plastically under prevailing load conditions. In the absence of practical mechanisms free of the Stribeck effect, a *minimum* velocity for stable motion is an intrinsic mechanical characteristic.




Figure 5.17 The Derivative of the Friction Model of Equation 5.3.

In this chapter the relationship between velocity and friction is explored. In the moderate to high velocity regimes friction is seen to be quite linear with velocity with no position dependence but with separate kinetic and viscous friction parameters in the positive and negative rotation directions. Low velocity friction is explored, and two tribological effects are shown to be present and important: the Dahl effect and the Stribeck effect, both of which bode ill for fine control.

Chapter 6

Demonstrations of Friction Compensation

To demonstrate the accuracy of the friction model, open-loop moves were undertaken. For the spatial motions demonstrated in sections 6.1 and 6.2, friction was predicted using a three component model:

1. Kinetic Friction ;
2. Viscous Friction ;
3. A Table Lookup to Compensate for Position-Dependent Friction.

6.1 Open Loop Motion of One Joint

Figure 6.1 is a plot of the desired and actual robot motion during a 20 second oscillating motion. This motion was conducted in open loop: there is no feedback correction. The motion of figure 6.1 is in fact the *first* reciprocating motion conducted, friction compensation is based on a model developed as described in chapters 4 and 5; *not* upon a model identified on the trajectory shown. The only state-dependent portion of the torque was a position-dependent table lookup correction. In addition to the table lookup correction, a four parameter friction model was used to generate the torques for the plotted motion; the parameters are shown in table 6.1.




Figure 6.1 Desired and Actual Position during an OPEN-LOOP Motion made using Parameters Identified on different Trajectories.

Table 6.1 Parameters used for the motions of figures 6.1 and 6.2.

	The Motion of Figure 6.1	The Motion 2 of Figure 6.2	
Mass;	4.900	5.115	(Kg-M ²)
Viscous, positive rotation;	4.938	5.070	(N-m/rad/sec)
Viscous, negative rotation;	3.446	3.772	(N-m/rad/sec)
Kinetic, positive rotation;	8.432	8.121	(N-m)
Kinetic, negative rotation.	-8.262	-7.907	(N-m)

The friction parameters used for the motion of figure 6.1 are shown in table 6.1. They are those derived in chapter 5 above, which is to say that they are not specific to this trajectory. It is generally true that a set of parameters derived from a specific motion will afford better accuracy on that specific motion; the motion of figure 6.2 demonstrates this. The motion of figure 6.2 was conducted using parameters derived from data collected during the motion of figure 6.1. The final

position error at the end of the motion of figure 6.2 was only 1.02 deg, or 0.14% of the total motion distance. The motion of figure 6.1 provides, however, a better indication of the global accuracy of the open-loop friction correction; this motion shows an accumulated position error of 4.1% - still not bad for open-loop control.




Figure 6.2 Desired and Actual Position during an OPEN-LOOP Motion made using Parameters Identified on the same Trajectory.

The total torque applied to bring about the motion of figure 6.2 is shown in figure 6.3. In figure 6.4 this torque is broken down into the inertial and friction components. The RMS inertial torque is 3.62 Newton-meters, the RMS friction torque is 10.72 Newton-meters, nearly 3 times greater. In figure 6.5 the friction torque is further broken down into the kinetic, viscous and table lookup torques.




Figure 6.3 Total Torque Applied During the Motion of Figure 6.2.




Figure 6.4 Total Torque Broken Down into Inertial and Friction Torque Components. Note that the Inertial Torque is never great enough to overcome Static Friction.




Figure 6.5 Friction Torque Broken Down into Kinetic, Viscous and Table Lookup Components.

6.2 Open Loop Motion of Three Joints

To extend the friction model to joints 2 and 3, the break away experiment of chapter 4 was carried out to build a position correction table; and the constant torque gliding experiment, chapter 5.1, was carried out to measure the viscous and kinetic friction parameters. The positive direction lookup table corrections applied to joints 2 and 3 are shown in figures 6.6 and 6.7. Roughly 60,000 data samples were taken to build each table.

Data to measure friction as a function of velocity were collected at four torque levels for joint 2 and three torque levels for joint 3. At each torque level five measurements were made, allowing the variance to be estimated. These data are shown in figures 6.8 and 6.9.




Figure 6.6 Lookup Table Correction to Torques Applied at Joint Two of the PUMA 560.





Figure 6.7 Lookup Table Correction to Torques Applied at Joint Three of the PUMA 560.



Negative Direction

Viscous 8.53 N-m/rad/sec

Kinetic -11.34 N-m


Positive Direction

Viscous 7.67 N-m/rad/sec

Kinetic 12.77 N-m

Figure 6.8. Friction torque as a function of velocity, Joint 2.

Compensation for position dependencies in the friction data for joints 2 and 3 was not as successful as that for joint 1. Figure 6.10 shows a bang-coast-bang move of joint 3 comparable to those of figure 4.4. The RMS velocity error is 0.054 rad/sec in the motion of figure 6.10, vice 0.008 rad/sec in the motions of joint 1 shown in figure 4.4. One thing confounding the identification of friction at joints 2 and 3 is the gravity loading on these joints. During the break away experiment the gravity load was compensated by a configuration-dependent calculation [Armstrong, Khatib and Burdick 86] that provided a gravity balancing torque added to the experimental torque. The RMS velocity error of 0.054 rad/sec corresponds roughly to an RMS torque error of 0.018 N-m, or 0.14% of the maximum gravity load. The uncertainty



Negative Direction

Viscous 3.02 N-m/rad/sec

Kinetic -5.57 N-m

Positive Direction

Viscous 3.27 N-m/rad/sec

Kinetic 5.93 N-m

Figure 6.9. Friction torque as a function of velocity, Joint 3.

in the gravity compensation is substantially larger than this. I originally believed that the gravity compensation could be detected in the break-away data with extraordinary sensitivity, by comparing the positive and negative rotation break-away data. But other sources of difference between the positive and negative rotation break away friction upset this measurement of gravity loading. For this work the gravity compensation was manually tuned to roughly the 0.1 N-m level by tuning gliding motions for minimum velocity variation: a tedious and imprecise process.




Figure 6.10 Velocity Profile Recorded During an Open-Loop Motion of Joint 3 with Position Dependent Correction Applied.

To conduct three joint motions the model becomes substantially more complex than that required for one joint motion: a more complex inertial model must be used and gravity compensation is required. Thus the model has three parts:

- | | |
|----------------------------------|----------------------------|
| A. Model of rigid body dynamics, | 9 Parameters ; |
| B. Gravity compensation, | 4 Parameters ; |
| C. Friction Compensation | 12 Parameters + 6 Tables . |

The rigid body model used is the simplified model presented in [Armstrong, Khatib and Burdick 86]. The inertial, Coriolis and centrifugal forces are computed; this is feasible because the forces are pre-computed and applied open-loop. The lumped inertial parameters of the model depend upon the simple inertial parameters of each

link, which are affected by sensors mounted on the arm. A difficult identification was anticipated, as the added equipment will affect all of the nine inertial parameters. The identification proved unnecessary; it was sufficient to measure the mass, and locate the center of mass of each of the added instruments. The additional mass was added into the simple inertial constants and the nine lumped inertial constants were recomputed. The required mass and location data are presented in table 6.2.

Table 6.2 Mass and Center of Gravity Data for Items Added to the Arm.

Joint	Mass (Kg)	Position X axis	Position Y axis	Position Z axis	Item Name
6	0.471	0	0	0.0762	Force Wrist
2	0.128	-0.089	-0.159	0	Base of Accel Mount
2	0.426	-0.089	-0.190	0	Remainder of Accel Mount
2	0.247	-0.089	-0.225	0	Rotational Accel
6	0.457	0.0	0	0.140	Zebra Invader
6	0.173	0	0	0.200	Force Fingers
3	0.318	-0.083	-0.222	0	Force Finger Box.
3	0.199	-0.083	-0.222	0	Wrist Electronics
3	0.071	-0.089	-0.4064	0	Wrist Cable
6	0.210	0	0	0.127	Shashank Cube
3	0.247	0	0	-0.0584	Rotational Accel on Joint 3

Because software to compute the parameters was available from the investigations of [Armstrong, Khatib and Burdick 86], this entire effort, from borrowing the postal scale to using the parameters during motion took less than a morning ! Experience carefully tuning the diagonal elements of the inertia matrix during single joint motions of each joint suggests that the pre-computed parameters are accurate to 0.1 N-m^2 , or 2%.

The gravity parameters were tuned manually, as described above, and the friction parameters were those presented with figures 5.1, 6.8 and 6.9. Taken together they form a 25 parameter model that was used to produce the open-loop motions of figure 6.11. The cumulative position error at the end of the three joint, open-loop motion is less than 10%.




Figure 6.11 Position Profile Recorded During an Open-Loop Motion of Joints 1, 2 and 3.

6.3 Friction Compensated Force Control

To push the fidelity of force control to a level far below the magnitude of the static friction of the mechanism an impulsive control system was implemented. Friction at very low velocities works against good control in three ways:

1. When the mechanism is stuck, the force across the friction junction is indeterminate;
2. Because of the Dahl effect, break away is not simple;
3. Because of the Stribeck effect, motion at low velocities is unstable.

When the velocity is across a rubbing junction is non-zero, the friction force takes on a value that is determined by the position, velocity and possibly other factors of the mechanism. But when the velocity is zero, the friction force is whatever is required to prevent motion. The implication for control is that knowledge of the applied forces, especially during the transition from stuck to unstuck, may be unavailable; and this will confound efforts to predict the behavior.

The transition from static to kinetic friction is a complex of displacement and velocity-dependent processes. These processes are not thoroughly studied in this report, though conclusive evidence for their presence on an externally observable scale exists in the data of section 5.4. From the viewpoint of this report the Dahl effect is simply a source of uncertainty, a hidden state. It should be pointed out that in other work, [Walrath 84], a parameter of a Dahl-like model was adaptively identified, enabling the researcher to predictively compensate for friction during a rapid transition through zero velocity. For this report the implication of the Dahl effect is that the exact force required to achieve motion from a stuck point is unknown.

The manipulation of delicate objects in hard contact often occurs at low velocities; in many practical situations the velocities will lie in the range of the Stribeck effect. For the mechanism studied, no feedback control is possible which would stabilize the system in the regime of mixed lubrication, from 0.002 to 0.015 radians per second. This implies that any control that must cross this regime must do so unstably, that is bang through.

The motivation and design of the impulsive controller are heuristic. The motivation is based on the arguments above: uncertainty exists in the level of force required to achieve motion, and steady motion at extremely low velocities is unachievable. The impulsive controller operates with a lookup table of torque sequences. Associated with each torque sequence is a force step, measured prior to the exercise of control. The desired force step is determined calculating the force error - the desired force less the measured force - and scaling by a gain. The stored torque sequence providing the maximal force step smaller than the desired step is selected and played out. Because the impulsive action excites the flexibilities of the mechanism, the force measurements are low pass filtered and a space of time is allowed to pass before the next control action is taken. A block diagram is shown in figure 6.12. A feed-forward term is also included; this is important when the desired force is not small. The impulse table employed by the controller is shown in table 6.3.

The impulsive controller was used to insert a piece of wire wrap wire, #30 copper, 0.25 mm diameter, into a hole in a plate of glass, 0.75 mm diameter. The insertion poses a substantial control challenge because the buckling strength of the wire is only 0.2 Newtons, one 60th of the 12 Newton static friction of the mechanism. The configuration of the apparatus is shown in figure 6.13. Joints 2




Figure 6.12 A Block Diagram of the Impulsive Controller.

Table 6.3 Impulse Table for the Impulsive Controller.

Force Step (Newtons)	Torques in Sequence	Wait Time (sec)	First Torque (N-m)	Second Torque (N-m)	Third Torque (N-m)
0.000	1	0.005	0		
0.020	1	0.12	13		
0.035	2	0.12	13	10	
0.055	2	0.12	13	13	
0.076	2	0.12	14	14	
0.096	3	0.12	15	15	2

and 3 of the PUMA were position controlled by the standard industrial controller; they were used to move the wire along the 'V' slot between the glass slides. Joint 1 was aligned with the force control direction and was controlled by the impulsive controller.




Figure 6.13 Experimental Setup for the Wire Wrap Insertion.

The measured force, control action and measured position during a successful insertion are shown in figures 6.14, 6.15 and 6.16. The control action plotted in figure 6.15 is the force step corresponding to the impulse selected by the controller. With a new wire, the insertion was successful about half of the time. The surface being tracked was canted with respect to the direction of motion, so that travel in the force controlled direction was required to maintain contact, as indicated by the change in position between $t = 0$ and $t = 5$ seconds. The wire was slanted into the direction of motion, and twice caught small features on the glass surface. These two points are indicated by the rise of the force error to positive values, at these points the control system reverses its action. At $t = 13$ seconds insertion occurs and the control system drives the wire through the hole.

The commanded force during this motion was 0.1 Newtons, one 120th of the static friction; the RMS contact force error is 0.038 Newtons, one 316th of the static

friction. The RMS control action is 2.76 Newtons, 73 times greater than the RMS applied force error.




Figure 6.14 Measured Force During Wire Wrap Wire Insertion.

To account for the variation in friction with position, table lookup friction compensation was applied during impulsive control. For this application, a table low pass filtered to 500 cycles per radian was used. The contribution of the table lookup compensator to the motion of figures 6.14, 6.15 and 6.16 is shown in figure 6.17. Compensation for the position dependence is especially important for the




Figure 6.15 Control Action During Wire Wrap Wire Insertion.




Figure 6.16 Measured Position During Wire Wrap Wire Insertion.

smallest force steps: 13; 13, 10; and 13, 13 in table 6.2. A diminishment in these impulses of 1 N-m corresponds to a reduction in the output of a factor of 3; an increase of 1 N-m corresponds to an overshoot of a factor of 2. The regions of rapid reversal near six and eleven seconds correspond to reversals in control action and arm direction evident in figures 6.15 and 6.16. A reversal in arm direction causes a switch from the positive direction table to the negative direction table.




Figure 6.17 Table Lookup Action during the Force-Controlled Insertion.

In this chapter three demonstrations of friction-compensated motion are presented. The demonstration of open-loop single joint motion shows that the friction model can be made quite accurate. The demonstration of three joint motion shows

that the modeling techniques developed for joint one of the PUMA mechanism extend to joints two and three, and points the way toward maximum performance motion of the mechanism. It should be noted that during the motions presented, the inertial forces at no time exceed the static friction; that is to say that if the same motions had been undertaken with torques pre-computed using only an inertial model, no motion whatever would have occurred. A demonstration of force control is also presented, using a special 'impulsive' controller. The control law of this controller specifies very hard, quick actions to overcome friction. This control law is applicable only near zero velocity and suffers several limitations; but in its special application - high fidelity control of forces during low speed motions of a mechanism with substantial static, Dahl and Stribeck friction - it is at least an order of magnitude better than the best previously demonstrated.

Chapter 7

Recommendations to the Engineer

"Therefore always when you wish to know the quantity of the force that is required in order to drag the same weight over beds of different slope, you have to make the experiment and ascertain what amount of force is required to move the weight along a level road, that is to ascertain the nature of its friction."

Leonardo da Vinci (1452-1519),
The Notebooks, F II 106 r

Friction is a performance limiting factor for many manipulation mechanisms. For this reason it will be of interest to the engineer designing mechanisms or control for manipulation. The use to which a friction model is to be put will determine the demands upon its completeness and accuracy. The broad categories of application of friction models are:

- A. To determine the performance limits of an existing mechanism;
- B. To design control for a mechanism with friction;
- C. To design mechanisms for high performance in the presence of friction.

The three categories of application are ordered from least demanding to most; with the third representing largely uncharted waters.

This work points to a number of properties of friction which should be measured and suggests experimental procedures for obtaining the measurements.

The properties to be measured are:

1. Repeatability of the Friction Forces;
2. Magnitude of the Kinetic Friction;
3. Magnitude of the Static Friction ;
4. Critical Velocity of the Stribeck Effect;
5. Linearity of Friction as function of velocity and the Magnitude of the Viscous Friction;
6. Magnitude and Character of Position Dependence in the Static and Kinetic Friction.

7.1 Recommendations on Experimental Technique

With a force-actuated mechanism, it is straightforward to measure properties 2, 4 and 5. A stiff velocity servo should be implemented, as in section 5.2, and constant velocity motions conducted at a range of speeds; the average required torque should be measured. Several samples should be collected with each starting point and velocity in order to estimate the repeatability; such issues as warm up and independence from load should be examined. By reducing the commanded speed in repeated trials of constant velocity motion, the speed at which stick-slip motion begins can be observed. This velocity is not V_c of equation (5.3), but the experiment is straight forward, and will provide an estimate useful in establishing the limits to performance.

All of the analytic work studying limit cycles induced by static friction assumes a measurement of the difference between the static and kinetic friction. This difference is the magnitude of the Stribeck effect, and will be a consequence of the material properties of the drive train components and the lubricant. Because the static friction can vary with position at very high spatial frequency, it is difficult to measure. The break away experiment of section 4 was highly successful, and can be implemented on any apparatus with force-commanded actuation and position sensing. The key is sampling at adequate spatial frequency, and binning according to position. If the sampling is not sufficiently dense, problems of undersampling and aliasing arise. If the data are examined according to the natural distribution rather than according to evenly-spaced bins, the effects of non-uniform friction will bias the estimate of mean static friction.

Position dependence is introduced into the static friction of electric mechanisms with gear transmissions by the non-uniformity of the motor, especially brush type motors, and by the cyclic interaction of gear teeth. The latter effect was dominant in the mechanism studied here. A mechanism that employed a uniform actuator and transmission, such as pneumatic actuation through a four bar linkage, might not exhibit important spatial dependency. DC servo motors and high ratio gear drives can, however, exhibit a 30% change in static friction over a distance small in relation to the smallest feature of the actuator or drive train. Sampling with a spatial frequency 20 times higher than the finest gear pitch or the expected cogging in the motor is recommended.

The break away experiment will provide measures of the static friction and position dependence in the friction, items 3 and 6 above. This completes the list of basic measurements of friction. The performance limits imposed by the Stribeck

WD-A190 732


DYNAMICS FOR ROBOT CONTROL: FRICTION MODELING AND
ENSURING EXCITATION DUR. (U) STANFORD UNIV CA DEPT OF
COMPUTER SCIENCE B S ARMSTRONG MAY 88 STAN-CS-88-1205


273

UNCLASSIFIED

F/G 20/11

NL





effect are only weakly mapped out by this experimental procedure: the minimum stable velocity is measured. Knowledge of the controller bandwidth might be used in conjunction with this minimum velocity to estimate the minimum position or force step that can be taken. Often in practical controllers this estimate is embedded in the form of a deadband in integral control.

7.2 Recommendations on Control

From the perspective of the control engineer friction is divided into two regimes: near zero velocity and away from zero velocity. Away from zero velocity the problem is a straightforward one: the friction should be modeled and the model used in feedback design or for feed-forward compensation. The experience of this research shows friction in the moderate to high velocity regime to be well behaved and simple to model. There is some impetus to neglect friction in this velocity regime because it is benign: it is never destabilizing. This is nonsense, the damping should be understood and under the control of the controller. For control system designs that operate in coordinate spaces other than joint space, such as operational space control or Salisbury stiffness control, it is even more important that friction be compensated: these systems are designed to provide independent motion in the output coordinates rather than the joint coordinates; the unmodeled friction in joint coordinates will translate to an unmodeled coupling of motions in the output coordinates.

Control of mechanisms near zero velocity is another thing altogether. Two factors conspire to challenge the use of a friction model directly in control: uncertainty introduced by the Dahl effect, and the destabilizing character of the Stribeck effect. The Dahl effect has not been thoroughly studied in this effort,

casual examination of its repeatable behavior in trials such as that shown in figure 5.4 suggests that the effect could be captured in a predictive model. The challenge to predicting Dahl friction will arise in ascertaining the state to the required accuracy. The break-away transitions studied here have been carefully constructed for repeatability. In an active control situation, such factors as time at zero velocity and acceleration near zero velocity will be important. The challenge of estimating the state near zero velocity is made greater by the fact that the system is unstable. The instability introduced by the Stribeck effect will all but eliminate hope of accurately estimating the velocity in the traditional fashion. Direct and accurate velocity sensing is strongly indicated.

The impulsive control demonstrated relies on the system coming to rest between impulsive actions, thus skirting the difficult challenge of accurately estimating velocity. The controller takes advantage of information other than that directly available from the sensors: knowledge that the mechanism will be stuck in the friction at the time a control action is taken. This special requirement limits the velocity and bandwidth of the controller, but within these limitations the controller is shown to perform well. I believe that the challenge of determining the state with sufficient accuracy to effectively predict friction limits the feasibility of using control to compensate for low velocity, non-linear friction.

Joint torque sensing is being pursued as a way to directly sense the friction and compensate for it by feedback control [Pfeffer, Khatib and Hake 86]. This approach requires torque or force sensing devices at the output of the transmission or drive train of each link of a mechanism. Accurate joint torque sensing and high bandwidth feedback control will compensate for some of the effect of non-linear friction at zero and low velocities.

7.3 Recommendations on Design

The principal contribution of this work toward mechanism design is to bring to light the Stribeck effect and show it to be a repeatable, tractable and important factor for control and ultimate mechanism performance. That is to say to show that attention to performance must begin with the mechanical design. The best solution to many difficult problems is to avoid them, which, in this case, translates to the design of mechanisms that achieve motion without rubbing. The IBM research micro manipulator is an example of such a device: the three axis micro manipulator is comprised of flexible beams set at right angles and actuated electro-magnetically. Other design alternatives, such as forced lubrication homogeneous drive technology, and more compliant mechanisms may also offer improved performance.

Chapter 8

Conclusions

This work has shown that friction is large, that it is quite repeatable and that, during motion, it can be effectively predicted and compensated. In this particular mechanism, a dc servo motor driven rotary joint with grease lubricated gears and ball bearings, it has been found that:

- Motion friction depends upon position and is about 99% repeatable;
- Break away friction, when measured in a controlled way, is about 97% repeatable;
- Above a minimum velocity, the friction is very linear with velocity;
- The Stribeck effect plays a destabilizing role below a minimum velocity;
- The Dahl Effect plays an important role in the transition from rest to motion
- The friction is different in the positive and negative rotation directions.

Demonstrations of open-loop control and high fidelity force control show the friction modeling to be quite accurate.

Though the direct results of this work apply only to a particular mechanism, our PUMA 560, the methodology and general findings are more widely applicable.

The measurement and modeling of the Stribeck effect, never previously discussed in the controls literature, points the way toward the development of analytic and experimental techniques to study and measure this effect and its impact on control. And the unavailability of friction models has been a hindrance to the effective off-line tuning of control systems, the implementation of optimal control strategies and the accurate simulation of motion. The results presented should prove useful in these applications.

The challenge of the introduction, that is the challenge of force control in hard contact, is not yet met however. As discussed below in the appendix, neither impulsive control nor dither show considerable advantage over simple feedback during force control in hard contact. But study of friction may yet point the way: all of the control schemes tested in force control - linear feedback, impulsive control and control with dither - either neglect or attempt to jump past the break-away and low-velocity friction effects. The key, in hard contact, may be to govern these low velocity effects, especially the Dahl effect, with control that correctly models the forces of small motions.

Appendix A

Small Studies

In this appendix the results of a number of small studies are reported. These small studies investigate friction properties that are not considered in the main thrust of the research but may none-the-less be of interest to some readers. The studies reported here are preliminary examinations made quickly with apparatus that was available.

A.1 Friction as a Function of Motor Angle

The break-away data of figure 4.2 show a position dependence in the static friction. Whether the observed structure in the static friction is connected with motor position can be tested by grouping the break-away data according to motor angle, as shown in figure A.1. To make figure A.1 the break-away data from a full rotation of joint 1, 52 motor revolutions, have been regrouped according to the motor angle. That is to say that the value plotted at each point in figure A.1 is the average of data collected at 52 different arm positions, all corresponding to one motor position. If the break-away friction were uncorrelated with motor angle, any structure in figure A.1 would be coincidental.




Figure A.1 Break-Away Torque as a Function of Motor Rotation Angle, merged data from 52 Motor Revolutions.

Figure A.1 shows twelve peaks in the friction per rotation of the motor. The pinion gear on the motor shaft is twelve pitch, suggesting that the periodic friction in figure A.1 occurs at one cycle per gear tooth.

The drive train of joint 1 of the PUMA 560 robot consists of two intermediate gears acting in parallel, the motor pinion gear and a large bull gear, as shown in figure A.2. If the friction plotted in figure A.1 is taken to be the friction in the motor and the motor/intermediate gear interface; it can be subtracted from the total friction to yield the residual friction. This was done and the residual friction was grouped according to intermediate gear rotation angle. The result, friction as a function of intermediate gear angle, is shown in figure A.3.

The friction signal of figure A.3 shows a component at one cycle per intermediate gear revolution and another at 48 cycles per gear revolution. The one cycle




Figure A.2 Schematic Illustration of the arrangement of gears in joint 1 of the PUMA 560 robot.

per revolution component is due to eccentricity of the intermediate gear. The intermediate gear is 48 pitch at the motor pinion/intermediate gear interface, accounting for the 48 cycle per revolution signal. The mean value of static friction in figure A.3 has been arbitrarily assigned to 2 Newton-meters: the experiment can not distinguish the DC value of static friction associated with each rubbing interface.

The 20% variation in friction (motion friction as well as static, see figure 4.1) occurring with the passage of each motor tooth is a very substantial disturbance to control. Consider that in a standard controller the integral control term will attempt to track the varying friction of figure A.1. The friction disturbance spans the spatial frequency spectrum from one cycle per arm revolution to hundreds of cycles per radian. This disturbance is a strong impetus toward homogeneous drive mechanisms, such as that proposed and examined in [Townsend 88].




Figure A.3 Break-Away Torque as a Function of Intermediate Gear Rotation Angle, merged data from 13 Gear Revolutions.

A.2 Joint 2 Motor Alone and Joint 2 Link Alone

During a maintenance operation the motor of joint 2 was detached from link 2. This opportunity was used to measure the friction in motor 2 alone and in link 2 alone. The break-away experiment was used to measure the static motor friction, as described in chapter 4.1. With the rotational accelerometer attached, link 2 was lifted and allowed to swing under the influence of gravity. The velocity and position were estimated by integrating the acceleration signal and the friction parameters were estimated as described in chapter 5.1.

The break-away friction of the motor alone is compared to the break-away friction of the motor and joint in figure A.4. The mean static friction of the motor alone is 8.0 N-m (reflected to the joint) compared with a mean static friction of 12.7

N-m for the motor with drive train and link. The dominant spatial frequency of the motor with link curve in figure A.4 is 65 cycles per revolution of the motor. This frequency does not correspond to any known drive train feature.




Figure A.4 Static Friction as a Function of Motor 2 Angle. Static Friction Measured with Link 2 Attached and Motor 2 Alone.

While link 2 was swinging freely, the link bearings and the intermediate gear were turning. The acceleration profile recorded during the swinging motion is shown in figure A.5; note the jumps in acceleration that occur when the velocity reverses. The kinetic and viscous friction parameters of the link alone are presented in figure A.6. The link bearings and bull gear/intermediate gear interface contribute 5% of the total kinetic and viscous friction.




Figure A.5 Acceleration and Velocity Profile of Link 2 Swinging under the Influence of Gravity.




Figure A.6 Comparison of the Kinetic and Viscous Friction Parameters for the Assembled Joint 2 and for Link 2 without the Motor.

A.3 Trials with Dither

Dither is a high frequency signal added to the control signal; it is used to reduce the effect of static friction. Dither is commonly used in linear control of hydraulic actuators, where static friction can be very substantial. Dither was applied during force control in compliant contact and provided roughly a factor of three improvement in the fidelity of force control.

The apparatus used to test dither incorporated linear position + derivative control of joint 1. The finger force sensors were used to measure contact force. A single pole lag compensator was used, and a proportional gain of 8 Newton-meters of per Newton of force error. Integrated acceleration was used to generate a rate signal which multiplied a velocity gain of -30 Newton-meters per radian/second. The controller sampled 200 times per second. The control parameters were taken from the COSMOS control system [Khatib and Burdick 86] and correspond to critical damping with an environment stiffness of 2000 Newton-meters per radian. A stiff spring was used to give an environmental stiffness of 2,800 N-m per rad. The apparatus was configured as shown in figure 5.3.

In each trial the robot arm was commanded to track a triangular desired force command. Initial Contact occurred under active force control and was stable. The commanded and actual force from a trial with linear control and 4 Newton-meters of dither is shown in figure A.7. As seen in the figure, the actual force follows the desired force with an offset that is dependent on the derivative of desired force. The force error times the proportional gain is roughly the level of static friction,




Figure A.7 Desired and Actual Force During Active Force Control

which averages 9.4 Newton-meters. Note that the use of integral control is limited in situations with non-linear friction because of the tendency to induce limit cycling.

Two Newton-meter dither was applied at a range of frequencies, the RMS force error at each frequency is shown in figure A.8. With a controller sample rate of 200 Hz, the maximum - and most effective - dither frequency is 100 Hz. Dither was next applied at 100 Hz and a range of amplitudes. The efficacy of dither at each of the amplitudes tried is shown in figure A.9.

Figure A.9 shows dither to be quite effective, even at amplitudes substantially greater than the static friction of the mechanism. The motor currents applied to achieve force control with 12 N-m of dither at 100 Hz are shown in figure A.10. The robot sang audibly with dither of 12 N-m or more. The fidelity of force control with dither and the non-linear impulsive control of section 6.3 are compared in figure A.11. At the higher force rate, dither and the impulsive control give comparable




Figure A.8 RMS Force Error During Force Control with 2 Newton-meter Dither Applied at each Frequency.




Figure A.9 RMS Force Error During Force Control with 100 Hz Dither Applied at each Amplitude.

performance. In a second trial, the desired force was scaled down by an order of magnitude. Here the impulsive controller held the RMS force error to values reduced by a factor of two relative to dither.




Figure A.10 Applied Motor Current during Force Control with 12 N-m of Dither at 100 Hz.

Neither dither nor the impulsive controller provided any improvement in force control fidelity when applied during hard contact. The frequency of the first bending mode goes from 20 rad/sec to 50 rad/sec when going from the spring contact of the trials here to hard contact. It is apparent from the success of applying very large dither that the low pass filter affect of the mechanism and environment compliance is important. Which leaves unsolved the challenge of the introduction: to control hard contact in mechanisms with static friction.




Figure A.11 Comparison of the RMS Force Error during Force Control Trials with Dither and the Impulsive Controller.

A.4 Friction as a Function of Load

With the gliding experiment described in chapter 5.1, friction was measured in joint one under three different load conditions. The results are presented in table A.1. The applied torque was fourteen Newton-meters, table lookup compensation was used and the mean gliding velocity was measured five times under each load condition. The load torque presented is the torsional load on joint 1 due to gravity, it acts around an axis orthogonal to the direction of rotation, and thus does not effect the rotation torque directly: the load is borne by the ball bearings. The equivalent additional friction is a computed quantity equal to the amount of torque required to cause the measured change in velocity given the joint viscous friction of 4.94 N-m per rad/sec.

Orthogonal loading gives a small but preceptible effect. The load of 4 Kg at full arm extension is nearly twice the manufacturer's specified maximum carrying

Table A.1 Measurements of Glide Velocity at Three Arm Loads.

Load Condition	Load Torque (N-m)	Glide Velocity (rad/sec)	Equivalent Additional Friction (N-m)
Arm Upright	2.3	1.029	-
Arm Extended	61	1.014	0.003
Arm Extended, 4 Kg Load	92	0.989	0.200

capacity of five pounds and results in an increase of 1.4% in the measured friction.

We may conclude that bearing load is a minor contribution to friction.

A.5 Creep

In the study of materials, creep is the slow deformation of a material under a load too small to cause rapid failure. Ordinary glass, being a fluid, will creep under its own weight - at a rate of milli-meters per century. In control, the phenomenon of creep is of theoretical interest because it determines whether or not the friction force is discontinuous as velocity goes through zero. A discontinuous friction force may violate the Lipschitz condition, which is a prerequisite for several important results in control theory, such as the small gain theorem.

The arm was allowed to come to rest and torque was applied to joint 1 in gradually increasing levels up to 8.0 N-m. The break-away data indicated that at 0.0022 radians, the position of the arm, break-away should occur at 9.1 N-m. As the torque was applied windup was observed. Over several trials, the mean windup was 0.0013 radian during the transition from 0 to 8.0 N-m of torque. From 8.0 N-m, the torque was increased in 0.1 N-m steps, one step per hour. The motor shaft

encoder was used to detect motion; its resolution is 9964 counts per radian of arm motion. The results are presented in table A.2.

Table A.2 Raw Data Collected during 29 hours of Creep Investigation.

Torque (N-m)	Motion (radians)	Remarks
8.0	0	Torque Applied for 18 hours, 12 minutes.
8.1	0	
8.2	0	
8.3	0	
8.4	0	
8.5	0.0001	Motion occurred 30 minutes into the 1 hour trial.
8.6	0	
8.7	0	Motion at moderate velocity to end of range, i.e. full break-away. Motion began 1 minute into trial.
8.8	1.875	

One shaft encoder count was recorded prior to full break-away: that observed during the hour at 8.5 N-m of torque. Windup was expected, but the occurrence of the count after 30 minutes of torque application suggests creep as the motion process. The absence of motion during the hours at 8.6 and 8.7 N-m shows the creep rate to be less than 1.1×10^{-8} radians per second with the applied torque greater than 97% of the break-away torque. The absence of motion for 18 hours at 8.0 N-m shows the creep rate to be less than 1.5×10^{-9} radians per second with an applied torque equal to 90% of the break-away torque. These results indicate that creep is microscopically small. A stability theory that requires a bounded derivative of friction with respect to velocity is not practically supported by these data.

The effect of rapping was investigated briefly. With an applied torque of 8.2 N-m and the same initial position as above, the base of the robot was rapped lightly with a 1 oz. brass mallet. The strokes were applied in a way that would induce

vibration in the drive train without directly applying torque. The rapping induced a steady motion of 1 shaft encoder count (0.0001 rad.) for each 10 raps. After 60 raps full break-away occurred and .2447 radians of motion. Note that the applied 8.2 N-m is substantially less torque than the expected 9.1 N-m break-away torque. The new stuck position corresponded to a location of high friction. Steady rapping apparently induces creep like behavior. Local vibration sources, such as equipment fans, may have contributed to the motion observed at 8.5 N-m of applied torque.

A.6 Effects that were not Observed

That which is not observed may be as important as that which is. In this section a few phenomena are discussed that were expected but did not make themselves apparent. The first of these is torque dependent friction. Because the normal force across some rubbing interfaces, notably the individual gear teeth, is affected by the applied torque, it was expected that a friction component proportional to motor torque would be in evidence. Such a friction component was long sought but never observed, though friction of this type has been identified in the Stanford/JPL three finger hand [Fearing 87].

Magnetic Cogging is also observed in the Stanford/JPL hand, and was sought in the PUMA 560 experiments. But it was not observed. The experiment most sensitive to magnetic cogging is the measurement of static friction as a function of motor angle, shown in figure A.4. The spatial fourier transform of this data, shown in figure A.12, shows a prominent feature at 31.25 cycles per motor revolution. The motor manufacturer (Magnetic Technologies) specifies a magnetic ripple at 25 cycles per revolution with a magnitude not greater than 4% of the applied torque. The feature at 31.25 cycles per revolution has a peak magnitude of 0.26 N-m per $\sqrt{\text{Hz}}$,

and a square root power between the half maxima equivalent to a torque signal of 0.34 N-m. The magnitude of 0.34 N-m is 4.4% of the average torque of 8.0 N-m applied during the break-away experiment, putting the measurement in the range of the 5% upper bound set by the manufacturer. But the shift in frequency from the expected 25 cycles per revolution to 31.25 is unaccounted for. The fractional value for cycles per revolution may be an artifact of the use of the discrete fourier transform.




Figure A.12 Spatial Fourier Transform of the Static Friction of the Joint 2 Motor while Disconnected from Link 2.

Bibliography

Annotated Bibliography

This annotated bibliography is provided in response to a lack of communication between the specialties concerned with friction and across the decades over which friction work has been done. The need for greater communication is shown by the complete absence of any reference to Bowden and Tabor in the controls literature and the absence of reference to the describing function work of the fifties by roboticists doing similar work today. The author's categorization of each work is shown in italics at the upper right; an asterick indicates the author's recommendation.

Machine Design; Tribology: Theoretical, Experimental

Biel, C. 1920.

"Die Reibung in Gleitlagern bei Zusatz von Velloil zu Mineralöl und bei Veränderung der Unlaufzahl und der Temperatur," *Zeitschrift des Vereines Deutscher Ingenieure* 64(1920)449:83.

Biel was the first to point out that the type of friction curve measured by Stribeck, [02], and described theoretically by Sommerfeld, [04], is likely to describe the general behavior of lubricated surfaces (see chapter 5.5). The work includes experimental study of friction for several bearings as a function of lubricant and temperature, but is most important for its contribution to popularizing the generalized Stribeck curve.

*Tribology: Theoretical, Experimental ***

Bowden, F.P. and Tabor, D. 1956.

Friction and Lubrication, New York: John Wiley and Sons Co.

*Tribology: Theoretical, Experimental ***

Bowden, F.P. and Tabor, D. 1982.

Friction, Malabar: Krieger Pub. Co.

These two short books, 150 and 170 pages respectively, are the *principia* of tribology, the study of rubbing. They are not the same, the former is much more technical and is concerned with the details of friction mechanisms, models and experiments. It provides an excellent view of the state of understanding of tribology circa 1956 and is universally cited in the friction literature. The latter is more discursive, little direct experimental data is presented, and interesting quotations are used to introduce each chapter. The books should probably be read latter then former. For any student of friction these books are required reading.

Adaptive Control: Theoretical, Experimental

Canudas, C., Astrom, K.J. and Braun, K. 1986 (April 7-10).

"Adaptive Friction Compensation in DC Motor Drives," *Proc. of the 1986 Inter. Conf. on Robotics and Automation*, San Fransisco: IEEE.

Canudas, Astrom and Braun provide a consise description of a system which adaptively identifies kinetic and viscous friction parameters. A brief theoretical discussion substantiates the algorithm employed. The careful experimental work and results are described. Though reference is made to several prior experiments in adaptive friction compensation, the relationship of this work to other works is not developed.

Tribology: Survey

Czichos, H. 1978.

Tribology, New York: Elsevier Scientific Pub. Co.

Czichos provides an excellent survey of the field of tribology. The book is very dense: Czichos is not shy about assuming knowledge of chemistry and physics on the parts of his reader. The book can be used as a starting point for many detailed investigations of friction; the references are very complete and liberally scattered throughout the discussion. There are over five hundred references.

*Control: Simulation, Theoretical, Experimental **

Dahl, P.R. 1968 (May).

"A Solid Friction Model," TOR-158(3107-18), The Aerospace Corporation, El Segundo, California.

Control: Theoretical, Simulation

Dahl, P.R. 1976 (December).

"Solid Friction Damping of Mechanical Vibrations," *AIAA Journal*, 14(12)1675:82.

Control: Theoretical, Experimental

Dahl, P.R. 1977.

"Measurement of Solid Friction Parameters of Ball Bearings," *Proc. of 6th Annual Symp on Incremental Motion, Control Systems and Devices*, U. of Illinois.

Dahl is important for bringing experimental data to the study of control of mechanisms with friction. The Dahl model (see chapter 5.4) treats the transition from static to kinetic friction as a function of distance. It is comparable in this way to [Rabinowicz 51]. Dahl cites [Rabinowicz 65], but does not make any use of the reference, rather depending entirely on experimental data acquired by H. Shibata at the Aerospace Corp. The technical report [Dahl 68, 24 pages] provides the clearest exposition of the model and its justification. In [Dahl 76] the model is applied to flexible structures, and in [Dahl 77] additional experimental work with ball bearings is discussed.

The Dahl model has been used in the work of a number of researchers toward control of mechanisms with friction and was found to be important in the mechanism studied here.

Tribology: Historical Perspective

Dowson, D. 1979.

"*History of Tribology*," New York: Longman.

History of Tribology (660 pages) starts with archaeological evidence for the use of bearings and lubrication in ancient times, and proceeds to a detailed account of modern tribology. The book may appeal to two separate audiences, or to two separate faculties in the same reader. The historical accounts are quite descriptive and present a fascinating chapter in the history of technology. As the subject approaches current research, Dowson changes hat from that of narrator to that of scientist, and the work becomes substantially more technical. The citations are

good, and for cases when chronology provides an important research tool, *History of Tribology* will be a useful source of pointers into the literature of tribology.

*Adaptive Control: Experimental **

Gilbert, J.W. and Winston, G.C. 1974.

"Adaptive Compensation for an Optical Tracking Telescope," *Automatica* vol 10, pp. 125:31.

Gilbert and Winston demonstrate a model reference adaptive control scheme that includes a coulomb friction parameter. The implementation, which is discussed at length, was done with an analog computer. Non-linear circuit components were used to compute the coulomb friction signals. A factor of 6 improvement in tracking performance is shown.

Robotics: Simulation

Gogoussis, A. and Donath, M. 1987 (March 31 - April 3).

"Coulomb Friction Joint and Drive Effects in Robot Mechanisms," *Proc. 1987 Inter. Conf. of Robotics and Automation*, Raleigh: IEEE, pp. 883:9.

Gogoussis and Donath are concerned principally with the simulation of motions of mechanisms with non-linear friction. They discuss a friction model which includes negative viscous friction at low velocities, which they modify around zero velocity to make it suitable for simulation. Oscillations introduced by the negative viscous friction are discussed, as are a number of design issues that arise from consideration of friction. Simulation is shown to be possible by incorporating friction forces into the joint constraining forces computed during the recursive Newton-Euler calculation of dynamics.

*Tribology: Theoretical, Experimental **

Rabinowicz, E. 1951.

"The Nature of the Static and Kinetic Coefficients of Friction," *Journal of Applied Physics*, 22(11):1373:79.

*Tribology: Theoretical, Experimental **

Rabinowicz, E. 1958.

"The Intrinsic Variables affecting the Stick-Slip Process," *Proc. Physical Society of London*, vol 71, pp. 668:75.

In "The Nature of the Static and Kinetic Coefficients of Friction" Rabinowicz studies the transition from static to kinetic friction. His experiment is simplicity itself, a block on an inclined plane is struck by a marble. The release height of the marble, declination of the plane and travel of the block are related to basic phenomena of break away. Several materials and lubricants are studied. The integral of the friction curve is explored by increasing the impulse delivered by the marble and recording the increase in traveled distance. "The Intrinsic Variables affecting the Stick-Slip Process" develops a related thought. By studying the curves drawn by slider-surface contact during stick-slip motion, Rabinowicz is able to determine the acceleration of the slider, which is born by a flexible mount. Asymetry between the acceleration and deceleration is used to support a friction model incorporating a characteristic distance. Rabinowicz does not seem aware of Stribeck, the Stribeck effect might also be used to account for the data.

Tribology: Experimental

Rabinowicz, E. 1956.

"Autocorrelation Analysis of the Sliding Process," *Journal of Applied Physics*, 27(2)131:5.

By studying the autocorrelation of friction measured during sliding Rabinowicz explores the size and properties of junctions formed between the sliding surfaces. The character of the interface junctions is important for the Dahl effect, and thus important for control.

Tribology: Descriptive

Rabinowicz, E. 1956 (May).

"Stick and Slip", *Scientific American*, 194(5)109:17.

Rabinowicz is credited with coining the phrase stick-slip. This paper surveys historical and contemporary knowledge of the phenomenon and its role in mechanical motion as well as every day affairs.

*Tribology: Theoretical, Experimental ****Rabinowicz, E. 1965***Friction and Wear of Materials* New York: John Wiley and Sons.

This book is [Bowden and Tabor 56] brought up to date. Coverage of basic friction and wear processes and the affects of lubrication is excellent. The book is somewhat more detailed than [Bowden and Tabor 56], and deserves to be brought out of the remote annex of the library.

*Tribology: Experimental***Stribeck, R. 1902.**

"Die Wesentlichn Eigenschaften der Gleit- und Rollenlager - The Key Qualities of Sliding and Roller Bearings," *Zeitschrift des Vereines Seutscher Ingenieure* 46(1902)1342:48, 1432:37 and 1463.

Stribeck does notably careful measurements of bearing friction across a range of bearing designs and lubricants. His work appears at a time when demands on bearings were increasing rapidly. The paper is important for its identification of negative viscous friction at low velocities. Stribeck's data do not extend to the boundry lubrication regime; the form of what is now called the generalized Stribeck curve came later from a synthesis of the work of Stribeck, Sommerfeld and Biel.

*Tribology: Theoretical***Sommerfeld, A. 1904.**

"Zur Hydrodynamischen Theorie der Schmiermittelreibung," *Zeitschrift Fur Mathematik und Physik*, 50(1904):97-155.

Sommerfeld solves the Reynolds equation for the lubricant flow in journal bearings. Sommerfeld is aware of Stribeck's work, [02], and makes a point of the agreement between his theoretical predictions and the data of Stribeck. The paper is important to applied mathematics for its contribution of an original solution method for a class of partial differential equations; and important to tribology for the contribution of the curve that later became known as the generalized Stribeck curve.

Control: Theoretical **

Tou, Julius. 1953.

"Coulomb and Static Friction in Servo-Mechanisms," *PhD Thesis*, Electrical Engineering Dept., Yale University.

Control: Theoretical

Tou, J. and Schultheiss, P.M. 1953.

"Static and Sliding Friction in Feedback Systems," *J. of Applied Physics*,

24(9)1210:17.

Tou applies describing function analysis to study the closed loop control of mechanisms with static and kinetic friction. He assumes an instantaneous transition from static to kinetic, as shown in figure 1.1. The treatment is quite thorough - very little that is contained in the score or more of subsequent articles applying similar analysis is not to be found in Tou. The journal article is condensed from the thesis; space available to the journal article does not permit presentation of all of the work of the thesis, and the extra effort of obtaining the thesis is well rewarded.

Tou's conclusion, obtained many times since, is that a static friction will cause a controller with an integral control term to limit cycle. For a particular system, the analysis and analog computer simulation show minor loop (velocity) feedback to be stabilizing. Tou provides three references to work in describing function analysis. Like Tustin, he makes no reference to the literature of friction.

Robot Control: Theoretical

Townsend, W. and Salisbury, J.K. 1987 (March 31-April 3).

"The Effect of Coulomb Friction and Sticktion on Force Control," *Proc. 1987 Inter. Conf. of Robotics and Automation*, Raliegh: IEEE, pp. 883-9.

Townsend and Salisbury apply describing function analysis to the study of limit cycling exhibited by systems with static friction and integral control. They choose the static plus kinetic plus viscous friction model of figure 1.2. Townsend and Salisbury are apparently unaware of Tou or others who have applied similar analysis to the same problem. The paper is good however, and makes an interesting contribution by employing the Routh criteria to examine input-dependent stability that can arise when control system parameters are experimentally tuned. As control

system parameters are virtually always experimentally tuned in robot controllers, this is an important contribution.

Control: Theoretical

Tustin, A. 1947

"The effects of Backlash and of Speed-Dependent Friction on the Stability of Closed-Cycle Control Systems," *IEE Journal*, v. 94, part 2A:143-51.

In this fascinating article that anticipates by an hour the introduction of describing function analysis, Tustin explores the implications for stability of an assumed exponential friction model (see chapter 5.2). The paper is important because it explains stick-slip behavior in systems that do not have integral control. The impact of backlash and speed-dependent friction are studied by vector graphic methods, more modern methods may achieve the same result and be easier to apply.

Adaptive Control: Experimental

Walrath, C.D. 1984.

"Adaptive Bearing Friction Compensation Based on Recent Knowledge of Dynamic Friction," *Automatica*, 20(6):717:27.

Walrath presents a thorough piece of experimental work with the object of improving the accuracy of telescope pointing by adaptive friction compensation. He starts with the Dahl model, but finds a higher correlation with velocity than with position, and makes empirically justified modifications to the model. An adaptive compensator is implemented on a digital computer and shown to give a factor of 5 improvement in RMS line of sight tracking error.

Additional References

Armstrong, B., Khatib, O. and Burdick, J. 1986 (April 7-10).

"The Explicit Dynamics Model and Inertial Parameters of the PUMA 560 Arm," *Proc. 1986 Inter. Conf. of Robotics and Automation*, San Francisco: IEEE.

Asada, H. and Youcef-Toumi, K., 1984.

"Analysis and Design of a Direct-Drive Arm with a Five-Bar Parallel Drive Mechanism," *ASME Journal of Dynamic Systems, Measurement and Control*, vol. 106, pp. 225-230.

Bohacek, P.K. and Tuteur, F.B. 1961 (May).

"Stability of Servomechanisms with Friction and Stiction in the Output Element," *IRE Transactions on Automatic Control*, AC-6(61)222:27.

Bowden, F.P., Moore, A.C., and Tabor, D. 1943.

J. of Applied Physics, 14(80).

Bowden, F.P. and Leben, L. 1939.

"The nature of Sliding and the Analysis of Friction," *Proc. of the Royal Society, Series A* vol 169, pp. 371:91.

Bowden, F.P. and Tabor, D. 1939.

"The Area of Contact Between Stationary and Between Moving Surfaces," *Proc. of the Royal Society, Series A* vol 169, pp. 391:413.

Cannon, R. and Schmitz, E. 1984 (Fall).

"Initial Experiments on the End-Point Control of a Flexible One-Link Robot," *International Journal of Robotic Research*, 3(3)62:75.

Craig, J.J. 1986.

"Adaptive Control of Mechanical Manipulators," *PhD Thesis*, Electrical Engineering Dept., Stanford University.

Dokos, S.J. 1946 (June).

"Sliding Friction Under Extreme Pressures," *J. of Applied Mechanics*, vol 13, pp. 148:56.

Dudley, B.R. and Swift, H.W. 1949.

"Frictional Relaxation Oscillations," *Philosophical Magazine*, Vol. 40, Series 7, pp. 849:61.

Fearing, R. 1987.

Private Communication.

Gogoussis, A. and Donath, M. 1987 (March 31 - April 3).

"Coulomb Friction Joint and Drive Effects in Robot Mechanisms," *Proc. of the 1987 Inter. Conf. on Robotics and Automation*, Raleigh: IEEE.

Khatib, O. and Burdick, J. 1986 (April 7-10).

"Motion and Force Control of Robot Manipulators," *Proc. 1986 Inter. Conf. of Robotics and Automation*, San Fransisco: IEEE.

Kubo, T., Anwar, G. and Tomizuka, M. 1986 (April 7-10).

"Application of Nonlinear Friction Compensation to Robot Arm Control," *Proc. 1986 Inter. Conf. of Robotics and Automation*, San Fransisco: IEEE.

Maples, J. 1985(June).

"Force Control of Robotic Manipulators with Structural Flexibility," *PhD Thesis*, Electrical Engineering Dept., Stanford University.

Pfeffer, L., Khatib, O. and Hake, J. 1986 (June).

"Joint Torque Sensory Feedback in the control of a PUMA manipulator," *Proc. 1986 American Control Conference* Seattle.

Rabinowicz, E. and Tabor, D. 1951.

"Metallic transfer between sliding metals: an autoradiographic study," *Proc. of the Royal Soc. of London, A*, vol. 208, 455:75.

Rabinowicz, E. Rightmire, B.G., Tedholm, C.E. and Willians, R.E. 1955.

"The Statistical Nature of Friction," *Trans. of the ASME*, vol 22, October, pp.981-4.

Rice, J.R. and Ruine, A.L. 1983 (June).

"Stability of Steady Frictional Slipping," *J. of Applied Mechanics*, vol 50, pp. 343:49.

Roberts, R.K. 1984.

"The Compliance of End Effector Force Sensors for Robot Manipulator Control," *PhD Thesis*, Electrical Engineering Dept., Purdue University.

Shen, C.N. 1962.

"Synthesis of High Order Nonlinear Control Systems with Ramp Input," *IRE Trans. on Automatic Control*, AC-7(2)22:37.

Shen, C.N. and Wang, H. 1964.

"Nonlinear Compensation of a Second- and Third-Order System with Dry Friction," *IEEE Trans. on Applications and Industry*, 83(71)128:36.

Satyendra, K.N. 1956 (September).

"Describing Functions Representing the Effect of Inertia, Backlash and Coulomb Friction on the Stability of an Automatic Control System," *AIEE Transactions*, vol 75, pt. II, pp. 243:249.

Silberberg, M.Y. 1957 (January).

"A Note on the Describing Function of an Element with Coulomb, Static and Viscous Friction," *AIEE Transactions*, vol 75, pt. II, pp 423:25.

Townsend, William 1988.

"The Effect of Transmission Design on the Performance of Force-Controlled Manipulators," *PhD Thesis*, Mechanical Engineering Dept., Massachusetts Institute of Technology.

Wang, D. and Vidyasagar, M. 1987 (March 31-April 3).

"Control of a Flexible Beam for Optimum Step Response," *Proc. 1987 Inter. Conf. of Robotics and Automation*, Raliegth: IEEE, pp. 1567-72.

Chapter 9

Introduction

9.1 Introduction




Figure 9.1 A General Model Inverse Control Structure.

From the perspective of a controller, a robot arm is a machine to which forces are applied with the result of motion. The relationship between the applied forces and the resulting motion comprises the dynamic model; the applied forces are sometimes known as the forces of motion. Unlike time invariant systems, such as an elevator, when the robot arm moves this relationship changes, in some aspects the relationship between forces and motion may change by as much as 300% in

a fraction of a second. This configuration dependence within the dynamic model introduces a level of complexity into robot control that is absent from the challenge of controlling time invariant systems; complexity which motivates the use of inverse model control schemes such as is shown in figure 9.1. The *forward* dynamic model, which determines motions from applied forces, is represented by $f(q, \dot{q}, \ddot{q})$. The *inverse* dynamic model, $f^{-1}(q, \dot{q}, \ddot{q})$, determines the required torques from desired motions. In principle, if $f^{-1}(\cdot)$ were known exactly it would be possible to achieve error free motion. In practice feedback control structures, such as that shown the figure, are used to correct motion errors. The model inverse control scheme, called computed torque control when applied to robots, has received considerable research attention because it promises to correctly account for the changing forces of motion. Computed torque control is beginning to appear in industrial robots.

To employ computed torque control, the inertial parameters of the robot arm must be known. These are simple physical properties: mass, location of center of gravity and inertia; but a robot arm may have many parameters and they may be difficult to measure directly. For these reasons parameter identification methods are attractive. Parameter identification is the process of experimentally determining fixed[†] characteristics of a system which affect the input-output relationship. In our case the system is a robot arm, the input is applied torques and the output is motion. Parameter identification has appeal because it can be applied to assembled robots, possibly as they perform their intended task.

[†] In the general parameter identification problem the parameters may be slowly varying rather than fixed. It is in this case that the parameter identification algorithms actually have the greatest relevance. But for this analysis of robot dynamic parameter identification, it is presumed that the parameters are fixed.

Uncertain measurements and systematic errors will introduce error into any experiment: Heisenberg has relieved us of any uncertainty in this. In some cases the error may be insubstantial; but in the identification of robot parameters noise can be very considerable. This thesis is focused on the problem of the error in parameter estimates that arises from experimental noise; the relationship between noise and error will be investigated and a means of minimizing the error by careful trajectory selection developed. Three noise sources will be considered: measurement noise (uncertainty in determining what motion occurred); motion noise (uncertainty in determining what torques were applied), and systematic error (errors in modeling or biased behavior of sensors). These noise sources will give rise to two kinds of error in the estimate of the robot parameters: bias (a systematic shift in the measured parameters), and variance (a random shift in the measured parameters that will vary from trial to trial of the experiment). Bias and variance are shown schematically in figure 9.2. The connections between noise sources and parameter estimate error are shown in figure 9.3.

In many situations bias is the most grave estimation problem; thus sensor noise and systematic error are the most troublesome noise sources.

The inverse dynamic model of the robot arm, $f^{-1}(\cdot)$, and the noise models used in this analysis are presented in chapter 10. Using these models and standard parameter identification algorithms, the precise relationships between noise sources and parameter estimate errors are derived in chapter 11. The analysis will show the magnitude of the errors to depend upon the magnitude of the noise and the level of *excitation* in the experiment. The magnitude of the errors will be seen to depend upon the magnitude of the noise and the level of *excitation*, as is shown schematically in figure 9.4.




Figure 9.2 The True Parameter Vector, Bias and Variance in Parameter Space

	Bias	Variance
Sensor Noise (Random)	X	X
Motion Noise (Random)		X
Systematic Error (Unmodeled Dynamics)	X	

Figure 9.3 The Contributions of Three Noise Sources to Bias and Variance.




Figure 9.4 Schematic Relationship between Noise Sources, Excitation and Parameter Error

In many situations the noise power is determined by the sensors and mechanism used and may be affected only at great cost. The dynamic model, $f(\cdot)$, is largely determined by the mechanism design. The engineer is left some freedom in the choice of model complexity. While a change in model complexity will affect the sensitivity of the experiment to noise, such a change does not affect the nature or analysis of the impact of noise upon parameter accuracy. For now we shall assume that a full model identification is desired. The issue of identifying a reduced model is considered in chapter 14.3. Improved excitation is left as the principle means of improving the accuracy of the estimated robot parameters. Good choice of trajectory is the principle means of improving the excitation.

Excitation in a multivariate estimation problem is a subtle issue. The notion of excitation might be approached by imagining standing in a dark space, attempting to explore it with a pencil beamed flashlight, sweeping the flashlight about the space seeing one spot at a time. The robot experiment yields up information about one combination of parameters (one direction in parameter space) at a time, the 'illuminated direction' changing as the robot moves along its trajectory. The

path of our flashlight in the dark (parameter) space corresponds to the sequence of parameter combinations illuminated by the identification experiment. The excitation of an experiment is the amount of light shown in the least illuminated direction. The least illuminated direction may be illuminated better in two ways: we may buy a bigger flashlight, or, if the light is unevenly distributed, we may distribute the light more evenly. This thesis is about distributing the light more evenly. As we shall see, in practical situations some directions in parameter space can be 5,000 times less excited (illuminated) than others, offering considerable room for improvement.

A method is presented in chapter 12 for selecting an experimental trajectory offering maximally even excitation. The method employs the calculus of variations to increase the illumination in the least illuminated direction. The method has been applied to two experiments reported in the literature; the results of this application are presented in chapter 13. And in chapter 14 the implications of poor excitation and the prospects for attaining good excitation in practical situations are discussed.

Two broad issues arise in consideration of manipulator parameter identification: the importance of obtaining the 'true' parameters and the merits of alternative means of parameter measurement. As seen in figure 9.3, bias arises from two of the three noise sources considered. From the standpoint of control, it can be argued that the biased parameter set is in fact the most desirable parameter set: the biased parameters set gives the minimum squared prediction error. This issue is taken up more fully in chapter 14.2. For the purposes of this investigation it is taken for granted that the mass parameters are true, physical quantities and that accurate knowledge of these quantities is desired.

While working through the lengthy development of adaptive control means and their error processes it is possible to lose sight of the fact that we are measuring physical quantities and that alternative means of measurement exist. The analysis of error arising during parameter identification will show that it is difficult to achieve an accurate measurement of the parameters: a few percent of noise can result in tens of percent of error. Direct measurement of the parameters, for all its rigors, and limited identification, in spite of its lesser generality, should be considered and possibly applied in practical situations.

9.2 Prior Work

The literature of parameter identification is very large. In the past few years, since the manipulator dynamic models themselves became available, attention has been directed toward identification of manipulator dynamic parameters. In work toward developing methods for robot dynamic parameter identification, Khosla and Kanade [85] present on-line and off-line methods for identifying inertial parameters. Neuman and Khosla [85] treat a similar problem in detail, and give some consideration to identifying the kinematic parameters as well. Mukerjee and Ballard [85] make the interesting suggestion of instrumenting each joint of a robot arm with multi-axis force and torque sensors. This sophisticated instrumentation would allow measurement of otherwise unobservable inertial parameters. Olsen and Bekey [85, 86] also consider robot inertial parameter identification. In their first paper, [85], they mention the issue that is the substance of this paper: the problem of excitation. In their second paper, [Olsen and Bekey 86], they give brief consideration to the impact of additive noise and propose filtering to limit this impact. Acceleration estimation is one substantial noise source that has been addressed in the theoretical study of manipulator parameter identification; [Hsu

et.al. [87] and [Middleton and Goodwin 88] have proposed identification methods which do not directly employ estimates of acceleration, however both proposed algorithms require that a lead filter be applied to the velocity signal, a step comparable to direct acceleration estimation.

Experimental work is important in demonstrating the utility of parameter identification algorithms. Craig [86; 87] has demonstrated adaptive control, incorporating on-line identification, of the Adept 1 arm; and Atkeson, An and Hollerbach [86] have demonstrated the use of batch least squares estimation to determine the inertial parameters of the MIT/Asada Direct Drive arm. These two reports of experimental work will provide a vehicle to this thesis for the study of noise processes during identification and the connection between excitation and noise immunity. Details of excitation and sensor noise will be taken from these experiments and used to show the importance of the trajectory selection in practical situations.

Slotine and Li [87a, 87b, 86] have also presented experimental results in manipulator parameter identification. Their algorithm combines sliding mode analysis with an output error update. With some over simplification, in linear identification an equation error algorithm employs estimates of the true system state while an output error algorithm employs estimates of the system state derived from the model and inputs. This distinction substantially affects the analysis of noise and error. The analysis presented in this thesis will apply to the equation error formulation: Slotine and Li are unique among those presenting identification methods for manipulator parameters in presenting an output error formulation. An analysis should be carried out of the noise and error processes for the output error

formulation; in particular, such an analysis should address the impact of unmodeled dynamics.

Addressing noise in linear parameter identification, Mareels *et.al.* [87] discuss the connection between excitation and the sensitivity of the parameter estimates to noise. Kalaba and Spingarn [82] discuss the problem of achieving good excitation for non-linear systems. Widrow and Stearns [85] thoroughly discuss the LMS algorithm, which was used by Craig [86], including attention to the variance and convergence properties. Mendel [73] and Ljung [87] address several identification algorithms with attention given to identification in the presence of noise. Finn [74] provides a good treatment of the estimation problem from a statistical point of view. Kendall and Stuart [73] discuss the estimation problem at a more advanced level and address the problem which arises when there is uncertainty in the knowledge of the state variables, as is introduced here by considering sensor noise. Their section on functional and structural models, (*The Advanced Theory of Statistics*, volume 2, pages 391-435) is one of the most readable passages in the advanced statistics literature.

Work toward determining the mass parameters of a manipulator from engineering data is reported in [An, Atkeson and Hollerbach 85] and [Tarn, Bejczy, Han and Yun 85]. Direct experimental measurement of the mass and inertial parameters of a 6 degree of freedom robot arm is reported in [Armstrong, Khatib and Burdick 86].

The original contributions of this thesis lie in the application of noise analysis to the identification of manipulator parameters, with particular attention to the problems that arise in robotics. A measure of experimental sensitivity to unmodeled dynamics is developed and demonstrated, and the bias introduced by sensor noise (output error) is brought to light and shown to be substantial. The level of

excitation provided by practically achievable trajectories is, for the first time, addressed directly through the use of practical examples. Taken together these developments lead to the means to evaluate whether robot parameters may be accurately determined through parameter identification.

Chapter 10

The System and Noise Models

10.1 A Model of Manipulator Rigid Body Dynamics

The rigid body model of manipulator dynamics may be written [Craig 86]:

$${}_n\tau_k = {}_n\phi'_k(q, \dot{q}, \ddot{q}) \theta^* \quad (10.1)$$

where ${}_n\tau_k$ is the torque at joint n at sample time k ;
 ${}_n\phi_k$ is the regressor vector, comprising functions which
constitute the rigid body model;
 ϕ' is the transpose of ϕ .
 $(q, \dot{q}, \ddot{q})$ are the joint position, velocity and acceleration;
 θ^* is the true parameter vector;
 n is the index on the joint number;
 k is the index on the sample time;

Several things are to be noted here: a discrete time model is used, and the notation is that of linear identification and adaptive control, rather than that used elsewhere in robotics. Additionally, ${}_n\tau_k$ and ${}_n\phi_k$ are indexed on the joint number, n , whereas θ^* is not. Many inertial parameters appear in the equations of more than one joint, and data from several joints may be combined to produce an improved estimate of θ^* . This merging of data can only be variance weighted if noise analysis

is undertaken. Finally, some of the functions within ϕ include accelerations which are often difficult to measure. Identification methods have been presented which do not employ measured accelerations, e.g., [Slotine and Li 86], [Hsu et.al. 87] and [Middleton and Goodwin 88]. These methods may avoid the problems attendant with estimating acceleration, but noise and excitation remain issues. In particular, it is shown that a well-conditioned $\mathcal{R}$ is necessary to minimize the effects of unmodeled dynamics.

10.2 The Noise Model

The choice of noise model is, to an extent, up to the investigator. To provide a model sufficient to capture noise processes relevant to robot parameter identification experiments, I will introduce three noise components:

- | | | | |
|---|--|---|--------|
| 1. Sensor Noise
(Random) | $\phi^* = \hat{\phi} + \tilde{\phi},$ | $\tilde{\phi} \in \mathcal{N}(0, C_{\tilde{\phi}});$ | |
| 2. Motion Noise
(Random) | $\tau = \phi^{*'}\theta^* + \tilde{\tau},$ | $\tilde{\tau} \in \mathcal{N}(0, \sigma_{\tilde{\tau}}^2);$ | |
| 3. Systematic Error
(Unmodeled Dynamics) | $\tau = \phi^{*'}\theta^* + \tilde{\tau} + \nu,$ | $\nu \in L_{\infty}.$ | (10.2) |

where ϕ^* is the true regressor vector;
 $\hat{\phi}$ is the estimated regressor vector;
 $\tilde{\phi}$ is noise introduced by sensor noise;
 $\tilde{\tau}$ is a random contribution to the torques at the joint;
 ν is a systematic error contribution;
 $\mathcal{N}(0, \sigma^2)$ is a random variable with zero mean and σ^2 variance;
 L_{∞} is the set of bounded functions.

The regressor vector, $\phi_k'(q, \dot{q}, \ddot{q})$, depends upon the position, velocity and acceleration of the mechanism. Noise in the measurement of these quantities will introduce noise into the calculated value of $\hat{\phi}$. This noise, $\tilde{\phi}$, is taken to be normally distributed and uncorrelated with ϕ^* , $\tilde{\tau}$, or ν . The system model, equation (10.1), provides only for torques predicted by the rigid body model; terms $\tilde{\tau}$ and ν provide random and systematic torque signals. The random torques may arise from driver noise, external disturbances or a random component of friction. The systematic contributions may be physical torques, such as unmodeled friction or flexibility; or may be virtual torques which arise in systematic error in equation (10.1).

The system model, eqn (10.1), can be augmented to include the three noise sources:

$$\tau_k = (\hat{\phi}_k' + \tilde{\phi}_k')\theta^* + \tilde{\tau}_k + \nu_k . \quad (10.3)$$

By evaluating the expected value of $\tilde{\theta}$, the parameter misadjustment, using the augmented model of equation (10.3), the bias and variance in the parameter estimates can be determined. The bias is an offset between the true parameter values and those that will be estimated by the algorithm given infinite data. Note that this definition of bias is somewhat different from the formal definition of statistics (Fisher's definition): here bias is defined with respect to the true parameter values, θ^* . The variance describes a region in which the estimated parameters may be expected to lie. The region's shape is determined by the parameter covariance matrix, $C_{\tilde{\theta}}$, the size is determined by the required confidence.

10.3 A Note On Notation

In this paper many variables have all or several of four degrees of freedom. The degrees of freedom are:

- N the number of joints;
- R the number of basis functions;
- K the number of stages in the identification trajectory;
- M the number of iterations over the identification trajectory during the parameter identification experiment.

For consistency, an upper case letter N, R, K, M will indicate the maximum count in a dimension; and a lower case letter n, r, k, m will serve as the index on the respective dimension. Thus

$${}^r_n\phi_k^m$$

which is a scalar, is the r^{th} regressor vector element, contributing at the n^{th} joint, evaluated at the k^{th} stage of the trajectory on the m^{th} iteration over the trajectory during the identification experiment. R is the number of individual functions in the model. the vectors ϕ and θ are both $(R \times 1)$. The trajectory stage number, k , runs from 0 to K and repeats. This notation arises out of the assumption that a fixed trajectory is repeated M times.

The absolute time (clock time) of a sample is given by:

$$\text{time} = (m * K + k) * \Delta T.$$

where ΔT is the sample time period.

For reference, the definitions of the principal variables used in this paper are provided here.

${}_n\tau_k$	is the torque at joint n and trajectory stage k ;
T	is the concatenated vector of torques: $[\tau_0 \dots \tau_K]'$, it is $(K \times 1)$;
ϕ	is the regressor vector, the basis functions of the model;
ϕ'	the transpose of ϕ ;
Φ	is the concatenated vector of ϕ' , it is $(K \times R)$;
θ	is the parameter vector;
θ^*	is the true parameter vector;
$\hat{\theta} \hat{\phi}$	are the estimated parameter and regressor vecs, see (11.1) and (10.2);
$\tilde{\theta} \tilde{\phi}$	are the parameter misadjustment and regressor noise vectors;
$\nu \tilde{\tau}$	are systematic error and torque noise, see eqn (10.2);
$\mathcal{N}$	is the concatenated vector of systematic error contributions $[\nu_0 \dots \nu_K]$;
μ	is the bias susceptibility, see eqn (11.18);
$\mathcal{R}$	is the input correlation matrix, see eqn (11.10).

Definitions pertinent to the study of the on-line identification algorithm:

Γ	is the gain matrix, chosen by the engineer;
I	is the identity matrix;
Q	is the eigenvector matrix of $\mathcal{R}$;
$\Lambda_{\mathcal{R}} \Lambda_{\Gamma}$	are the eigenvalue matrices of $\mathcal{R}$ and Γ ;
C	is the parameter covariance matrix, see eqn (11.10).

Definitions pertinent to the optimization algorithm:

$B \ C$	are weighting matrices assigning cost to control (acceleration) and jerk;
$J \ H$	are the cost function and the augmented cost function to be minimized;
L_k	is the cost at stage k ;
λ	is the vector of Lagrange multipliers.

Operations:

$\kappa[\cdot]$ is the condition number of matrix $(\cdot)$;

$\underline{\sigma}[\cdot]$ $\bar{\sigma}[\cdot]$ are the smallest and largest singular values of matrix $(\cdot)$;

$E\{\cdot\}$ is the expected value of $(\cdot)$;

Π is the product operator.

Chapter 11

Noise Analysis of the LMS and Least Squares Identification Algorithms

11.1 The Basic LMS Parameter Identification Algorithm

To study the LMS parameter identification algorithm we may represent the parameter misadjustment by writing

$$\theta^* = \hat{\theta} + \tilde{\theta} \quad (11.1)$$

where $\hat{\theta}$ is the estimate of the parameter vector;
 $\tilde{\theta}$ is the parameter misadjustment vector.

The LMS parameter update equation is [Craig 86]:

$$\hat{\theta}_{k+1} = \hat{\theta}_k + \sum_n \Gamma_n \phi_k e_k \quad (11.2)$$

where e_k is the equation error, $e_k = (r_k - \phi_k' \hat{\theta}_k)$;
 Γ_n is a gain matrix, set arbitrarily and usually diagonal.

The Γ_n matrix is subscripted because it is possible that different gains would be applied to data from the several joints, especially if the noise or $\|\phi_k\|$ vary markedly

between joints. For the remainder of this study, summation over the joints is implied, and the subscript n is dropped. Substituting equations (10.1) and (11.1) into (11.2) gives the error transition equation:

$$\tilde{\theta}_{k+1} = (\mathbf{I} - \Gamma \hat{\phi}_k \hat{\phi}_k') \tilde{\theta}_k . \quad (11.3)$$

For the parameter estimate to converge to the true parameters, (that is $\tilde{\theta} \rightarrow 0$), the matrix $[\Gamma \hat{\phi} \hat{\phi}']$ must be positive definite over each sufficiently long portion of the trajectory, leading to the customary persistent excitation condition [Craig 86; Sastry 84]:

$$\begin{aligned} \exists \alpha, \beta > 0 \quad \text{s.t.} \quad \forall l \quad \alpha \mathbf{I} < \Gamma \sum_{k=l}^{l+\rho} \phi_k \phi_k' < \beta \mathbf{I} \\ \rho \geq R \end{aligned}$$

where $\mathbf{I}$ is the identity matrix;

R is the number of basis functions in the model.

The requirement for accurate identification in the presence of noise we will see to be stronger: α must not only be positive, it must be reasonably large; this is the requirement of *sufficient* excitation.

11.2 The Consideration of Noise and Evaluation of Bias

The parameter misadjustment resulting from experimental noise may be studied by augmenting the error transition equation, (11.3), with the system model including noise, (10.3), giving:

$$\tilde{\theta}_{(k+1)} = (\mathbf{I} - \Gamma \hat{\phi}_k \hat{\phi}_k') \tilde{\theta}_k - \Gamma \hat{\phi}_k (\tilde{\phi}_k' \theta^* + \tilde{\tau}_k + \nu) . \quad (11.6)$$

By taking the expected value of each side of (11.6) and using the assumption that the system has reached steady state, that is $(E\{\tilde{\theta}_{(k+1)}\} = E\{\tilde{\theta}_k\})$, the expected bias can be determined. The covariance of $\tilde{\theta}$, i.e., the variance in the parameter measurement, can be determined by taking the expected value of the covariance matrix under the assumption that the covariance has reached steady state. The existence of a steady state and asymptotic convergence of the algorithm to that state have been demonstrated in [Craig 86] and elsewhere.

$$E\{\tilde{\theta}_{(k+1)}\} = E\{(\mathbf{I} - \Gamma \hat{\phi}_k \hat{\phi}_k') \tilde{\theta}_k - \Gamma \hat{\phi}_k (\tilde{\phi}_k' \theta^* + \tilde{\tau}_k + \nu)\} . \quad (11.7)$$

Achieving steady state implies that

$$\begin{aligned} \sum_{k=0}^K E\{\tilde{\theta}_{(k+1)} - \tilde{\theta}_k\} &= 0 = \\ \sum_{k=0}^K E\{-\Gamma \hat{\phi}_k \hat{\phi}_k' \tilde{\theta}_k - \Gamma \hat{\phi}_k (\tilde{\phi}_k' \theta^* + \tilde{\tau}_k + \nu)\} . \end{aligned} \quad (11.8)$$

Using the linearity of the expected value operation to move the summation in under the expected value operation and expanding $\hat{\phi}$ into $(\phi^* - \tilde{\phi})$ gives:

$$0 = -E\left\{\sum_{k=0}^K \Gamma \phi_k^* \phi_k^{*'} \tilde{\theta}_k\right\} - E\left\{\sum_{k=0}^K \Gamma \tilde{\phi}_k \tilde{\phi}_k' \hat{\theta}_k\right\} - E\left\{\sum_{k=0}^K \phi_k^* \nu\right\} \quad (11.9)$$

where many terms have been eliminated using the fact that $\tilde{\phi}$ and $\tilde{\tau}$ are uncorrelated with ϕ^* , θ^* or ν .

$E\{\tilde{\theta}\}$, and therefore $E\{\hat{\theta}\}$, are not stationary in equation (11.9). They move as the experimental trajectory progresses. In order to determine the average bias, which

is independent of the rate gain, Γ , the assumption of slow adaptation is invoked. If Γ is sufficiently small the expected $\tilde{\theta}$ does not depend upon the stage of the trajectory and $\tilde{\theta}_k$ may be replaced by $\tilde{\theta}$. Defining two important matrices, the input correlation matrix and the regressor covariance matrix:

$$\begin{aligned} \text{Input Correlation Matrix } \mathcal{R} &= \left(\frac{1}{K}\right) \sum_{k=0}^K \phi_k^* \phi_k^{*'} ; \\ \text{Regressor Covariance Matrix } \mathcal{C}_{\tilde{\phi}} &= \left(\frac{1}{K}\right) E \left\{ \sum_{k=0}^K \tilde{\phi}_k \tilde{\phi}_k' \right\} . \end{aligned} \quad (11.10)$$

equation (11.9) may be reduced to:

$$0 = -\Gamma \mathcal{R} E\{\tilde{\theta}\} - \Gamma \mathcal{C}_{\tilde{\phi}} E\{\tilde{\theta}\} - \Gamma \sum_{k=0}^K \phi_k^* \nu_k . \quad (11.11)$$

Expanding $\hat{\theta}$ into $(\theta^* - \tilde{\theta})$ and solving for $E\{\tilde{\theta}\}$ yields:

$$E\{\tilde{\theta}\} = (\mathcal{R} + \mathcal{C}_{\tilde{\phi}})^{-1} \mathcal{C}_{\tilde{\phi}} \theta^* - (\mathcal{R} + \mathcal{C}_{\tilde{\phi}})^{-1} \sum_{k=0}^K \phi_k^* \nu_k . \quad (11.12)$$

11.3 Bias Due to Random Sensor Noise

In equation (11.12) we see two sources of bias in the estimated parameters, one due to random noise in the regressor vector, $\mathcal{C}_{\tilde{\phi}}$, and the other due to unmodeled dynamics, ν . Considering only the effect of sensor noise, the bias is given by (see [Mendel 73], chapter 5.3):

$$E\{\tilde{\theta}\} = (\mathcal{R} + \mathcal{C}_{\tilde{\phi}})^{-1} \mathcal{C}_{\tilde{\phi}} \theta^* . \quad (11.13)$$

The contribution of random noise in the regressor to the bias is not often discussed; consideration of the cost function minimized by the LMS algorithm, $E\{e^2\}$, shows how the bias arises:

$$e_k = (\tau_k - \hat{\phi}_k \hat{\theta}_k) = (\phi_k^* \tilde{\theta}_k - \tilde{\phi}_k \hat{\theta}_k) , \quad (11.14)$$

where only parameter misadjustment, $\tilde{\theta}$, and regressor noise, $\tilde{\phi}$, have been considered. The LMS algorithm will seek to choose the value of $\hat{\theta}$ that will minimize the magnitude of $E\{e^2\}$. The $\phi^* \tilde{\theta}$ term is minimized by making $(\hat{\theta} \rightarrow \theta^*)$. However $\tilde{\phi} \hat{\theta}$ is minimized by making $(\hat{\theta} \rightarrow 0)$!! The error signal due to $\tilde{\phi}$ bucks the error signal due to $\tilde{\theta}$ and tends to cause the LMS algorithm to underestimate the parameters.

11.4 Bias Due to Unmodeled Dynamics and Bias Susceptibility

Considering only the effect of systematic error, the bias is given by:

$$E\{\tilde{\theta}\} = -[\mathcal{R}]^{-1} \left(\frac{1}{K} \right) \sum_{k=0}^K \phi^* \nu_k . \quad (11.15)$$

The identification algorithm attempts to fit the available model parameters to the unmodeled dynamics. The outcome will depend upon the correlation of the model basis functions and the unmodeled dynamics, and upon the conditioning of the $\mathcal{R}$ matrix.

An upper bound on the expected parameter misadjustment vector may be determined from equation (11.15); this bound, the *bias sensitivity*, is given by:

$$\frac{\text{Max} \left\{ \frac{\|E\{\tilde{\theta}\}\|}{\|\mathcal{N}\|} \right\}}{\|\mathcal{N}\|} = \frac{1}{\sqrt{\sigma[\mathcal{R}]}} = \text{bias sensitivity} \quad (11.16)$$

where $\mathcal{N}$ is the concatenated vector of ν , i.e., the elements ν_k stacked one atop another: $\mathcal{N} = [\nu_1 \cdots \nu_K]'$.

Note that $\|\mathcal{N}\|$ equals the root mean squared average of ν , the L_2 norm is used.

Bias sensitivity has the units of *experiment sensitivity*, which is the ratio of the size of the true parameter vector to the size of the applied torque, or output per input of the identification experiment:

$$\frac{\|\theta^*\|}{\|\mathbf{T}\|} = \text{experiment sensitivity} \quad (11.17)$$

where $\mathbf{T}$ is the concatenated torque vector, $\mathbf{T}$ is $(K \times 1)$.

The ratio of bias and experiment sensitivity is a dimensionless quantity, bias susceptibility, that will be a coarse but useful indicator of the robustness of the experiment to systematic error.

$$\mu = \frac{\text{bias sensitivity}}{\text{experiment sensitivity}} = \frac{(1 / \sqrt{\sigma[\mathcal{R}]})}{(\|\theta^*\| / \|\mathbf{T}\|)} \quad (11.18)$$

where μ is the Bias Susceptibility.

The bias susceptibility is the maximum amplification of error due to unmodeled dynamics possible in a particular experiment. As a measure, bias susceptibility will be most useful if the ratio of modeled and unmodeled dynamics is approximately known, in which case a bound on parameter bias may be determined by

$$\|E\{\tilde{\theta}\}\| \leq \mu \|\theta^*\| \frac{\|\mathcal{N}\|}{\|\mathbf{T}\|} . \quad (11.19)$$

Bounds determined by matrix norm operations are often quite conservative. Equation (11.19) will over estimate parameter bias when $\mathcal{N}$ has very little contribution in the direction of weakest identification, as will be the case when $\mathcal{N}$ is the consequence of higher order bending modes separated in frequency from the modes represented in ϕ , the model. But manipulators often exhibit unmodeled dynamics of other sorts, such as friction, actuator non-linearities or ripple, and imperfections in the kinematic model underlying the dynamic model. All of these contribute torques, $\mathcal{N}$, which will be interdependent with the elements of the model, ϕ . Consider that the structure of the ARMA model can represent the true dynamics of any linear system of the correct order, but the structure of a manipulator dynamic model can represent the dynamics of only one design of manipulator. An example in chapter 13.2 shows that $\|E\{\tilde{\theta}\}\|$ can approach the bound of (11.19) in a plausible circumstance.

11.5 The Consideration of Noise and Evaluation of Variance and Convergence Rate

Analytically evaluating the convergence rate of the sequential algorithm when the adaptation is fast is a solvable problem, [Armstrong 87], but leads to a complicated matrix product that offers little insight. Furthermore, the parameter covariance matrix, $\mathcal{C}_{\tilde{\theta}}$, is non-stationary in this case, and thus difficult to analyze. To permit analysis that will offer insight into the convergence and covariance behavior of the sequential identification algorithm, stationarity can be artificially introduced by choosing Γ small enough to give slow adaptation, that is, $\bar{\sigma}(K\Gamma\mathcal{R}) \ll 1$. This assumption, that Γ is small enough to give slow adaptation is invalid in many

experimental situations. For example, in the Adept 1 arm identification experiment, [Craig 86], Γ tuned to give a steady state random parameter misadjustment of $\pm 5\%$ gives $\bar{\sigma}(K\Gamma\mathcal{R}) \approx 10$. Reducing Γ to give $\bar{\sigma}(K\Gamma\mathcal{R}) \approx 0.1$ would result in a prohibitively long experiment, roughly 25 hours for 95% convergence. None-the-less there is value in the limited analysis: it will give useful insight and provide the analytic handle needed for experiment optimization. Once a trajectory has been chosen using the slow adaptation analysis, it is possible to verify that the true behavior will closely follow the predicted behavior by evaluating the exact convergence rate. In most cases the true behavior will closely follow the predicted behavior because the performance of the identification algorithm is dominated by the smallest singular value of the gain scaled excitation matrix, which will virtually always be much less than one. The assumption of slow adaptation is applied widely, often implicitly, in the study of adaptive control and is a requirement for averaging analysis (see [Anderson et.al. 86]).

By repeated application of the error transition equation, (11.23), we may determine the change in the parameter misadjustment that will occur with each iteration over the entire experimental trajectory:

$$\tilde{\theta}_0^{m+1} = \prod_{k=0}^K (\mathbf{I} - \Gamma \hat{\phi}_k \hat{\phi}_k') \tilde{\theta}_0^m \quad (11.20)$$

where m is the index on trajectory iteration.

The polynomial expansion of (11.20) may be formed and, if the adaptation rate is sufficiently slow, all terms of Γ^2 and higher order may be eliminated, leaving

$$\tilde{\theta}_0^{m+1} = (\mathbf{I} - \Gamma \sum_{k=0}^K \hat{\phi}_k \hat{\phi}_k') \tilde{\theta}_0^m = (\mathbf{I} - K\Gamma\mathcal{R}) \tilde{\theta}_0^m \quad (11.21)$$

If $K\Gamma\mathcal{R}$ is positive definite and not too large, (11.21) will converge. Using the assumption of slow adaptation, (11.21) may be approximated by a continuous time system which will decay toward zero error with exponential time constants given by, [Widrow and Stearns 85]:

$$\tau_r = \frac{\Delta T}{\|\mathbf{r}\lambda\|_2} \quad (11.22)$$

where τ_r is the r^{th} exponential error decay time constant;
 $\|\mathbf{r}\lambda\|_2$ is the magnitude of the r^{th} eigenvalue of $\Gamma\mathcal{R}$;
 ΔT is the sample period.

Note that because the $\mathbf{r}\lambda$ are the eigenvalues of $\Gamma\mathcal{R}$, they are not necessarily real.

Augmenting (11.21) with the noise model of equation (10.3), as was done in equation (11.6), and again retaining only zeroth or first order terms in Γ gives

$$\tilde{\theta}_0^{m+1} = (\mathbf{I} - \Gamma \sum_{k=0}^K \hat{\phi}_k \hat{\phi}_k') \tilde{\theta}_0^m - \Gamma \sum_{k=0}^K \hat{\phi}_k (\tilde{\phi}_k' \theta^* + \tilde{\tau}_k + \nu_k) \quad (11.24)$$

The covariance of the parameter misadjustment vector may be found by taking the expected value of the outer product of $\tilde{\theta}$. We are concerned here with only the random portion of the parameter misadjustment, so we will remove the parameter bias from $\tilde{\theta}$ by redefining $\tilde{\theta}$ according to:

$$\tilde{\theta} = \theta^* - \hat{\theta} - \{(\mathcal{R} + C_{\tilde{\phi}})^{-1} C_{\tilde{\phi}} \theta^* - (\mathcal{R} + C_{\tilde{\phi}})^{-1} \sum_{k=0}^K \phi_k^* \nu_k\}. \quad (11.25)$$

The hypothesis of slow adaptation is used to treat $\tilde{\theta}$ as a function of m only. The error transition equation becomes:

$$\begin{aligned} \tilde{\theta}_{(k+1)} = \tilde{\theta}_k - \Gamma(\mathcal{R} + C_{\tilde{\theta}})\tilde{\theta} + \Gamma A(\tilde{\theta} + B) - \Gamma \phi^* \tilde{\phi}' \theta^* \\ + \Gamma \tilde{\phi} \nu - \Gamma \phi^* \tilde{\tau}; \end{aligned} \quad (11.26)$$

where for compactness two lumped terms are used:

$$A = (\tilde{\phi} \phi^{*'} + \phi^* \tilde{\phi}')$$

$$B = \text{the bias: } \{(\mathcal{R} + C_{\tilde{\theta}})^{-1} C_{\tilde{\theta}} \theta^* - (\mathcal{R} + C_{\tilde{\theta}})^{-1} \sum_{k=0}^K \phi_k^* \nu_k\}.$$

By taking the expected value of the matrix outer product of equation (11.26) the covariance transition equation is formed (transposes are eliminated from symmetric matrices and terms in $\tilde{\tau}$ to the first power are omitted because $E\{\tilde{\tau}\} = 0$).

Terms Governing the Covariance Transition Behavior:

$$\begin{aligned} C_{\tilde{\theta}(m+1)} &= C_{\tilde{\theta}m} - C_{\tilde{\theta}m}(\mathcal{R} + C_{\tilde{\phi}})\Gamma - \Gamma(\mathcal{R} + C_{\tilde{\phi}})C_{\tilde{\theta}m} \\ &\quad + \Gamma(\mathcal{R} + C_{\tilde{\phi}})C_{\tilde{\theta}m}(\mathcal{R} + C_{\tilde{\phi}})\Gamma \end{aligned}$$

Terms That will go to zero Under the Expected Value Operation:

($\tilde{\phi}$ to the first power)

$$\begin{aligned} &+ \tilde{\theta}(B + \tilde{\theta})'A'\Gamma - \tilde{\theta}\theta^{*'}\tilde{\phi}\phi^{*'}\Gamma + \tilde{\theta}\nu\tilde{\phi}'\Gamma \\ &+ \Gamma A(\tilde{\theta} + B)\tilde{\theta}' - \Gamma\phi^*\tilde{\phi}'\theta^*\tilde{\theta}' \end{aligned}$$

Terms That will go to zero Under the Expected Value Operation:

($\tilde{\phi}$ to the third power)

$$\begin{aligned} &- \Gamma(\mathcal{R} + C_{\tilde{\phi}})\tilde{\theta}(B + \tilde{\theta})'A'\Gamma + \Gamma(\mathcal{R} + C_{\tilde{\phi}})\tilde{\theta}\theta^{*'}\tilde{\phi}\phi^{*'}\Gamma \\ &- \Gamma(\mathcal{R} + C_{\tilde{\phi}})\tilde{\theta}\nu\tilde{\phi}'\Gamma - \Gamma A(\tilde{\theta} + B)\tilde{\theta}'(\mathcal{R} + C_{\tilde{\phi}})\Gamma \\ &+ \Gamma\phi^*\tilde{\phi}'\theta^*\tilde{\theta}'(\mathcal{R} + C_{\tilde{\phi}})\Gamma + \Gamma\tilde{\phi}\nu\tilde{\theta}' - \Gamma\tilde{\phi}\nu\tilde{\theta}'(\mathcal{R} + C_{\tilde{\phi}})\Gamma \end{aligned} \tag{11.27}$$

Noise Terms:

$$\begin{aligned} &+ \Gamma A(\tilde{\theta} + B)(B + \tilde{\theta})'A'\Gamma - \Gamma A(\tilde{\theta} + B)\theta^{*'}\tilde{\phi}\phi^{*'}\Gamma \\ &+ \Gamma A(\tilde{\theta} + B)\nu\tilde{\phi}'\Gamma - \Gamma\phi^*\tilde{\phi}'\theta^*(B + \tilde{\theta})'A'\Gamma \\ &+ \Gamma\phi^*\tilde{\phi}'\theta^*\theta^{*'}\tilde{\phi}\phi^{*'}\Gamma - \Gamma\phi^*\tilde{\phi}'\theta^*\nu\tilde{\phi}'\Gamma + \Gamma\tilde{\phi}\nu(B + \tilde{\theta})'A'\Gamma \\ &- \Gamma\tilde{\phi}\nu\theta^{*'}\tilde{\phi}\phi^{*'}\Gamma + \Gamma\tilde{\phi}\nu\nu\tilde{\phi}'\Gamma + \Gamma\tilde{\phi}\tilde{\tau}^2\tilde{\phi}'\Gamma \end{aligned}$$

Equation (11.27) is not as daunting as it seems, by using the assumption that the covariance has reached steady state, eliminating the terms that go to zero under the expected value operation, and combining the noise terms equation (11.27) reduces to the matrix Lyapunov equation:

$$0 = -C_{\tilde{\theta}}(\mathcal{R} + C_{\tilde{\phi}})\Gamma - \Gamma(\mathcal{R} + C_{\tilde{\phi}})C_{\tilde{\theta}} + \Gamma(\mathcal{R} + C_{\tilde{\phi}})C_{\tilde{\theta}}(\mathcal{R} + C_{\tilde{\phi}})\Gamma + \Gamma\Sigma\Sigma'\Gamma \quad (11.28)$$

$$\text{where } \Sigma = \{\phi^* \tilde{\phi}' \theta^* + \tilde{\phi} \nu + \tilde{\phi} \tilde{\tau} - (\tilde{\phi} \phi^{*'} + \phi^* \tilde{\phi}')(\tilde{\theta} + (\mathcal{R} + C_{\tilde{\phi}})^{-1}(C_{\tilde{\phi}} \theta^* - \sum_{k=0}^K \phi_k^* \nu_k))\}.$$

The parameter covariance matrix, $C_{\tilde{\theta}}$, satisfying (11.28) may be found using the Kronecker product. Equation (11.28) can be simplified in a way that provides useful physical insight if two assumptions are made: $\mathcal{R} \gg C_{\tilde{\phi}}$ and $\tau \gg \nu$. These assumptions will be valid when the noise sources are small in relation to the signal power of the identification experiment. In the practical examples studied in chapter 13 the first assumption is found to be invalid; but the analysis to follow will provide a coarse but insightful measure of the steady state covariance of the LMS algorithm. Equation (11.28) should be evaluated when an exact measure is required. When the two assumptions hold, equation (11.28) reduces to :

$$0 = -\Gamma \mathcal{R} C_{\tilde{\theta}} - C_{\tilde{\theta}} \mathcal{R}' \Gamma' + K \Gamma \mathcal{R} C_{\tilde{\theta}} \mathcal{R}' \Gamma + v^2 \Gamma \mathcal{R} \Gamma' \quad (11.29)$$

$$\text{where } v^2 = E\{(\tilde{\phi}' \theta^*)^2\} + E\{\tilde{\tau}^2\}.$$

This matrix Lyapunov equation may be solved explicitly when the experimental trajectory (and thus $\mathcal{R}$) is known. If Γ is chosen which commutes with $\mathcal{R}$, it is possible to find an analytic solution to equation (11.29). Borrowing from [Widrow and Stearns 85], we may rotate the coordinate system of the parameter space from the natural coordinates into the normal coordinates of the gain scaled excitation

matrix. Relying on the fact that $\mathcal{R}$ must be positive definite and symmetric, we may define Q and $\Lambda_{\mathcal{R}}$ such that

$$\mathcal{R} = Q \Lambda_{\mathcal{R}} Q' \quad (11.30)$$

where Q is the orthonormal eigenvector matrix of $\mathcal{R}$;
 $\Lambda_{\mathcal{R}}$ is a diagonal matrix made up of the eigenvalues of $\mathcal{R}$.

Because Γ has been chosen to commute with $\mathcal{R}$,

$$Q' \Gamma Q = \Lambda_{\Gamma} \quad (11.31)$$

where Λ_{Γ} is a diagonal matrix of the eigenvalues of Γ .

Note that the elements of Λ_{Γ} are ordered according to the order of eigenvectors in Q , not according to size. In the light of equations (11.30) and (11.31), we see that $\Gamma \mathcal{R}$ may be written

$$\Gamma \mathcal{R} = Q \Lambda_{\Gamma} Q' Q \Lambda_{\mathcal{R}} Q' = \gamma Q \Lambda Q' \quad (11.32)$$

where $\gamma \Lambda$ is a diagonal matrix of the eigenvalues of $\Gamma \mathcal{R}$;
 γ is a scalar chosen so that Λ_{11} equals one.

Now defining C_Q to be the covariance matrix in rotated coordinates,

$$C_Q = Q' C_{\tilde{\theta}} Q ,$$

equation (11.29) may be reduced to:

$$0 = -\gamma \Lambda C_Q - \gamma C_Q \Lambda + \gamma^2 \Lambda C_Q \Lambda + \gamma^2 v^2 \Lambda Q' \Gamma' Q. \quad (11.33)$$

Because Λ is diagonal it is straightforward to expand the matrix multiplications of equation (11.33), giving a scalar equation for each of the elements of C_Q :

$$C_{Qij} = \frac{\gamma v^2 \Lambda_{jj}}{(\Lambda_{ii} + \Lambda_{jj}) - \gamma \Lambda_{ii} \Lambda_{jj}} (Q' \Gamma Q)_{ij} = \frac{\gamma v^2 \Lambda_{jj}}{(\Lambda_{ii} + \Lambda_{jj}) - \gamma \Lambda_{ii} \Lambda_{jj}} \Lambda_{\Gamma ij} \quad (11.34)$$

where Λ_{ii} is the i^{th} eigenvalue of $\Gamma \mathcal{R}$;
 v^2 is the average noise power;
 Q is the eigenvector matrix of $\mathcal{R}$ which serves as the rotation from normal to natural coordinates;
 Γ is the gain matrix, chosen by the engineer;
 γ is a scaling parameter, determined by the engineer's choice of Γ .

Λ_{Γ} is, of course, diagonal, thus C_Q is diagonal.

There are two special choices of Γ which commute with $\mathcal{R}$: the identity matrix and $\mathcal{R}^{-1}$. In linear identification these two choices correspond to the Least Mean Squared Error (LMS) and Recursive Least Squares (RLS) algorithms respectively. Owing to slow adaptation $\gamma \Lambda_{ii} \ll 2$. C_{θ} may therefore be approximated by:

$$\begin{aligned} \Gamma = \mathbf{I} & : C_{\theta} \simeq \frac{\gamma v^2}{2} ; \\ \Gamma = \mathcal{R}^{-1} & : C_{\theta} \simeq \frac{\gamma v^2}{2} \mathcal{R}^{-1} . \end{aligned} \quad (11.35)$$

In a sense these are extreme choices of Γ along an axis from extreme disparity of convergence rates and uniform parameter variance to uniform convergence rates and extreme disparity of parameter variance, as shown in table 11.1. The direction of weakest identification is that of the eigenvector of $\mathcal{R}$ corresponding to the smallest eigenvalue of $\mathcal{R}$.

Table 11.1 Convergence Rate and Error Resulting from Special Choices of Γ .

Choice of Γ	Convergence	Variance, $C_{\hat{\theta}}$	Systematic Error
$\gamma \mathbf{I}$ (LMS)	Most Weakly Identified Direction $\kappa[\mathcal{R}]$ Slower.	Uniform in All Directions.	Systematic Error is not Affected by the Choice of Γ .
$\gamma \mathcal{R}^{-1}$ (RLS)	Uniform in All Directions.	Most Weakly Identified Direction $\kappa[\mathcal{R}]$ Noisier.	

Note: $\kappa[\mathcal{R}]$ is the condition number of $\mathcal{R}$.

11.6 Making the Choice of Γ

It has been suggested that special choices of Γ will result in an estimator performance substantially better than that predicted above. One never gets something for nothing. If the slow convergence in the weakly identified direction is compensated with a large gain in Γ , an increase in the noise sensitivity of the experiment results. For the case of slow adaptation, and Γ chosen to commute with $\mathcal{R}$, the ratio of convergence rate to RMS parameter deviation is fixed by $\underline{\sigma}[\mathcal{R}]$.

If adaptation in the parameter space direction of weakest identification is decoupled from adaptation in the directions of strong identification, equation (11.35) may be extended to rapid adaptation. That is, $\underline{\sigma}[K\Gamma\mathcal{R}] \ll 1$, rather than $\bar{\sigma}[K\Gamma\mathcal{R}] \ll 1$, is sufficient for the analysis above to apply.

In general fast adaptation will introduce coupling: the directions and rates of adaptation rotate and change as the manipulator progresses along the experimental trajectory. In the two experimental situations studied this coupling was small.

11.7 Noise Analysis of the Batch Least Squares Identification Algorithm

Starting with equation (11.1), it is possible to identify the parameters, θ^* , by solving the weighted normal equations [Atkeson, An and Hollerbach 86]:

$$\begin{aligned} \text{The Model:} \quad \mathbf{T} &= \Phi^* \theta^* ; \\ \text{The Parameter Estimate:} \quad \hat{\theta} &= [\hat{\Phi}' W \hat{\Phi}]^{-1} \hat{\Phi}' W \mathbf{T} . \end{aligned} \tag{11.36}$$

where Φ^* is the true concatenated regressor vector, Φ^* is $(K \times R)$;
 $\hat{\Phi}$ is the estimate of the concatenated regressor vector;
 $\mathbf{T}$ is the concatenated torque vector, $\mathbf{T}$ is $(K \times 1)$;
 W is a weighting matrix that may reflect varying degrees of confidence in the elements of $\hat{\Phi}$ and $\mathbf{T}$.

Two standard results of least squares estimation theory, also called multiple regression, are [Finn 74]:

1. Introducing systematic error (unmodeled dynamics) introduces bias:

$$\text{if} \quad \mathbf{T} = \Phi \theta^* + \mathcal{N} , \tag{11.37}$$

where $\mathcal{N}$ is the vector of torques produced by systematic errors;

$$\text{then} \quad E\{\tilde{\theta}\} = -[\Phi' W \Phi]^{-1} \Phi' W \mathcal{N} . \tag{11.38}$$

When $W = I$, equation (11.38) is exactly equation (11.15), which gives the bias of the LMS algorithm.

2. If the experiment is corrupted by additive noise of power σ_r^2 , that is

$$\tau_k = \phi_k' \theta^* + \tilde{\tau}_k, \quad E\{\tilde{\tau}^2\} = \sigma_r^2, \quad E\{\tilde{\tau}\} = 0, \quad (11.39)$$

the parameter misadjustment covariance, $C_{\tilde{\theta}} = E\{\tilde{\theta}\tilde{\theta}'\}$ is given by:

$$C_{\tilde{\theta}} = \sigma_r^2 [\Phi' W \Phi]^{-1}. \quad (11.40)$$

When $W = I$

$$C_{\tilde{\theta}} = \sigma_r^2 [\mathcal{R}]^{-1}. \quad (11.41)$$

The analysis above employs a purely additive noise model; (11.39) does not provide for noise in the regressor vector, Φ , as does the noise model of equation (10.3). As the regressor depends on state, it reflects sensor noise. A complete description of the effect of regressor noise is substantially more complex (see [Kendall and Stuart 73]), evaluating the exact variance is beyond the scope of this thesis. It is important however, to see the origins of bias introduced by sensor noise. For this reason the expected value of the parameter misadjustment will be evaluated with a noise model which includes only sensor noise will be considered.

Introducing sensor noise into the regressor vector:

$$\Phi^* = \hat{\Phi} + \tilde{\Phi} \quad (11.42)$$

where Φ^* is the true concatenated regressor vector;

$\hat{\Phi}$ is the estimated concatenated regressor vector;

$\tilde{\Phi}$ is the noise in $\hat{\Phi}$ regressor vector, introduced by sensor noise.

Using the model of equation 11.36 and expanding Φ^* according to equation (11.42) gives:

$$T = \hat{\Phi} \theta^* + \tilde{\Phi} \theta^* \quad (11.43)$$

The normal equation solves:

$$\hat{\Phi}' W \hat{\Phi} \hat{\theta} = \hat{\Phi}' W T \quad (11.44)$$

Multiplying equation (11.43) on the left by $\hat{\Phi}'$ and combining with (11.44) gives

$$\hat{\Phi}' W \hat{\Phi} \hat{\theta} = \hat{\Phi}' W T = \hat{\Phi}' W \hat{\Phi} \theta^* + \hat{\Phi}' W \tilde{\Phi} \theta^* \quad (11.45)$$

Expanding $\hat{\Phi}$ and applying the expected value operation gives:

$$\begin{aligned} E\{(\Phi^{*'} W \Phi^* - \Phi^{*'} W \tilde{\Phi} - \tilde{\Phi}' W \Phi^* + \tilde{\Phi}' W \tilde{\Phi}) \hat{\theta}\} = \\ E\{\hat{\Phi}' W \hat{\Phi} \theta^*\} + E\{\Phi^{*'} W \tilde{\Phi} \theta^*\} - E\{\tilde{\Phi}' W \tilde{\Phi} \theta^*\} \end{aligned} \quad (11.46)$$

The expected value on the left is difficult to evaluate for finite data lengths because $\hat{\theta}$ is dependent upon, and thus correlated with, $\tilde{\Phi}$ in a complicated way. But $[\Phi^* \Phi^{*'}]^{-1}$ multiplies each contribution of $\tilde{\Phi}$ to $\hat{\theta}$; thus, in the limit of infinite data:

$$\lim_{K \rightarrow \infty} [\Phi^* \Phi^{*'}]^{-1} \rightarrow 0; \quad E\{\tilde{\Phi}' \tilde{\Phi} \hat{\theta}\} \rightarrow E\{\tilde{\Phi}' \tilde{\Phi}\} E\{\hat{\theta}\} \quad (11.47)$$

Equation (11.47) is a assertion of the fact that as the data set becomes long, the effect of any *single* noise sample upon the parameter estimate vanishes. Applying this to equation (11.46) gives:

$$\begin{aligned} (E\{\Phi^{*'} W \Phi^*\} + E\{\tilde{\Phi}' W \tilde{\Phi}\}) E\{\hat{\theta}\} = \\ (E\{\Phi^{*'} W \Phi^*\} + E\{\tilde{\Phi}' W \tilde{\Phi}\}) \theta^* - E\{\tilde{\Phi}' W \tilde{\Phi}\} \theta^* \end{aligned} \quad (11.48)$$

Incorporating the weighting matrix, W , into the definitions of $\mathcal{R}$ and $C_{\tilde{\Phi}}$ for the analysis of the batch least squares algorithm, that is $\mathcal{R} = E\{\Phi^{*'} W \Phi^*\}$ and $C_{\tilde{\Phi}} = E\{\tilde{\Phi}' W \tilde{\Phi}\}$; gives:

$$\begin{aligned} E\{\hat{\theta}\} &= (\mathcal{R} + C_{\tilde{\Phi}})^{-1} (\mathcal{R} + C_{\tilde{\Phi}}) \theta^* - (\mathcal{R} + C_{\tilde{\Phi}})^{-1} C_{\tilde{\Phi}} \theta^* \\ &= \theta^* - (\mathcal{R} + C_{\tilde{\Phi}})^{-1} C_{\tilde{\Phi}} \theta^* ; \end{aligned} \quad (11.49)$$

or

$$E\{\tilde{\theta}\} = (\mathcal{R} + C_{\tilde{\Phi}})^{-1} C_{\tilde{\Phi}} \theta^* .$$

The bias of the batch least squares algorithm, shown in equation (11.49), is exactly that of sequential LMS algorithm, shown in equation (11.13). This is encouraging: the two algorithms will converge to the same result.

11.8 Summary of Noise Results

The method of this chapter has been to study the expected value of the error transition equation using a system model augmented with experimental noise. The noise sources considered were:

*Sensor Noise,
Motion Noise,
Systematic Error.*

In making the analysis it was necessary to employ several assumptions:

For the LMS Algorithm:

*Slow Adaptation,
Convergence to Steady State,
Independence of the sensor and motion noise from other signals.*

For the Batch Least Squares Algorithm:

Infinite Data Length; (comparable to Slow Adaptation).

The parameter estimation errors which result from the three noise contributions are summarized in figure 11.1. In each square of figure 11.1 we see the inverse of the input correlation matrix, $\mathcal{R}^{-1}$. If this matrix is poorly conditioned moderate noise contributions may have a large effect on the estimated parameters. In the next chapter a method is described to improve the condition number of $\mathcal{R}$. In the subsequent chapter the results of applying the improvement method to two practical experiments are presented

	Bias	Variance
Sensor Noise (Random)	$E\{\tilde{\theta}\} = (\mathcal{R} + C_{\tilde{\theta}})^{-1} C_{\tilde{\theta}} \theta^*$	$C_{\tilde{\theta}} \simeq \frac{\gamma \sigma_{\epsilon}^2}{2} \mathcal{R}^{-1}$
Motion Noise (Random)		$C_{\tilde{\theta}} \simeq \frac{\gamma \sigma_{\epsilon}^2}{2} \mathcal{R}^{-1}$
Systematic Error (Unmodeled Dynamics)	$E\{\tilde{\theta}\} = -[\mathcal{R}]^{-1} (\frac{1}{K}) \sum_{k=0}^K \phi^* \nu_k$	

Figure 11.1 The Contributions of Three Noise Sources to Bias and Variance.

Chapter 12

Generation of Very Exciting Identification Trajectories

To maximize the convergence rate or noise immunity of an identification experiment, it is necessary to maximize the minimum singular value of the input correlation matrix, $\underline{\sigma}[\mathcal{R}]$. This is a non-linear path optimization problem that may be solved with the calculus of variations. The cost function is most naturally stated as a terminal cost: $(1/\underline{\sigma}[\mathcal{R}])$. But singular value cannot be evaluated knowing only the terminal manipulator state; it is necessary to know each of the $\mathcal{R}$ matrix elements. To evaluate $\underline{\sigma}[\mathcal{R}]$ as a function of terminal state, it would be necessary to augment the state vector with the elements of $\mathcal{R}$, or, because $\mathcal{R}$ is symmetric $(R * (R + 1)/2)$ states. This would change the MIT/Asada arm problem from one of 6 states to one of 126 states. Alternatively, the cost may be cast as a path cost summed along the trajectory:

$$\begin{aligned} J &= \mathcal{F}[\mathcal{R}] \\ \mathcal{F}[\mathcal{R}] &= \sum_k L_k(x, u) \end{aligned} \tag{12.1}$$

where J is the cost function to be minimized;
 $\mathcal{F}[\mathcal{R}]$ is the cost to be minimized, either $(1/\underline{\sigma}[\mathcal{R}])$ or $\kappa[\mathcal{R}]$;
 L_k is the cost as a function of state and control evaluated at stage k .

To Determine L in closed form is unnecessary: it is sufficient to determine the derivative of L such that

$$\frac{dJ}{d\phi} = \sum_k \frac{dL_k}{d\phi} \quad (12.2)$$

Exploiting the fact that $\mathcal{R}$ is formed by taking a sum along the trajectory, it is straightforward to find dL . The derivative of $\mathcal{F}[\mathcal{R}]$ can be formally written:

$$dJ = \sum_i \sum_j \frac{\partial \mathcal{F}[\mathcal{R}]}{\partial \mathcal{R}_{ij}} \sum_k \sum_r \frac{\partial \mathcal{R}_{ij}}{\partial^r \phi_k} d\phi \quad (12.3)$$

where $\mathcal{R}_{ij}$ is the $(ij)^{th}$ element of the $\mathcal{R}$ matrix;
 $\frac{\partial \mathcal{F}[\mathcal{R}]}{\partial \mathcal{R}_{ij}}$ is the partial derivative of the cost function with respect to
the $(ij)^{th}$ element of the $\mathcal{R}$ matrix.

The partial derivatives of the $\mathcal{R}$ matrix elements are given by:

$$\frac{\partial \mathcal{R}_{ij}}{\partial^r \phi_k} = \begin{cases} 0 & : r \neq i, j; \\ (\frac{1}{K})^j \phi_k & : r = i; \\ (\frac{1}{K})^i \phi_k & : r = j; \\ (\frac{2}{K})^i \phi_k & : r = i = j. \end{cases}$$

Bringing the summation along the trajectory of equation (12.3) through the summations over i and j gives dL_k which satisfies (12.2) :

$$\begin{aligned} dL_k &= \sum_i \sum_j \frac{\partial \mathcal{F}[\mathcal{R}]}{\partial \mathcal{R}_{ij}} \sum_r \frac{\partial \mathcal{R}_{ij}}{\partial^r \phi_k} d\phi \\ J &= \int \frac{\sum_k dL_k}{d\phi} d\phi + \sum_k L_k^0 \end{aligned} \quad (12.4)$$

where L_k^0 is a constant of integration.

Adjoining the Lagrange multipliers and setting their values to remove the state derivative term from dJ gives the Hamiltonian [see Bryson and Ho 75, chap. 1]:

$$H_k = L[X_k, \bar{q}_k] + \lambda'_{k+1} f(X_k, \bar{q}_k)$$

where f is the discrete time system dynamics: $X_{(k+1)} = f(X_k, \bar{q}_k)$.

The trajectory is specified as a sequence of accelerations, $\bar{q}$. The problem of enforcing the specified trajectory, or mapping from $\bar{q}_{desired}$ to τ_k is non-trivial, but is not addressed here. The issue at hand is trajectory selection, and it is presumed that the experimenter has some means to bring about the selected trajectory, $(q, \dot{q}, \ddot{q})_{desired}$. The Hamiltonian is minimized by gradient search in the space of trajectories, the operational equations are [see Bryson and Ho 75, chap. 2]:

$$\lambda'_k = \frac{\partial L(X, \bar{q})}{\partial X_k} + \lambda'_{k+1} \frac{\partial f(X, \bar{q})}{\partial X_k} \quad (12.5)$$

$$\lambda_K = 0$$

$$\bar{q}_k^{m+1} = \bar{q}_k^m - \mu_{cv} \left(\frac{\partial L}{\partial \bar{q}_k} + \lambda_k \frac{\partial f}{\partial \bar{q}_k} \right)' \quad (12.6)$$

An iteration toward reducing J is taken by first sweeping the existing trajectory, $\bar{q}^m$, forward, calculating the $\mathcal{R}$ matrix, calculating $\partial \mathcal{F}[\mathcal{R}] / \partial \mathcal{R}_{ij}$ for each element of the $\mathcal{R}$ matrix and sweeping the λ backward using equations (12.5). The next trajectory is found by the control update equation, (12.6).




Figure 12.1 Four Steps of the Optimization Procedure.

The optimization process is shown schematically in figure 12.1. The four step process is repeated until the trajectory ceases to improve, which is expected to occur at a local optimum.

12.1 Quadratic Cost Applied to Acceleration and Jerk

The optimization procedure of equation (12.5) will minimize $\mathcal{F}[\mathcal{R}]$ without regard to acceleration limits or the undesirability of rapid accelerations. To incorporate these factors into the optimization, cost is assigned to acceleration and jerk; in each case a quadratic cost. Adding quadratic cost is approximate: in the case of acceleration, the limit may come from actuator torque limits which are reflected in a complex way in hard, state-dependent acceleration limits; in the case of jerk, a complex relationship exists between exciting unmodeled dynamics and the bias introduced into the identification. Evaluating the true constraints would be quite complex, and for the moment there is more to be gained by addressing the issue of gross trajectory selection.

The quadratic cost functions may be introduced by writing:

$$L_k = \int \left(\frac{dL_k}{d\phi} \right) d\phi + L_k^o + \ddot{q}' B \ddot{q} + \ddot{q}' C \ddot{q} \quad (12.7)$$

where B and C are positive definite matrices applying a cost to acceleration and jerk;

$\ddot{q}$ is the jerk, the third derivative of position.

The partial derivatives of equation (12.5) are evaluated as before. To set the cost of acceleration and jerk relative to the cost of poor excitation, 'Bryson's rules' may be used. We may set B and C according to:

$$\frac{\|B\|}{(\text{Acceleration Limit})^2} \approx \frac{d\mathcal{F}[\mathcal{R}]/d\ddot{q}}{(\text{Tolerable } \mathcal{F}[\mathcal{R}])^2}$$

If the relationship between jerk and excitation is known or assumed, the cost of jerk may be set by balancing the increase in bias due to jerk, with the decrease in bias due to greater excitation:

$$\frac{\|C\|}{(\text{Tolerable Jerk})^2} \approx \frac{d\mathcal{F}[\mathcal{R}]/d\ddot{q}}{(\text{Tolerable } \mathcal{F}[\mathcal{R}])^2}$$

Determining the tolerable $\mathcal{F}[\mathcal{R}]$ or jerk is a matter of judgment. Furthermore, the derivative of the $\mathcal{F}[\mathcal{R}]$ with respect to the trajectory specification is difficult to determine and not constant. When $\mathcal{F}[\mathcal{R}] = (1/\underline{\sigma}[\mathcal{R}])$, the relative costs of excitation, acceleration and jerk must be balanced, otherwise increasing $\underline{\sigma}[\mathcal{R}]$ will tend toward greater accelerations indefinitely. In this case the algorithm will press actively against the costs of acceleration and jerk, and the final trajectory will be strongly influenced by the relative weights of the several costs. Minimizing $\kappa[\mathcal{R}]$ is technically much easier because the condition number is a dimensionless quantity,

and its use reduces the need to carefully balance the cost of $\mathcal{F}[\mathcal{R}]$ with the costs of acceleration and jerk. For a poorly conditioned matrix, the derivative of the condition number with respect to the smallest singular value is much greater than the derivative with respect to the largest, and the two cost functions are equally effective.

12.2 Implementation Issues

The optimization was implemented with the starting point and trajectory length fixed, the path was constrained as described above, and the end point was unconstrained. These choices are not essential to the result. It would be interesting to allow the trajectory length to vary, its selection based on a measure of *information per time*.

Beyond selecting the cost function, the art of this procedure lies in the selection of the update gain, μ_{cv} . The system that implements equation (12.5) incorporates an automatic gain adjustment function. The matrix $\partial\mathcal{F}[\mathcal{R}]/\partial\mathcal{R}_{ij}$ is evaluated numerically. dJ is then evaluated in two ways: directly, by evaluating $J^{m+1} - J^m$; according to the gradient with equation (12.4). If the two measures of dJ are nearly equal, the performance surface is taken to be fairly smooth and μ_{cv} is increased. If the two are not close together, a smaller step is taken. The gain adjustment depends in essence on a measure of the second derivative of the performance surface, and causes the algorithm to behave more like Newton's method than simple gradient search. In the cases studied, the performance surface was quite convoluted, with long flat regions and sharp drops in multiple, irregular combinations. Prior to implementation of the automatic gain adjustment, optimization was impractically slow.

There is a subtle implementation issue in the selection of units. Because some of the elements of θ have the dimension of inertia and others the dimension of moment (see [Atkeson, An and Hollerbach 86]) it is possible to change the relative scale of the elements of θ and thus of ϕ by changing the choice of units. A change in the relative scale of the elements of ϕ will result in a change in the condition number of the input correlation matrix, $\kappa[\mathcal{R}]$. The experimental error is *not* affected by a change in $\kappa[\mathcal{R}]$ that is due to a change in units, because there is a corresponding change in the relative accuracy of the elements of ϕ . But the optimization process does not incorporate any information about the relative accuracy of the elements of ϕ and would be affected by a shifted $\kappa[\mathcal{R}]$. This issue was investigated for the MIT/Asada arm and the Adept 1 arm experiments. It was found that MKS units are suitable. Because $\kappa[\mathcal{R}]$ is in each case quite large due to true poor excitation, there is a broad range of choice of units that will have little detrimental impact on $\kappa[\mathcal{R}]$. In order for $\mathcal{R}$ to be well conditioned, the distances which multiply the moments should be nearly unity, angstroms or light years would be poor choices of units for current manipulators.

Optimization was attempted using trajectories that could be represented as end points of fifth order polynomial splines rather than the full arbitrary trajectory representation. The effort was not successful, and I do not believe that the very exciting trajectories can be adequately represented in a parametric way. A trajectory of the MIT/Asada arm experiment has 900 degrees of freedom: 3 joints and 300 stages. A trajectory described by two end points, a via point and two time parameters has 11 degrees of freedom. We may imagine a 900 dimensional trajectory space with the space of parametric trajectories existing as an 11 dimensional manifold within the trajectory space. In order for the gradient

search to find a good parametric trajectory, the good parametric trajectory must exist and a continuous path must exist on the 11 dimensional manifold of parametric trajectories from the starting trajectory to the good trajectory. Experience with trajectory optimization suggests that good trajectories are rare, rather locally good, and that the performance surface is very highly convoluted. All three of these factors argue against discovering very exciting trajectories using parametric trajectory representations. In fact, each of the three trajectories used by Atkeson, An and Hollerbach [86] was found to be near a local optimum.

12.3 Local Minima, the Global Minimum and Path Constraints

All that has been said about optimization comes with the caveat that only local optima may be found. In general, a different final trajectory will result from a change in the starting trajectory. No practical method is known for finding the globally optimal trajectory or even determining how good that trajectory might be. Taking as starting points the three trajectories of [Atkeson, An and Hollerbach 87], the optimizer produced markedly different results: one optimization yielded a condition number of 1,000, another 50. Engineering judgment, patience and good luck are required to find a good trajectory.

Experience suggests that hard path constraints introduce local minima into the performance surface and thereby pose a substantial problem for the algorithm. During optimization, trajectories used on the Adept 1 arm tended to migrate outside of the mechanism's joint limits. Constraint enforcement was implemented by adjoining additional Lagrange multipliers [Bryson and Ho 75, chap. 2], but this limited the improvement achieved: the optimizer would get stuck in local minima near the starting trajectory. To achieve better improvement, it was necessary to

relax the path constraints and repair the trajectory manually. I believe that a substantially better trajectory for the Adept 1 arm experiment could be found if a larger space of starting trajectories were examined or the joint limits enforced in a better way.

The quadratic cost added to acceleration and jerk in equation (12.7) did a satisfactory job of keeping the required accelerations within bounds. These additional cost components did not seem to complicate the problem of local minima, possibly because they were smooth cost functions, rather than sharp path constraints.

Chapter 13

Results

13.1 Comparison of Original and Very Exciting Trajectories

The trajectory optimization procedure was applied to the identification experiments reported in [Atkeson, An and Hollerbach 86] describing identification done on the MIT/Asada direct drive arm, and [Craig 86] describing adaptive control applied to the Adept 1 arm. The $\kappa[\mathcal{R}]$ was optimized with the procedure of section 12; the results are summarized in tables 13.1 and 13.2. The rms jerk, $\dot{\ddot{q}}$, is included in these tables as a measure of the degree to which unmodeled flexibilities will be excited. Table 13.1 shows that $\underline{\sigma}[\mathcal{R}]$ for the MIT/Asada arm identification is improved by a factor of 34. An, Atkeson and Hollerbach choose a minimum jerk trajectory, so it is not surprising that the jerk is increased by the optimizer. Table 13.2 reports the results of applying the optimizer to the Adept 1 arm experiment. The $\underline{\sigma}[\mathcal{R}]$ is improved by a factor of 27. The trajectory jerk is substantially reduced in this case because Craig employed the bang-coast-bang trajectory generator of the Adept 1 control system.

Table 13.1. Measures of the Original and Optimal Trajectories, MIT/Asada Arm Identification.

	$\kappa[\mathcal{R}]$	$\underline{\sigma}[\mathcal{R}]$	RMS Jerk
MIT/Asada Arm Identification Trajectory 2	5,710	0.15	0.026
Optimized Trajectory	48.7	5.2	0.041

Table 13.2 Measures of the Original and Optimal Trajectories, Adept 1 Arm Identification.

	$\kappa[\mathcal{R}]$	$\underline{\sigma}[\mathcal{R}]$	RMS Jerk
Adept 1 Arm Identification Trajectory	4,576	0.0074	1.17
Optimized Trajectory	76	0.20	0.14

To assess the impact of noise, quantization noise affecting the regressor estimate, $\hat{\phi}$, are presented in tables 13.3 and 13.4. For the MIT/Asada arm experiment, tachometers were the primary sensors, with acceleration estimated off-line by symmetric first difference. A high order low pass filter with a cutoff frequency of 30 Hertz was applied to the acceleration data, reducing the quantization noise power. For the Adept 1 arm experiment, shaft encoders were the primary sensors; velocity and acceleration were estimated by first and second difference.

Table 13.3. Sensor and Estimate Noise, Assuming the use of First Difference. ΔT is the sample period.

	Shaft Encoder	Tachometer
precision, q	0.010 - 0.0003 rad	0.04 - 0.005 full scale
σ_{pos}^2	$(q^2/12)$	-
σ_{vel}^2	$(2/\Delta T^2) (q^2/12)$	$(q^2/12)$
σ_{accel}^2	$(4/\Delta T^4) (q^2/12)$	$(2/\Delta T^2) (q^2/12)$

Table 13.4. Standard Deviations of State Measurements in the Adept 1 Arm and MIT/Asada Arm Identification Experiments.

	Adept 1 Arm Identification	MIT/Asada Arm Identification
$\Delta T(\text{seconds})$	0.004	0.005
σ_{pos} (radians)	1.6×10^{-5}	-
σ_{vel} (rad/sec)	0.0057	0.0036
σ_{accel} (rad/sec ²)	2.0	0.54

The increase in excitation achieved by trajectory optimization is shown in figure 13.1, $\sigma[\mathcal{R}]$ is shown for the original and optimized trajectories.




Figure 13.1 Level of Excitation, $\sigma[\mathcal{R}]$, of the original and optimized trajectories in the experiments studied.

If the LMS algorithm is employed to estimate the parameters and the gain is set to give a uniform 5% parameter estimate covariance, the convergence rate is determined by the noise and excitation according to equation 11.22; the 95% convergence times for the MIT/Asada arm and Adept 1 arm identification experiments are shown in figure 13.2. Because of the lower noise power of the sensors, the choice of gain for the MIT/Asada arm experiment was actually stability limited (see [Widrow and Stearns 85] for a discussion of stability limitation).




Figure 13.2 95% Convergence time for the LMS algorithm applied to the two experiments studied, the gain is set to give 5% relative variance.

Convergence time is not an aspect of the batch least squares algorithm, which was employed in the MIT/Asada arm identification, the equivalent issue (linked to convergence time by the choice of gain in the LMS algorithm) is parameter covariance. In figure 13.3 the relative parameter deviation resulting from application of the batch algorithm to the data of one trajectory iteration is shown; the deviation

is computed with equation (11.41). The relative parameter deviation is a scalar percentage measure of the uncertainty in the parameter estimate in relation to the magnitude of the parameters themselves; it is given by:

$$\text{Relative Deviation} = \frac{\sqrt{\|C_{\phi}^{-1}\|}}{\|\theta^*\|} \quad (13.1)$$




Figure 13.3 Relative Deviation that would result from applying the batch least squares estimator to the two experiments studied.

Equation (11.13) gives the bias introduced into the parameter estimate by sensor noise. Using the sensor noise models presented in tables 13.3 and 13.4, the expected bias in the two experiments studied has been computed, and is presented in figure 13.4. The bias is presented as a percentage of the magnitude of the parameter vector:

$$\text{Relative Bias} = \frac{\|E\{\tilde{\theta}\}\|}{\|\hat{\theta}\|} \quad (13.2)$$




Figure 13.4 Expected Relative Bias introduced by Sensor Noise.

Notice that for the Adept 1 arm experiment the bias due to sensor noise goes up with the use of the optimized trajectory; this is the one case in which an optimized trajectory does not do as well as that chosen by the original investigator. The bias due to sensor noise is given by equation (11.13): $E\{\tilde{\theta}\} = (\mathcal{R} + C_{\tilde{\phi}})^{-1} C_{\tilde{\phi}} \theta^*$. Employing matrix norms to bound the bias shows:

	$\ (\mathcal{R} + C_{\tilde{\phi}})^{-1}\ $	$\ C_{\tilde{\phi}}\ $	Bound	Actual Value
Original	$\ 186\ $	$\ 4.0\ $	744	≥ 0.72
Optimal	$\ 5.3\ $	$\ 4.0\ $	21.2	≥ 1.29

The bound comes down by a factor of 35, yet the value itself increases by 80%; showing that matrix norm bounds must be treated with some care. In this case the disparity arises because the noise power, $C_{\tilde{\phi}}$, is dominated by acceleration noise, and the excitation problem arises in differentiating kinetic and viscous friction, a velocity dependent distinction; thus both $C_{\tilde{\phi}}$ and $\mathcal{R}$ are very highly skewed, and their skewdness lines up in a special way to give this effect.

Equation (11.15) gives the bias susceptibility, μ . Bias susceptibility is the ratio of the potential sensitivity of the identification experiment to unmodeled dynamics to the actual sensitivity to modeled dynamics. Because μ is always greater than one and is increased by poor excitation, a small contribution of unmodeled torques may give a large bias. For the experiments studied it is plotted in figure 13.5.




Figure 13.5 Bias Susceptibility in the MIT/Asada arm identification and Adept 1 arm identification Experiments.

Views from above of the original and optimal trajectories for the MIT/Asada arm identification experiment are shown in figure 13.6. In figures 13.7, 13.8 and 13.9, the acceleration curves for joints 1, 2 and 3 are presented. In figures 13.10 and 13.11 views from above of the original and optimal trajectories for the Adept 1 arm identification experiment are shown. And in figures 13.12 and 13.13 expanded portions of the acceleration curves for joints 1 and 2 of the Adept 1 arm identification experiment are presented. And figures 13.14 and 13.15 show the position curves for the original and optimal trajectories of the Adept 1 arm identification experiment.

The chief thing which may be said about the plots of original and optimized trajectories is that it is not at all obvious from looking at them that the optimized trajectories should be substantially better, much less a factor of ten to thirty better. Figures 13.14 and 13.15 in particular show that an apparently superficial change in the trajectory may give a tremendous change in the level of excitation.




Figure 13.6 MIT/Asada Arm Identification Experiment: Original and Optimal Trajectories in Cartesian Coordinates, Viewed from Above.




Figure 13.7 MIT/Asada Arm Identification Experiment: Original and Optimal Trajectories, Acceleration of Joint One.




Figure 13.8 MIT/Asada Arm Identification Experiment: Original and Optimal Trajectories: Acceleration of Joint Two.




Figure 13.9 MIT/Asada Arm Identification Experiment: Original and Optimal Trajectories, Acceleration of Joint Three.




Figure 13.10 Adept 1 Arm Identification Experiment: Original Trajectory in Cartesian Coordinates, Viewed from Above.




Figure 13.11 Adept 1 Arm Identification Experiment: Optimal Trajectory in Cartesian Coordinates, Viewed from Above.




Figure 13.12 Adept 1 Arm Identification Experiment: Original and Optimal Trajectories, an Expanded Portion of the Acceleration of Joint One.




Figure 13.13 Adept 1 Arm Identification Experiment: Original and Optimal Trajectories, an Expanded Portion of the Acceleration of Joint Two.




Figure 13.14 Adept 1 Arm Identification Experiment: Original and Optimal Trajectories, Position of Joint One.




Figure 13.15 Adept 1 Arm Identification Experiment: Original and Optimal Trajectories, Position of Joint Two.

13.2 Unmodeled Dynamics: A Case Study

Very little can be said generally about the effect of unmodeled dynamics upon the accuracy or utility of identified parameters. Equations (4.5) and (6.3) allow a bound on parameter bias to be ascertained in terms of bounds on the excitation and disturbance. But the actual effect of a systematic disturbance may be much less than this bound. To help give a feel for how sensitive to unmodeled dynamics an identification experiment may be, the parameter bias introduced by tachometer output ripple in the MIT/Asada arm identification experiment was evaluated. A model for this systematic disturbance and knowledge of the experimental trajectory will allow determination of the $\mathbf{T}$ and $\mathcal{N}$ vectors and the Φ and $\mathcal{R}$ matrices. All of these evaluations depend upon the particular choice of trajectory.

Electro-mechanical tachometers suffer from a multiplicative output ripple that is akin to magnetic cogging in dc motors. The tachometers used in the MIT/Asada arm identification experiment are low ripple devices: the ripple is reported to be less than 1%, [Asada, Youcef-Toumi 84]. It is straightforward to determine the error introduced into the measurement of velocity and estimate of acceleration by this ripple:

$$\begin{aligned}
\hat{\dot{q}} &= \dot{q} * (1 + mag * \sin(poles * q + phase)) \\
\tilde{\dot{q}} &= \dot{q} * mag * \sin(poles * q + phase) \\
\tilde{\ddot{q}} &= \ddot{q} * mag * \sin(poles * q + phase) \\
&\quad + \dot{q}^2 * mag * poles * \cos(poles * q + phase)
\end{aligned}
\tag{13.3}$$

where $\hat{\dot{q}}$ is the measured velocity;
 $\tilde{\dot{q}}$ is the velocity error introduced by the tachometer ripple;
 $\tilde{\ddot{q}}$ is the acceleration estimate error introduced by the tach ripple;
 $q, \dot{q}, \ddot{q}$ are the true joint angle, velocity and acceleration;
 mag is the magnitude of the tachometer ripple;
 $poles$ is the number of poles in the tachometer;
 $phase$ is the phase angle between the tachometer poles
and zero angle of the joint.

This error will corrupt ϕ , and we may calculate $\mathcal{N}$ according to:

$$\mathcal{N} = \tilde{\phi}' \theta^* \tag{13.4}$$

where $\tilde{\phi}$ is now defined by $\tilde{\phi} = \phi(q, \dot{q}, \ddot{q}) - \phi(q, \hat{\dot{q}}, \hat{\ddot{q}})$.

With Φ calculated in the normal way, we may evaluate (see eqn. (11.38))

$$E\{\tilde{\theta}\} = -[\Phi' W \Phi]^{-1} \Phi' W \mathcal{N} \tag{13.5}$$

Tachometer ripple is one possible source of systematic error. The bias that would be introduced by 1% ripple of a 6 pole tachometer into the identification of MIT/Asada arm parameters is shown in figure 13.16. The figure was made by evaluating equation (13.5); it shows the dependence of the *magnitude* of the bias upon the angle between the tachometer and the start position of the trajectory. The

bias affecting the original and optimized trajectories are shown. The bias peaks at 1.2 Kg-m^2 , or 7% of the magnitude of the parameter vector, $\|\theta^*\| = 17.7 \text{ kg-m}^2$.

In figure 13.17 the same effect is plotted as is in figure 13.16, except that the tachometer is presumed to have 3 poles. Because of a coincidental alignment between the disturbance and the eigenvector corresponding to the smallest eigenvalue of $\mathcal{R}$, the 1% ripple introduces a tremendous error in the parameter estimate. The bias is 67% as great as the parameter vector for some possible mechanical alignments of the tachometer (tachometer phase). The peak error introduced by this configuration is 59% as great as the maximum bias predicted by equation (11.15). Figure 13.17 shows that by choosing a very exciting trajectory the sensitivity of the experiment to systematic error is greatly reduced.




Figure 13.16 Magnitude of the Parameter Bias Vector, Plotted against the Phase Angle Between the Tachometer Poles and the Joint Angle. 6 Pole Tachometer.




Figure 13.17 Magnitude of the Parameter Bias Vector, Plotted against the Phase Angle Between the Tachometer Poles and the Joint Angle. 3 Pole Tachometer.

Chapter 14

Discussion

14.1 Why Seemingly Active Trajectories May Provide Poor Excitation

The basis functions, the columns of Φ , of a manipulator model are fundamentally different from those of linear systems. Manipulator models contain such elements as $\{1\bar{q}\}$ and $\{\cos^2({}_2q) * {}_1\bar{q}\}$, which can not be orthogonalized by any choice of input. The basis functions of linear systems are complex exponential functions which, if they are lightly damped or well separated in frequency, can be nearly orthogonal. The examples of poor excitation for linear identification presented in [Mareels et.al. 87] provide as much excitation as the most exciting optimized trajectories of this paper. It is not easy to choose an exciting manipulator trajectory - intuition can be wrong. A trial and error search of trajectories made up of harmonic functions and applied to the MIT/Asada arm experiment yielded $\kappa[\mathcal{R}]$ as high as 200,000, and none as low as 5,700, the best achieved by An, Atkeson and Hollerbach [85]. The researchers did well in their selection of trajectory. But trajectory space is huge and seems sparsely populated with good choices.

14.2 Trajectory Tracking versus Accurate Identification

Whether or not an experiment is sufficiently exciting, parameters identified using a particular trajectory will give a good model fit on that particular trajectory; if unmodeled dynamics are present, the estimator will strive to fit them. There is substantial evidence that good tracking performance may thus be achieved. Craig found that tracking error was minimized by setting the adaptation rate to a large value so that the parameter estimates swing wildly, [Craig 86]. Atkeson, An and Hollerbach [86] report that in each of three trials tracking error was reduced by using parameters identified on the experimental trajectory, even though large variations appear between parameter sets.

Fitting model components to unmodeled dynamics is a limited approach to good tracking: the model is valid only in an unknown neighborhood of the trajectory of the identification experiment; and the neighborhood can be quite small. Furthermore, because of poor excitation, a small systematic error can introduce a large bias into weakly identified parameter combinations. The weakly identified parameter combinations make little contribution to the torque along the experimental trajectory, but are precisely the combinations most sensitive to unmodeled dynamics. The existence of exciting trajectories shows that all combinations contribute torque along some trajectories. If a large misadjustment is introduced by fitting weakly excited parameter combinations to unmodeled dynamics, it may become a large torque error when that combination of parameters is excited along some different, possibly neighboring trajectory. Finally, there is no logical connection between the basis functions of the dynamic model and the unmodeled dynamics. It would be more reasonable, if one wished to fit the

unmodeled dynamics, to choose some simple, orthogonal functions rather than such arcane products of dynamic modeling as $\{2 \dot{q}_1 \dot{q}_3 \sin^2(q_2)(1 - 2\sin^2(q_3))\}$.

If tracking error is to be used as a measure of the quality of an identification algorithm, a trajectory different from the one over which the identification experiment was conducted should be used for demonstration. Otherwise unmodeled dynamics and the trajectory specific nature of the model must be addressed. For adaptive control algorithms, the transient error that occurs when a shift is made from one path to another, or the range of parameters identified when several different trajectories are used are far better measures of the quality of the model than is steady state tracking error on a repeated trajectory. It has been suggested that the parameters arriving at some steady state value in a sequential identification algorithm indicates that the algorithm has converged. The results presented here show that in practical situations the parameter convergence rates can vary from slow direction to fast direction by more than three orders of magnitude. No experiment presented has been run as long as the cases studied indicate is required for convergence.

14.3 Adaptive Control and Reduced Model Order

If the engineer is not free to choose the experimental trajectory, as will be the case with adaptive control in the work place, evaluating the conditioning of the $\mathcal{R}$ matrix will allow a determination of how many parameters may be adjusted. One might determine that a pick and place maneuver which is not sufficiently exciting to identify the parameters of the entire manipulator, might be sufficiently exciting to identify the mass of a new object in grasp.

One possible criterion for the elimination of parameters is whether the parameter itself is small. The explicit dynamic model of the serial Asada arm, [Armstrong 87; Atkeson, An and Hollerbach 86], shows that 6 of the inertial parameters depend only upon the off-diagonal inertia dyadic terms and center of gravity cross product terms. These parameters are zero in the CAD generated parameters, owing to the symmetry of the Asada arm. Eliminating these parameters from the model reduces the problem from 15 parameters to 9, and reduces the $\kappa[\mathcal{R}]$ of the best MIT/Asada arm identification trajectory from 5,700 to 154. If only nine parameters were identified, the best MIT/Asada arm identification trajectory would have provided sufficient data to estimate the parameters with a relative deviation of $\pm 1.2\%$, rather than $\pm 8\%$. More importantly, the bias susceptibility would be cut from 25 to 6.5. Even though eliminating off-diagonal inertial terms might introduce some unmodeled dynamic forces, if they are in fact non-zero, the reduction in bias susceptibility will almost certainly assure improved model accuracy over a large trajectory set.

14.4 Conclusions

Experimental noise poses a substantial challenge to the accurate identification of inertial parameters. We have seen that the accuracy of the identification is linked to the conditioning of the input correlation matrix and that this matrix is well conditioned only for very special choices of trajectory. The implication is that identification of all of the parameters of a manipulator dynamic model may not be experimentally feasible and that identification of reduced models and direct measurement of some parameters should be considered.

Bibliography

- An, C.H., Atkeson, C.G. and Hollerbach, J. M. 1985 (December).
"Estimation of Inertial Parameters of Rigid Body Links of Manipulators", *Proc. of the 24th Conf. on Decision and Control*, Fort Lauderdale, Florida: IEEE, pp. 990 - 1002.
- Anderson, B.D.O., Bitmead, R.R., Johnson, C.R. Jr., Kokotovic, P.V., Kosut, R.L., Mareels, I.M.Y., Praly, L. and Riedle, B.D. 1986.
Stability of Adaptive Systems, Cambridge: MIT Press.
- Armstrong, B., Khatib, O. and Burdick, J. 1986 (April 7-10).
"The Explicit Dynamics Model and Inertial Parameters of the PUMA 560 Arm," *Proc. 1986 Inter. Conf. of Robotics and Automation*, San Francisco: IEEE.
- Armstrong, B. 1987 (March 31 - April 3).
"On Finding Exciting Trajectories for Identification Experiments Involving Systems with Non-Linear Dynamics," *Proc. of the 1987 IEEE International Conf. on Robotics and Automation*, Raleigh, N.C.: IEEE, pp. 1131 - 1139.
- Asada, H. and Youcef-Toumi, K., 1984.
"Analysis and Design of a Direct-Drive Arm with a Five-Bar Parallel Drive Mechanism," *ASME Journal of Dynamic Systems, Measurement and Control*, v. 106, pp. 225-230.
- Atkeson, C.G. An, C.H. and Hollerbach, J. M. 1986 (Fall).
"Estimation of Inertial Parameters of Manipulator Loads and Links ", *Inter. Journal of Robotic Research* 5(3)101:119.

Bryson, A.E. and Ho, Y. 1975.

Applied Optimal Control, New York: Hemisphere Publishing Co.

Craig, J.J. 1986.

"Adaptive Control of Mechanical Manipulators," *PhD Thesis*, Electrical Engineering Dept., Stanford Univ.

Craig, J.J., Hsu, P. and Sastry, S.S. 1987 (Summer).

"Adaptive Control of Mechanical Manipulators," *International Journal of Robotics Research*, 6(2)16:28.

Finn, J.D. 1974.

A General Model for Multivariate Analysis, New York: Holt Rinehart and Winston, inc.

Hsu P., Bodsom, M., Sastry, S. and Paden, B. 1987 (March 31 - April 3).

"Adaptive Identification and Control for Manipulators without Using Joint Accelerations," *Proc. of the 1987 Inter Conf. on Robotics and Automation*, Raleigh, N.C.: IEEE, pp. 1210 - 1216.

Kalaba, R. and Spingarn, K. 1982.

Control, Identification, and Input Optimization, New York: Plenum Press.

Kendall, M. G. and Stuart, A. 1973.

The Advanced Theory of Statistics, New York: Hafner Publishing Co.

Kohsla, P.K. and Kanade, T. 1985 (December).

"Parameter Identification of Robot Dynamics," *Proc. of the 24th Conf. on Decision and Control*, Ft. Lauderdale: IEEE.

Ljung, L. 1987.

System Identification: Theory for the User, Englewood Cliffs, New Jersey: Prentice-Hall, Inc.

Mareels, I.M.Y, Bitmead, R.R., Gevers, M., Johnson, C.R. Jr., Kosut, R.L. and Poubelle, M.A. 1987 (January).

"How Exciting can a Signal Really Be?," *Systems & control Letters*, 8(3)197:205.

Mendel, J.M. 1973.

Discrete Techniques of Parameter Estimation, New York: Marcel Dekker, Inc.

Middleton, R.H. and Goodwin, G.C., 1988.

"Adaptive Computed Torque Control for Rigid Link Manipulations," *Systems & control Letters*, 10(1)9:16.

Mukerjee, A. and Ballard, D. 1985 (April 7-10).

"Self-Calibration in Robot Manipulators," *Proc. of the 1986 Inter. Conf. on Robotics and Automation*, San Francisco: IEEE.

Neuman, C.P. and Khosla, P.K. 1985 (May 29-31).

"Identification of Robot Dynamics: An Application of Recursive Estimation," *Proc. of the Fourth Yale Workshop on Applications of Adaptive Systems Theory*. Narendra, K.S. (ed.); Yale University, New Haven, Conn.

Olsen, H.B. and Bekey, G.A. 1986 (March 31 - April 3).

"Identification of Robots Dynamics," *Proc. of the 1986 Inter. Conf. on Robotics and Automation*, San Francisco: IEEE.

Olsen, H.B. and Bekey, G.A. 1985 (April 7-10).

"Identification of Parameters in Models of Robots with Rotary Joints," *Proc. of the 1985 Inter. Conf. on Robotics and Automation*, St. Louis: IEEE.

Press, J. 1972.

Applied Multivariate Analysis, New York: Holt, Rinehart and Winston, Inc.

Sastry, S.S. 1984.

"Model-Reference Adaptive Control-Stability, Parameter Convergence and Robustness," *IMA J. of Mathematical Control & Information*, 1984(1)27:66.

Slotine, J.J.E. and Li, W. 1986.

"On The Adaptive Control of Robot Manipulators," *Proc. A.S.M.E. Winter Annual Meeting, Anaheim, Ca.*

Slotine, J.J.E. and Li, W. 1987a.

"Adaptive Strategies in Constrained Manipulation," *Proc. of the 1987 Inter. Conf. on Robotics and Automation, Raleigh, N.C.: IEEE*, pp. 595-601.

Slotine, J.J.E. and Li, W. 1987b.

"Adaptive Manipulator Control: a Case Study," *Proc. of the 1987 Inter. Conf. on Robotics and Automation, Raleigh, N.C.: IEEE*, pp. 1392-1400.

Tarn, T.J., Bejczy, A.K., Han, S. and Yun, X. 1985 (September).

"Inertia Parameters of PUMA 560 Robot Arm," Robotics Laboratory Report, SSM-RL-85-01, Dept. of Systems Science and Mathematics, Washington Univ. St. Louis, Mo.

Widrow, B., McCool, J.J., Larimore, M.G. and Johnson, C.R. Jr. 1976 (August).

"Stationary and Non-stationary Learning Characteristics of the LMS Adaptive Filter," *Proceedings of the IEEE*, 45(8)1151:62.

Widrow, B. and Stearns, S.D. 1985.

Adaptive Signal Processing, Englewood Cliffs: Prentice Hall Inc.