

Robot Arm Control Exploiting Natural Dynamics

by

Matthew M. Williamson

M.Eng, University of Oxford (1993)
S.M., Massachusetts Institute of Technology (1995)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 1999

© Massachusetts Institute of Technology 1999

Signature of Author
Department of Electrical Engineering and Computer Science
May 18, 1999

Certified by
Rodney A. Brooks
Professor of Electrical Engineering and Computer Science, MIT
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Departmental Committee on Graduate Students

Robot Arm Control Exploiting Natural Dynamics

by

Matthew M. Williamson

Revised version of a thesis submitted to the
Department of Electrical Engineering and Computer Science
on May 18, 1999, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

This thesis presents an approach to robot arm control exploiting natural dynamics. The approach consists of using a compliant arm whose joints are controlled with simple non-linear oscillators. The arm has special actuators which makes it robust to collisions and gives it a smooth compliant motion. The oscillators produce rhythmic commands of the joints of the arm, and feedback of the joint motions is used to modify the oscillator behavior. The oscillators enable the resonant properties of the arm to be exploited to perform a variety of rhythmic and discrete tasks. These tasks include tuning into the resonant frequencies of the arm itself, juggling, turning cranks, playing with a Slinky toy, sawing wood, throwing balls, hammering nails and drumming.

For most of these tasks, the controllers at each joint are completely independent, being coupled by mechanical coupling through the physical arm of the robot. The thesis shows that this mechanical coupling allows the oscillators to automatically adjust their commands to be appropriate for the arm dynamics and the task. This coordination is robust to large changes in the oscillator parameters, and large changes in the dynamic properties of the arm.

As well as providing a wealth of experimental data to support this approach, the thesis also provides a range of analysis tools, both approximate and exact. These can be used to understand and predict the behavior of current implementations, and design new ones. These analysis techniques improve the value of oscillator solutions.

The results in the thesis suggest that the general approach of exploiting natural dynamics is a powerful method for obtaining coordinated dynamic behavior of robot arms.

Thesis Supervisor: Rodney A. Brooks

Title: Professor of Electrical Engineering and Computer Science, MIT

Acknowledgments

I would like to thank my advisor Rodney Brooks for his encouragement and support over these years. He has given me a great deal of intellectual freedom, and his ideas have been a source of inspiration for this work.

I would also like to thank Tomás Lozano-Pérez. A series of discussions with Tomás galvanized me into action and are responsible for this thesis being as comprehensive as is. Thank you.

I would also like to thank Stefan Schaal for originally suggesting that I try using oscillators.

The work in this thesis is from the Cog project, and I would like to thank Brian Scassellati, Cynthia Breazeal and Robert Irie for their support. Matt Marjanović was particularly helpful in the final stages of this thesis with the drumming task, figures, photos and general help. Charlie Kemp has also been kind and supportive. I would also like to thank all the people in the Leg Lab at MIT, particularly Jerry Pratt, Hugh Herr, Dan Paluska, Peter Dilworth, Jianjue Hu and Gill Pratt. Thank you for your friendship and encouragement.

Annika Pfluger kindly helped with the videos for the defense.

I am also indebted to the officemates that I have had at MIT: Tim Tuttle, Jessica Banks, and Una-May O'Reilly. Una-May proof read many papers and bits of this document, for which I am very grateful.

I would also like to thank Ulf Jonsson, Jorge Gonçalves and Alexandre Megretski for their help with the non-linear analysis section.

On a more personal note, I would like to thank my good friends Joel Minnich, Maura McCarthy and Fiona Percy for their love and support over the years. I would also like to thank Esther Chamberlain for her love, encouragement and patience.

This research was supported in part by the Defense Advanced Research Projects Agency and was monitored by the Office of Naval Research under contract numbers N00014-83-K-0125 and N00014-84-K-0099. The author was supported with a Engineering and Physical Sciences Research Council (EPSRC) scholarship.

Contents

1	Introduction	15
1.1	Synopsis	15
1.2	Approach	16
1.3	Comparisons	18
1.4	Contributions	20
1.5	Thesis outline	20
1.6	Note on data in thesis	21
2	Literature review	23
2.1	Introduction	23
2.2	Exploiting natural dynamics	23
2.3	Behavior-based robotics	25
2.4	Previous work using oscillators	26
2.5	Oscillator analysis	28
2.6	Biological central pattern generator evidence	29
2.7	Human arm control	29
2.8	Conclusion	30
3	Single degree of freedom motions	31
3.1	Introduction	31
3.2	Analysis methodology	32
3.3	Oscillator describing function	35
3.4	Predictive accuracy	36
3.5	Local stability	38
3.6	Effect of parameter changes (a): Design	40
3.7	Effect of parameter changes (b): Robustness	41
3.8	Design example: Resonance tuning	44
3.9	Design example: Juggling	47
3.10	Adaptation	51
3.11	Conclusion	51
4	Coupling through natural dynamics	55
4.1	Examples	56
4.2	Coupling model	62
4.3	Local stability	68
4.4	Summary of conclusions from the model	68
4.5	Crank turning	70
4.6	Self-organization to find resonant mode	73
4.7	Robustness	75
4.8	Quality of crank turning solution	81
4.9	Slinky toy	81
4.10	Conclusion	84

5	Coupling through explicit connections	85
5.1	Introduction	85
5.2	Case (a): Networks of oscillators	85
5.2.1	Connections and feedback	87
5.3	Case (b): Master-slave network	91
5.4	Case (c): Modifying the natural dynamics	95
5.4.1	Example: Circular motion of two joints	96
5.5	Conclusion	99
6	Oscillator analysis	101
6.1	Introduction	101
6.2	Boundedness of oscillator	101
6.3	Oscillator as piecewise-linear system	104
6.4	Oscillatory solution	107
6.5	Connecting oscillators to systems	108
6.5.1	Example: No oscillation	109
6.5.2	Example: Constant input	109
6.5.3	Example: Mass-spring system	110
6.6	Prediction of entrainment	111
6.6.1	Example: Non-minimum phase system	113
6.7	Local stability of the limit cycle	113
6.8	Local stability for the oscillator	117
6.8.1	Example: Mass-spring system	118
6.9	Global stability	118
6.10	Conclusion	119
7	Examples	121
7.1	Drumming	121
7.1.1	Counting	125
7.1.2	Hitting on the offbeat	126
7.1.3	Hitting more than one drum	127
7.2	Discrete from rhythmic	127
7.3	Throwing	129
7.4	Hitting	133
7.5	Conclusion	137
8	Conclusion	139
8.1	Further work	139
8.2	Vision	140
8.3	Summary	141
A	Experimental Apparatus	151
A.1	Overall arm design	151
A.2	Actuators	152
A.3	Spring design	154
A.4	Actuator control law	156
A.5	Arm design	158
A.5.1	General	158
A.5.2	Joint configurations	158
A.6	Arm technical data	164
A.7	Sensors	164

A.8 Electrical wiring	168
A.9 Control hardware	170
A.10 Hand design	177
A.11 Conclusion	177
B Arm and oscillator parameter tuning	179
B.1 Parameters	179
B.2 Automatic tuning	180
B.3 Procedure for new task	181
C Oscillator Inputs	183
D Oscillators and torque feedback	187
D.1 Torque feedback	187
D.2 Slinky example	188
D.3 Conclusion	190

List of Figures

1-1	The humanoid robot Cog	16
1-2	Joint level control of arm	17
1-3	Connections between oscillators and the arm	18
1-4	Comparison between traditional control and exploiting natural dynamics	19
1-5	Comparison between traditional control and oscillator control	19
3-1	Schematic of oscillator	33
3-2	SIM Oscillator behavior under changing tonic c and time constant τ_1	33
3-3	SIM Sample output from the oscillator	34
3-4	SIM Effect of increasing input signal on entrainment	34
3-5	SIM Minimum input for entrainment	35
3-6	Schematic of oscillator connected to a single degree of freedom system	36
3-7	SIM Oscillator Bode plot	37
3-8	SIM Plot showing intersection of $G(j\omega)$ and $1/N(j\omega, A)$ in the complex plane, predicting a limit cycle solution of the coupled system	37
3-9	SIM Power density spectrum of oscillator output signal	38
3-10	SIM Predictive accuracy of describing function analysis	39
3-11	REAL Phase plot of robot link motion	39
3-12	Oscillator gain schematic	40
3-13	SIM Effect of gain on describing function of the oscillator	42
3-14	SIM Robustness to changes in gain and time constants	43
3-15	SIM Robustness to changes in damping ratio	44
3-16	SIM Bode plot for Van der Pol oscillator	45
3-17	SIM Limit cycle prediction for Van der Pol oscillator	46
3-18	SIM Accuracy of limit cycle prediction, and resonance tuning effect for Van der Pol oscillator	47
3-19	REAL Robot results for resonance tuning with Van der Pol oscillator	48
3-20	Picture of juggling implemented on the robot	48
3-21	Schematic of juggling action	49
3-22	Oscillator driven juggling	49
3-23	SIM Robustness of juggling to changes in α	50
3-24	REAL Time traces from robot juggling	51
3-25	REAL Effect of external disturbance on oscillator entrainment	52
4-1	Configuration of oscillators and arm, where the oscillators are connected by mechanical coupling, not explicit connections	55
4-2	REAL Crank turning with one arm	56
4-3	REAL Effect of feedback on two degree of freedom crank turning	57
4-4	REAL Crank turning with two arms	58
4-5	REAL Crank turning using a redundant arm	59
4-6	REAL Transients of crank turning with a redundant arm	60
4-7	Cog and Slinky toy	61
4-8	REAL Two transients of Slinky operation	61

4-9		A simple model of coupling through the natural dynamics	62
4-10	SIM	Solutions to the damping constraint (4.6)	64
4-11	SIM	Solutions for the final system frequency	65
4-12	SIM	Eigenmodes of underlying mechanical system and oscillator driven system	66
4-13	SIM	Robustness of oscillator to finding mode of resonant system	67
4-14	SIM	Time transients from oscillator control of coupled system	67
4-15	SIM	Transient from Van der Pol oscillator control of the coupled system	68
4-16	SIM	Local stability of oscillator driven system	69
4-17		[.	70
4-18	REAL	Comparison of crank data with coupling model: imaginary parts	71
4-19	REAL	Comparison of crank data with coupling model: real parts	72
4-20	REAL	Transient of oscillator driven crank turning	74
4-21	REAL	Plot showing robustness of the oscillator driven crank turning to changes in arm stiffness and oscillator time constants	76
4-22	REAL	Crank robustness data for non-sensitive parameters	77
4-23	REAL	Robustness to extra masses attached to arm during crank turning	78
4-24	SIM	Robustness: effect of posture changes	78
4-25	REAL	Crank robustness data for arm posture parameters	79
4-26	REAL	Crank robustness data for joint amplitude changes	80
4-27	REAL	Real and reactive power during crank turning	82
4-28	REAL	Ratio between real and reactive power during crank turning	83
5-1		Three methods for coupling oscillators	86
5-2		Simple oscillator network	87
5-3	SIM	Plot of relative phase between oscillators in network as function of parameter γ	88
5-4		Two oscillators driving two independent masses using a network with internal connections	88
5-5	SIM	Transients of coupled system	89
5-6	SIM	Inputs to the oscillator while driving the coupled system	90
5-7		Master-slave oscillator network	91
5-8	REAL	Sawing using oscillators	92
5-9	REAL	Energy expended in sawing versus oscillator natural frequency	93
5-10	REAL	Sawing frequency versus oscillator frequency, varying the oscillator feedback	93
5-11	REAL	Energy expended in sawing versus sawing frequency	94
5-12	REAL	Percentage of wood cut per stroke plotted against average sawing frequency	95
5-13		Application of potential field for producing circular motion	97
5-14	REAL	Effect of potential field on arm motion	98
5-15	REAL	Effect of oscillator feedback and potential field on arm motion	98
6-1		Oscillator coupled to a linear system $G(s)$	104
6-2		Plot of linear regions and matrices for oscillator	105
6-3	SIM	Linear regions, fixed points and eigenvalues for oscillator	106
6-4	SIM	Intuitions for oscillator limit cycle behavior	107
6-5		Oscillator coupled to a linear system	108
6-6	SIM	Two transients of oscillator connected to $G(s) = (s + 0.1)/(2s + 0.1)$	110
6-7	SIM	Transients of the system with a constant input	111
6-8	SIM	Two transients of the oscillator driving a mass spring system	112
6-9		Oscillator tightly coupled to a driven system	112
6-10	SIM	Prediction of limit cycles for oscillator connected to non-minimum phase system	114
6-11	SIM	Two transients of oscillator driving non-minimum phase system	114
6-12		Schematic of Poincaré map	115

6-13		Schematic of limit cycle in a piecewise-linear system	116
6-14	SIM	Oscillator limit cycle plotted in the $x_1 - x_2$ plane	118
6-15	SIM	Plot of limit cycle for mass-spring system and oscillator	119
7-1		Schematic of drumming	122
7-2		Picture of Cog hitting a drum	122
7-3		Drumstick holder	123
7-4	REAL	Plot showing the effect of feedback on the drumming system	124
7-5	REAL	Two time plots for the oscillator entraining with an external beat	124
7-6	REAL	Plot showing examples of different rhythms possible using counting	125
7-7	REAL	Plot showing two stable limit cycles when the arms both hit on every other beat	126
7-8	REAL	Changing between auditory-motor limit cycles due to an auditory disturbance	127
7-9	REAL	Three different rhythms which use the motion of the dominant arm to inhibit the auditory processing	128
7-10	REAL	Switching between different rhythms	129
7-11	REAL	Four stills from a video of the robot drumming	130
7-12		Schematic of creating discrete motions from the oscillator output	130
7-13	REAL	Series of stills from sequence of the arm throwing a ball	131
7-14	REAL	Plot of joint angles during throwing	132
7-15	REAL	Stills from sequence of overarm throwing	133
7-16	REAL	Four stills from a video of the robot hitting the punchbag	134
7-17	REAL	Motion of the arm in the vertical plane during hitting	135
7-18	REAL	Height of the arm during hitting	135
7-19	REAL	Four stills from a video of the robot hammering a nail	136
A-1		Picture of robot arms	151
A-2		Arm kinematics	152
A-3		Schematic of actuator design	153
A-4		Cross shaped spring design	154
A-5		Mounting of strain gauges on spring	155
A-6		Comparison of cross-shaped spring with a flat plate in torsion	155
A-7	REAL	Comparison of cross-spring and flat plate stiffness	156
A-8	REAL	Comparison of cross-spring and flat plate maximum torque	157
A-9		Actuator control law	157
A-10		Joint level control of arm	158
A-11		Arm joint configuration I	160
A-12		Arm joint configuration II	161
A-13		Bearing configuration	162
A-14		Cable tensioning configuration	163
A-15		Alternative cable configurations	164
A-16		Robot shoulder joint	165
A-17		Upper arm section	166
A-18		Arm curled up	166
A-19		Lower arm design	167
A-20		Layout of sensors on robot joint	169
A-21		Sensor board	170
A-22		Connections to the sensor board	171
A-23		Schematic of sensor board	172
A-24		Wiring diagram for the robot arm	173
A-25		Physical routing of electrical cables	174
A-26		Motor board	175

A-27	Control hardware for the arm	176
A-28	Front End Processor, 68332 network and connections to other computers	176
C-1	SIM Effect of non-linear and linear inputs on oscillator output amplitude	184
C-2	SIM Describing function of oscillator with and without linear inputs	185
D-1	A simple model of coupling through the natural dynamics	187
D-2	SIM Local stability under torque feedback	188
D-3	Picture of Cog passing the Slinky toy from hand to hand	189
D-4	REAL Two examples of Slinky operation	190
D-5	REAL Comparison between Slinky data and coupling model: imaginary parts	191
D-6	REAL Comparison between Slinky data and coupling model: real parts	191
D-7	REAL Mode of Slinky operation	192

List of Tables

2.1	Table of literature which exploits natural dynamics	24
2.2	Table of literature of previous work using oscillators	27
4.1	REAL Slopes of lines to compare model and robot data for crank turning	73
A.1	REAL Technical data for the individual arm joints	164

Chapter 1

Introduction

1.1 Synopsis

This thesis presents an approach to robot arm control exploiting the natural dynamics of the arm and its environment. The idea is to “let the physics do the computation”. This results in a system which is versatile, robust, computationally simple and easy to implement.

The approach is motivated by related work in robotics. A number of researchers have presented examples of systems which exploit the natural dynamics of the system to perform tasks in a simple manner. These examples include the passive dynamic walkers of McGeer (1990), the dynamic running machines of Raibert (1986), and the open loop stable juggling machines of Schaal and Atkeson (1993). These approaches either use the dynamics directly, or augment the natural dynamics with simple controllers to provide robust performance.

My approach extends these ideas to a number of rhythmic and discrete tasks using a robot arm. The passive dynamics of the arm are manipulated to align them with the task, so that the natural behavior of the complete system is to perform the task. The motion is achieved using a controller to inject energy into the system.

In particular, my approach consists of using a compliant arm whose joints are controlled using simple non-linear oscillators. The arm has special actuators which make it robust to collisions and give it a smooth compliant motion. The actuators are used to implement low gain position control at each joint. This makes the robot links appear as if they were connected by springs and dampers, giving the whole arm a rich mass-spring behavior. The dynamics can be changed by altering arm posture, stiffness, damping and the manipulated object to match the passive arm dynamics with the task. Non-linear oscillators are used to inject energy and so generate the motion. The oscillators produce rhythmic commands at the arm joints which excite the arm dynamics. The oscillators are adaptive, using feedback from the arm joints to alter the frequency and phase of their outputs. The oscillator behavior is to adjust the commands relative to the arm and task dynamics; this appropriately adds energy to the arm, and produces the required motion.

Using the same compliant arm, oscillators and feedback, a wide variety of tasks have been implemented. These are all dynamic tasks where the spring-like properties of the arm are exploited to produce the motion. The tasks include tuning into the resonant frequencies of the arm itself, juggling, turning cranks in a variety of configurations using both one and two arms, playing with a Slinky toy, sawing wood, throwing balls, hammering nails and even drumming.

For most of these tasks, the oscillators at the joints are completely independent. The oscillators use the mechanical coupling through the arm to coordinate with one another and the task at hand. The variety of coordinated tasks that are possible with such a distributed system highlights the sensitivity of the individual oscillators to the arm dynamics. In fact, the thesis shows that coupling oscillators through the natural dynamics is more robust and more powerful than the more usual method of connecting them into networks.

This precise coordination between the oscillators and the arm comes hand in hand with remarkable robustness. The system is robust in two respects. The same task can be accomplished using a

wide variety of oscillator parameters. On the other hand, using a fixed set of oscillator parameters, the task can be accomplished for a wide variety of different system stiffnesses, inertias, and other properties. The robustness is a desirable property for a practical control scheme.

Researchers are sometimes reluctant to use oscillators to control their robots, because they do not have the tools to understand how they work, or how best to set the parameters. The complaint that oscillators of the type used in this thesis are difficult to tune is a common one. To answer these concerns, this thesis develops a variety of oscillator design and analysis techniques. These provide insight into the oscillator behavior, and give clear directions for tuning and design.

The following sections present the approach in more detail, comment on the relationship between this work and more traditional control methods, state the contributions of this thesis, and introduce the following chapters.

1.2 Approach




Figure 1-1: The humanoid robot Cog. The two arms used in this thesis are mounted on this robot. Each arm has 6 degrees of freedom arranged in a similar manner to a human arm. The arms are also approximately the same length as a human arm.

This section describes the approach taken in this thesis: the choice and use of a compliant robot arm, and the use of simple nonlinear oscillators to excite and exploit the natural dynamics of the system.

When designing a robot arm, one important consideration is the structural stiffness of the arm. The choice between a stiff robot and a more compliant one rests mainly on the types of tasks envisaged for the robot. For accurate position control and for high bandwidth force control, the robot should be stiff. For tasks which require a robust interaction with an uncertain world, the robot should be designed to be compliant. The dynamical interactions between the links of a compliant arm are more noticeable compared to a stiff arm, which makes the arm's position more difficult to control. For the work in this thesis, a compliant arm was constructed to be used for tasks which

required interaction rather than positional accuracy. The extra dynamics are exploited to perform tasks simply.

Two six degree of freedom arms were designed and used for this thesis, both mounted on the humanoid robot Cog (Brooks and Stein, 1994, Brooks et al., 1998), shown in figure 1-1. The arms use series elastic actuators at each joint, an actuator technology which incorporates a physical spring in series with the motor output. (Pratt and Williamson, 1995, Williamson, 1995). By measuring and controlling the deflection of the spring, the force output of the actuator can be controlled. The spring contributes to the smooth force output of the actuator by naturally filtering out the noise and backlash introduced by the gearbox. The spring also makes it easy to ensure that the actuator is passive and therefore stable when interacting with passive environments (Colgate and Hogan, 1989). Another important role for the spring is to absorb shock loads, protecting the motor gearbox from damage. These characteristics greatly contribute to the robustness of the arm design. The design of the arms is described in more detail in Appendix A.

The series elastic actuators provide force control, giving a force output at the joint which is close to the desired force. To control the position of the joint, a low gain proportional-derivative (PD) controller is used, whose output is the desired torque or force at the joints, as shown in figure 1-2. The stiffness and damping at the joints can be varied by changing stiffness K or damping B , and the arm moved around by changing the setpoint θ_v of the controller. The PD control makes the arm behave as if its links are connected by springs and dampers, but because the forces in the joints are accurately controlled by the series elastic actuators, the overall motion is smooth and compliant. The force control bandwidth of the actuators is fairly low, which forces the stiffness of the arm to be low. This gives robustness when interacting with objects.




Figure 1-2: The joint level control consists of a position control loop, with stiffness K , damping B , and setpoint θ_v organized around an inner torque control loop provided by the series elastic actuators. The inner loop controls the force output of the joint u to accurately track the desired force u_d . Organizing the control in this fashion results in smooth, compliant motion of the arm, where the overall stiffness and damping can be altered by changing the values of K and B .

To control the arms a set of simple non-linear oscillators are used to alter the setpoints for each joint. The oscillators are models of two biological neurons in mutual inhibition which form a resonant circuit (Matsuoka, 1985). The oscillators respond to the dynamical state of the arm using feedback of either the joint angle or the joint torque. The system is thus tightly coupled, as shown in figure 1-3, because the output from each oscillator drives the joint, and the feedback adjusts the oscillator output. In most of the work for this thesis, there are no connections between the oscillators. Instead the oscillators rely on mechanical coupling between the joints given by the natural dynamics to coordinate with one another.

The oscillators form a dynamical system which, without input, produces a rhythmic output. If a rhythmic input is applied to an oscillator, the oscillator will entrain with that input, producing an output signal of the same frequency. This occurs over a reasonably large range of frequencies and input amplitudes. The entrained behavior is complex, but can be roughly described as “resonant”




Figure 1-3: The oscillators are connected to each joint of the arm, and are tightly coupled so that the oscillator output θ_{vi} drives the joint, and the oscillator input is either the joint angle θ_i or joint torque u_i . There is generally no explicit connection between the oscillators. The oscillators rely on mechanical coupling through the physical structure of the robot, sensed using feedback to coordinate with one another.

where the oscillator drives the system at a frequency close to resonance, or in a mode similar to the resonant mode of the complete system.

The oscillator behavior makes them suitable as adaptive controllers for the arm. Once the arm has been set up so that the natural dynamics of the coupled arm-task system is appropriate for the task, the oscillator behavior is to respond to the system dynamics and inject energy to produce the required motion. The work in this thesis concentrates on demonstrating, explaining and understanding the behavior of the oscillators in conjunction with the arm natural dynamics. This includes design methods for the oscillator parameters, but does not address how to choose and design the arm parameters: arm posture, stiffness etc.. In most cases the initial arm configuration was chosen by hand, with automatic tuning used to modify that initial configuration. This is described in appendix B.

1.3 Comparisons

The main difference between traditional robot control and the approach taken in this thesis is the role of the robot dynamics. In traditional control, the robot is viewed as a general purpose manipulator which performs tasks independent of the robot configuration. The task is specified in terms of the desired motion (force, position, compliance) of the robot, and the robot control enforces that command. The robot dynamics are generally ignored or canceled, and certainly do not play a part in how the task is planned. The approach taken in this thesis is the opposite: the robot dynamics are crucial for the performance of the task as they determine the range of possible tasks, and also how the tasks are accomplished. The robot dynamics are specified so that the task motion is a passive behavior of the system, and the oscillators are used to inject energy into the arm and so create the motion.

This difference is illustrated in figure 1-4. The task illustrated is that of moving a mass backwards and forwards. In traditional control, the dynamics of the robot are removed, so the equivalent connection between the desired position of the mass x_d , and its actual position x is stiff. The x_d trajectory is required to move the mass backwards and forwards, and the controller needs to overcome the inertial and frictional forces on the mass. If the dynamics of the arm are exploited, represented here by a spring, the situation is somewhat simpler. The natural behavior of the mass is to vibrate on the spring and so move backwards and forwards. The role of the trajectory x_d is now to inject and remove energy to sustain the motion, not create the whole motion.




Figure 1-4: Comparison between traditional control **(a)**, and exploiting natural dynamics **(b)** for the task of moving a mass backwards and forwards. Under traditional control the dynamics of the robot are not used. The controller enforces a stiff structure which makes the actual mass position x track closely the desired position x_d . If the robot is made to be compliant, then its dynamics can be exploited to perform the same task in a different manner. The mass will naturally vibrate on the spring of the robot dynamics, and the role of the desired trajectory is to sustain the motion, rather than create the whole motion.

The traditional approach is more general, since the mass can be moved in any arbitrary trajectory x_d . However, for rhythmic tasks, the alternative has some advantages. One consequence of exploiting the dynamics is that the arm needs to be compliant. This has the benefit of giving robust interaction with objects and unexpected collisions. The traditional controllers need to be stiff to reduce tracking errors. This stiffness causes problems in practice: unexpected collisions are not dealt with robustly by high gain position controlled systems, and high gain force control is known to be difficult because of stability issues.




Figure 1-5: Comparison between traditional control **(a)**, and oscillator control **(b)**. In traditional control the plant's state x is controlled to follow the desired state x_d using a controller. This controller could have internal state as indicated by the ■, say from an integral term. The oscillator control has no explicit desired trajectory, this being generated internally by the oscillator dynamics. The oscillator modifies its control action u dependent on the system state x . This has the advantage that a low gain system can be used where the complexity of planning x_d is carried out by the oscillator dynamics rather than by some other system.

The second difference is in terms of the design of the controller, and is illustrated schematically in figure 1-5. While a traditional controller requires a desired trajectory x_d , the oscillator control generates that signal using its internal dynamics. As before, the traditional approach is more general, the oscillator system being restricted to the trajectories that are generated by the oscillator dynamics. Also as before, arranging the control in this way has some practical advantages.

The main advantage of using the oscillator is that the desired trajectory is reactive to the dynamics of the system. Referring to the mass-spring system in figure 1-4, the oscillator can generate a trajectory which complements the motion of the mass, by injecting and removing energy. The x_d generated by the oscillator is reactive since it is calculated within a tight loop, and is synchro-

nized with the system motion since it is generated relative to the state x . These characteristics are achieved without a separate system to calculate x_d , and without the extra sensing, modeling and computation that calculation of x_d would require.

In the case of oscillators controlling multiple joints of the arm, this internal generation of trajectories is even more advantageous. The oscillators are independent, coupled only through the arm dynamics. The trajectories for all the joints are thus generated in a distributed manner with coordination which is correct relative to the motion of the arm. This contrasts with the complexity of the system which would be needed to generate explicit trajectories for all the joints. This difference is accentuated by the versatility of the oscillator system. While calculating x_d for one task is relatively straightforward, repeating this for each joint and each new task would be tedious or require the extra complexity of kinematic modeling and calibration.

A further difference in the robotic case is that the oscillator control system does not deteriorate as the speed of the task increases and the dynamics of the arm become significant. Both position and force control for robots degrades at high speeds because of disturbances from the arm dynamics. If the arm dynamics are aligned with the task, and as the speed increases those dynamics remain aligned with the task, then the oscillator system will be robust to the change in speed.

1.4 Contributions

The contributions of this thesis are

- The demonstration of the versatility of exploiting natural dynamics by showing implementations of a wide variety of different tasks on a real robot arm. These tasks include both rhythmic tasks (e.g. juggling, crank turning, sawing) and discrete tasks (e.g. throwing, hitting, hammering).
- The demonstration of how seemingly complex tasks become simple when cast as resonances of the arm and task. They can then be performed both simply and robustly using oscillators.
- The demonstration that exploiting the natural dynamics using non-linear oscillators is robust and easy to implement.
- The demonstration that the principles behind exploiting the natural dynamics can be applied to create motions which are not completely dictated by the natural dynamics.
- The development of approximate analysis techniques for non-linear oscillators tightly coupled to mechanical systems, and to systems of oscillators coupled through mechanical coupling. These techniques predict the final motion, and so ease design and tuning. They also provide a powerful means of evaluating the robustness and applicability of oscillator based solutions.
- The presentation of exact theoretical results and analysis techniques which describe the oscillator capabilities, performance and stability. These results complement the approximate techniques.

1.5 Thesis outline

The thesis proceeds as follows:

Chapter 2 reviews the relevant literature in robotics. It considers other robotic approaches which exploit natural dynamics and other approaches using oscillators. It also provides evidence taken from neuroscience and motor psychophysics that humans also exploit their natural dynamics.

Chapter 3 introduces oscillator solutions for single degree of freedom motions, driving single links of the robot arm. The oscillator system uses feedback to coordinate its command with the arm,

responding on a cycle by cycle basis to the arm state. The solutions are also robust in two respects. A single set of oscillator parameters can accomplish a task for a wide range of system properties (stiffness, inertia, damping etc.), and for a fixed system, a wide range of oscillator parameters can be used to accomplish the task. The chapter introduces analysis techniques which can be used to predict the behavior of oscillator controlled systems, applied to different types of oscillator (e.g. Matsuoka, Van der Pol etc.), and applied to quite complex tasks such as juggling. The analysis is also used to explain the origin of the oscillator robustness.

Chapter 4 examines the behavior of oscillators driving the arm using multiple degrees of freedom, exploiting mechanical interactions through the arm itself to couple and coordinate the arm motion. The analysis tools introduced in chapter 3 are extended to model this multiple degree of freedom case. They indicate that the oscillators find the resonant mode of the underlying mechanical system. The oscillators find this mode automatically and can thus perform complex coordinated tasks in a distributed manner. The main example presented is crank turning. The crank attached to the arm creates a resonant mode which is the correct motion. The oscillators adapt to exactly coordinate with one another to turn the crank, independent of the arm configuration. This coordination behavior is also robust. The oscillators maintain coordination with the task in spite of wide variations in parameters or arm properties.

Chapter 5 moves away from a description of the oscillator properties to consider how to design motions using oscillators while maintaining their robustness properties. The chapter shows that connecting oscillators into networks with fixed connections is not as robust as using coupling them through the natural dynamics. This is because the mechanical coupling “frees up” the oscillator to be synchronized with the task. The explicit connections force the oscillator outputs to be synchronized with each other, rather than the task. The chapter shows how using a single oscillator to drive multiple degrees of freedom or, alternatively, manipulating the dynamics of the arm give an ideal compromise: the desired motion is achieved with robustness.

Chapter 6 returns to the question of analysis and design for oscillator driven systems. The analysis methods presented in chapters 3 and 5 are practically useful but only approximate the system behavior. This chapter presents exact results concerning the boundedness of the oscillator system, its behavior when coupled to systems, and the local stability of oscillator driven limit cycles. It also shows an alternative technique for describing the oscillator and driven systems as piecewise-linear systems which sheds light and intuition on the oscillator behavior. These results are important and useful in terms of design, giving tools for determining the oscillator behavior when coupled to a variety of systems.

Chapter 7 brings the thesis back to the practical implementation of tasks using oscillators. It demonstrates the versatility of oscillator solutions to single motion tasks such as throwing, hitting and hammering, as well as tasks which are a combination of rhythmic and discrete motions, such as drumming. The chapter shows that these tasks are easy to implement with oscillators and that the feedback provides useful behavior—powerful hitting and drumming at a steady rhythm. The chapter also shows the versatility of the oscillators in terms of the sensory signals that they can use to coordinate with the world, using both information related to the arm for throwing, and auditory signals for the drumming application.

Chapter 8 concludes the thesis, and provides suggestions for future work. Following that are a number of appendices, including information about the arm design.

1.6 Note on data in thesis

This thesis includes many results, some of which are simulated and some taken from the real robot. To differentiate the source of the data, the figure captions have been formatted differently. Figures with simulated data are marked SIM in the caption, and figures with data from the real robot are marked REAL.

Chapter 2

Literature review

2.1 Introduction

This chapter reviews the relevant literature in robotics, previous work using oscillators, and evidence that humans exploit natural dynamics to perform tasks.

2.2 Exploiting natural dynamics

There is a considerable range of literature on robotics research which exploits natural dynamics. The relevant literature has been summarized in table 2.1. This review divides the work into four categories: research which exploits the robot kinematics and statics, the dynamics of the robot itself, the static characteristics of the task, and the dynamic characteristics of the task. The literature is further divided into approaches where passive mechanisms are exploited at the design stage, approaches where dynamics are used to simplify control, and, lastly, where both the design and the control are motivated by exploiting the natural dynamics. The papers themselves are discussed in more detail below.

The idea of designing passive mechanisms that achieve tasks with minimal computation has a long history. Some of the first tasks which fit in this domain were the use of compliance for robot force control, (for a review of the early work see Mason (1982)). Adding compliance makes the task of controlling force significantly easier. The Remote Centered Compliance (Drake, 1977, Whitney, 1982) is a mechanism whose static behavior greatly reduced the difficulty of common assembly tasks. The properties of the mechanism allow robust assembly without requiring any extra sensing, control or computation. Ulrich (1990) is a good example of robot design using passive mechanisms to reduce gravity torques and improve actuator performance.

The passive dynamic walking machines of McGeer (1990), with later developments by Adolfsson et al. (1998), Fowlie and Kuo (1996) and Garcia et al. (1998) are good examples of design to exploit robot dynamics. These machines are completely passive, yet produce a coordinated walking powered by gravity. Another example is the masticating robot of Takanobu et al. (1995) where nonlinear springs were used to create a realistic biting action. The open loop control of force in the sheep shearing robots of Trevelyan (1992) is also relevant. The open loop behavior of the cutter kept it a fixed distance above the skin of the sheep, so achieving the goal of only cutting the fleece of the sheep.

Robot design where the dynamics of the task are exploited include the work of Shannon, whose passive juggling robot could juggle three balls completely passively (this robot is described in Schaal and Atkeson (1993)), the gymnastic robots of Playter and Raibert (1994) and Playter (1994), which exploited the springy dynamics of the arms of a doll to maintain stability while performing flips, and the drumming robot described by Hajian et al. (1997). The drumming robot used a low impedance drumstick holder and exploited the bounce of the stick on the drum to produce high frequency drum rolls.

	Design	Control	Both
Robot Statics	Compliance: Mason (1982) Remote Centered Compliance: Drake (1977), Whitney (1982) Mechanical design: Ulrich (1990)		
Robot Dynamics	Passive Dynamic Walkers: McGeer (1990), Adolfsson et al. (1998), Garcia et al. (1998), Fowble and Kuo (1996) Mechanical design: Ulrich (1990) Chewing robot: Takanobu et al. (1995) Sheep shearing: Trevelyan (1992)	Walking robot: Pratt and Pratt (1998) Humanoid: Kuniyoshi and Na- gakubo (1997) Control: Slotine (1988)	Running robots: Raibert (1986), Playter and Raibert (1992) Internal inertia: Brown (1994) Almost passive walking: der Linde (1998) Arm tasks: THIS THESIS
Task Statics		Force/position control: Raibert and Craig (1981) Impedance control: Hogan (1985a) Constrained tasks: Niemeyer and Slotine (1997)	
Task Dynamics	Shannon juggler: Schaal and Atkeson (1993) Robot gymnastics: Playter and Raibert (1994), Playter (1994) Drumming: Hajian et al. (1997)	Juggling: Rizzi and Koditschek (1994), Schaal and Atkeson (1993), Mason and Lynch (1993)	Juggling: Lynch et al. (1998) Arm tasks: THIS THESIS

Table 2.1: Table of literature which exploits natural dynamics

The second column in table 2.1 describes work which exploits natural dynamics in the control systems for robots. The work in this area which corresponds to exploiting the robot dynamics includes the walking robot of Pratt and Pratt (1998) which exploits passive leg dynamics to enable the use of a simple controller for walking, and the work of Kuniyoshi and Nagakubo (1997) which exploits the dynamic behavior of a simulated humanoid robot in the task of standing up from a sitting position. Slotine (1988) argued for the use of physical concepts such as energy in designing control systems

Work in this area which exploits the static aspects of tasks include hybrid force/position control Raibert and Craig (1981), where position and force constraints were used to modify the robot motion in orthogonal directions. Impedance control, Hogan (1985a), generalizes this approach by manipulating the impedance of the robot endpoint. A more recent example of exploiting constraints is the work of Niemeyer and Slotine (1997) where the constraint is used to determine the motion, so that the robot follows the direction of least resistance.

Task dynamics have been exploited for control by a number of authors, a common application being juggling a ball on a paddle. Rizzi and Koditschek (1994) used controllers that moved the robot paddle in a scaled mirror image of the ball motion. Schaal and Atkeson (1993) presented a variety of schemes which exploited the dynamics of the juggling task itself to achieve juggling with very simple controllers. For example, juggling of one ball could be achieved using a paddle driven with a constant frequency sine wave. The ball natural behavior is to entrain with the paddle motion, giving a stable juggling action with no explicit control. This contrasts with the control scheme of Rizzi and Koditschek (1994), where the paddle position was carefully controlled dependent on the ball state. Mason and Lynch (1993) described another juggling application, using a hand-coded controller to control the throw of a juggling club.

The final column of table 2.1 concerns approaches where the natural dynamics are exploited both in terms of design and control. A good example of this approach in terms of robot design are the running robots of Raibert (1986). Raibert used hydraulic actuators which gave his running robots a springy bouncing behavior, and then used three simple controllers to modify the speed, height and pitch stability of the bouncing motion. This approach yielded robots which could run fast as well as do gymnastic maneuvers (Playter and Raibert, 1992). Another example is the work of Brown (1994), who designed a robot to exploits its internal inertia for tasks such as tugging on a stuck door and weight lifting. The walking robot of der Linde (1998) augments a passive dynamic walker with small amounts of control to achieve an energy efficient walking robot.

An example of exploiting task dynamics is the robotic juggling robot presented by Lynch et al. (1998). This robot used a carefully designed hand and controller to manipulate the ball.

The work presented in this thesis fits into the final column of table 2.1. The arm used is specifically designed to be compliant and have spring-like dynamics. These dynamics are then exploited by the oscillator control. The work fits into two rows of the table: the robot dynamics are exploited to create efficient movement, and the task dynamics are exploited to coordinate the various joints of the arm with the task itself. The work extends previous research both in terms of the complexity and variety of applications demonstrated.

2.3 Behavior-based robotics

The work in this thesis is motivated by research on behavior-based robotics (Brooks, 1986). This has enjoyed considerable success in the mobile robotics community for tasks such as obstacle avoidance and navigation. The ideas have not been well accepted in the manipulation community and there are just a few examples of the ideas being applied to robot manipulation. For example, Smithers and Malcolm (1988) used reactive controllers with a traditional planner for assembly tasks.

The main reason for the lack of transfer is the complex kinematics of robot arms and tasks as compared to mobile robots. There is not the same obvious mapping from sensors to actions

(e.g. if the arm hits something it is not obvious how to move away from the obstacle). This has lead to solutions where the motion of the joints was directly mapped to the task (e.g. Asteroth et al. (1992)), very simple arms are used (Connell, 1988), or simple combinations of joints are used (Williamson, 1996, Marjanović et al., 1996). Connell (1994) decomposed the problem into a number of sensor-driven primitives which he combined to pick objects from cluttered work areas. Gershon (1990) described the decomposition of a sewing task into independent modules which were combined to robustly respond to the unpredictable behavior of the fabric. Combining behaviors is also more complicated for arms than for mobile robots. Beccari and Stramigioli (1998) suggested using impedance control to combine behaviors for assembly tasks.

Most behavior-based approaches build successively more complex behavior by combining simple, reactive controllers. The work in this thesis shows how a range of behaviors can be obtained with a single reactive controller. However, some characteristics of the oscillator control are analogous to behavior-based methods. The oscillator is tightly coupled to the environment, and relies on interaction with the environment to create complex behavior. The system also produces complex behavior without requiring explicit models of the controller or the environment.

2.4 Previous work using oscillators

A table summarizing the previous work using oscillators is included in table 2.2. The table compares the literature on a number of axes: whether the papers describe learning, the type of oscillator used, the task described, whether networks of oscillators are used, whether feedback is used, and whether the paper describes a simulated or real implementation. The papers are roughly ordered and grouped in terms of task complexity.

The first group of papers describe learning algorithms to set the connection strengths between groups of oscillators. Doya and Yoshizawa (1989) describes back-propagation for recurrent neural networks, Ermentrout and Kopell (1994) give a method to tune connection strengths in models of biologically plausible neurons, and Nishii (1998) provides methods for tuning connection strengths for oscillators described by phase dynamics. The task in all these papers was to produce a desired frequency and phase output from an oscillatory network. The work in this thesis for the most part exploits the natural dynamics to coordinate multiple oscillators and, so, does not use learning of connection strengths.

The next group of papers exploit the dynamics of an oscillator network to generate gait patterns similar to those observed in quadrapedal locomotion. The idea behind using an oscillator network for this is to produce a variety of different gaits under the control of one parameter. The network provides the coordination of the limbs, and the overall characteristics are controlled by other parameters. An ideal implementation would respond to an increase in frequency by changing from walking to trotting and then to a galloping rhythm. Kimura et al. (1993) presented an oscillator model for insect walking gaits, and Pribe et al. (1997) presented a different oscillator model (Elias-Grossberg) which also had gait changing properties. Collins and Richmond (1994) argued that the details of the network connectivity, rather than the detailed dynamics of the individual oscillators are important to produce patterns and transitions. Zielińska (1996) used a Van der Pol oscillator to generate the rhythm for a two joint leg, and Cohen et al. (1982) showed how coupling between simple oscillators could be used to generate an undulatory locomotion pattern. None of these papers take the dynamics of the task into account, which differs from the work in this thesis. The thesis work requires the natural dynamics in order to coordinate multiple degrees of freedom, and uses that dynamics to remove the need for explicit networks of oscillators.

The next two papers considered address the oscillator property of resonance tuning which is automatically driving a mechanical system at its resonant frequency. An oscillator system with this property was presented by Hatsopoulos et al. (1992), who subsequently used this task to argue for the use of oscillators coupled to systems (Hatsopoulos, 1996). This thesis further demonstrates the

Paper	Learning	Type of oscillator	Task	Connections	Feedback	Real
Doya and Yoshizawa (1989)	✓	S	Matching phases	✓		
Ermentrout and Kopell (1994)	✓	C	Matching phases	✓		
Nishii (1998)	✓	S	Matching phases	✓		
Kimura et al. (1993)		C	Gait pattern	✓	✓	
Pribe et al. (1997)		C	Gait pattern	✓		
Collins and Richmond (1994)		VDP, C	Gait pattern	✓		
Zielińska (1996)		VDP	Gait pattern	✓	✓	
Cohen et al. (1982)		S	Undulatory pattern	✓	✓	
Hatsopoulos et al. (1992)		C	Resonance tuning	Single DOF	✓	
Hatsopoulos (1996)		VDP	Resonance tuning	Single DOF	✓	
Doya and Yoshizawa (1992)	✓	C	Simple locomotion	✓	✓	
Nishii (1995)	✓	S	1 DOF hopping	Single DOF	✓	
Wadden and Ekeberg (1998)		C	Single leg	✓	✓	
Taga et al. (1991)		M	Bipedal walking	✓	✓	
Taga (1995a,b)		M	Bipedal walking	✓	✓	
Miyakoshi et al. (1998)		M	3D Bipedal walking	✓	✓	
Kimura et al. (1998)		M	Quadrupedal walking	✓	✓	✓
Ekeberg (1993)		C	Undulatory locomotion	✓	✓	
Ijspeert et al. (1998)	GA	C	Undulatory locomotion	✓	✓	
Hollerbach (1981)		S	Handwriting			
Miyakoshi et al. (1994)		M	Juggling	Single DOF	✓	
Schaal and Sternad (1998)		M	Arm control	✓	✓	
THIS THESIS		M	Arm control		✓	✓

Table 2.2: Table of literature of previous work using oscillators. The type of oscillator is indicated with S, meaning simple phase dynamics, VDP meaning Van der Pol oscillator, M meaning Matsuoka oscillator, and C meaning a different oscillator with complex dynamics.

power of coupling between oscillators and mechanical systems and extends their application to a broad range of tasks.

A large proportion of the oscillator literature is devoted to locomotion partly because it is an obvious rhythmic task and partly because of the biological origins of central pattern generators for locomotion (see section 2.6). Nearly all the work in this field is simulated, with rare examples of implementations on real robots. The next group of papers give a number of applications in locomotion.

There is a considerable range of applications, perhaps the most simple being the rolling example given by Doya and Yoshizawa (1992). An oscillator was tuned to move a mass on a disk, so causing it to roll.¹ Single degree of freedom control of a hopping robot was described by Nishii (1995), and the control of a single leg was described by Wadden and Ekeberg (1998). A task with more reasonable complexity was addressed by Taga et al. (1991). Taga presented a bipedal robot which walked using a network of neural oscillators to control the joints of the robot. The oscillators used feedback from the biped joint angles to entrain the oscillators with the motion. The biped was robust to perturbations and could walk up inclines due to this coupled oscillator-system dynamics. Later papers have added impedance control and finite state machines to control the walking (Taga, 1995a,b), and the latest papers describe obstacle avoidance using the dynamics of the oscillator system (Taga, 1998), and an application to 3D walking (Miyakoshi et al., 1998). These papers are responsible for the popularity of the Matsuoka oscillator used throughout this thesis and have inspired at least one implementation on a real robot (Kimura et al., 1998).

Undulatory locomotion is another application domain, inspired by work in the neuroscience and mathematical community on the lamprey (e.g. Cohen et al. (1982)). This primitive fish appears to use a set of coupled oscillators to produce its undulatory motion. Ekeberg (1993) described a computer model of lamprey swimming which included turning and noted the use of feedback to the oscillators to overcome perturbations. Ijspeert et al. (1998) used genetic algorithms to evolve a similar creature, although he also added legs to produce a creature that could both walk and swim.

As with previous work, the emphasis in these approaches is the connections between oscillators to produce the required coordination patterns for locomotion. Feedback has a minor and poorly understood role making the systems more robust to system perturbations. The work in this thesis has the opposite emphasis. Feedback is taken to be most important, being used to provide coordination and robustness, and connections are shown to be more difficult to tune, and not able to provide motion which adapts to the dynamics of the task at hand.

There are very few implementations of oscillators for arm control. Hollerbach (1981) described using simple oscillators to model and generate handwriting, and Miyakoshi et al. (1994) presented a simulated juggling task. Schaal and Sternad (1998) suggested a general framework of using oscillators to generate rhythmic movements of arms. The work presented in this thesis extends the use of oscillators to considerably more complex tasks than described in the literature. It also differs from most of the literature in that it is implemented on a real robot.

2.5 Oscillator analysis

Given the range of literature on oscillators, it is surprising that there is not more practical analysis of oscillator behavior. The work of Matsuoka (1985, 1987) examined stability of networks of neurons, but did not address the effect of feedback. The effect of inputs to neurons is often described using phase-response curves (e.g. Canavier et al. (1997), Dror et al. (1999)), although these appear to be appropriate for oscillators which are accurate models of biological neurons, producing spikes rather than sinusoidal-like outputs. Other work on neurons is more mathematical (for example Terman

¹This task was presented as a learning problem although simply connecting a Matsuoka oscillator to the same system resulted in rolling with minimal tuning. The reason for this may be the difference in the dynamic richness of the Matsuoka oscillator and the oscillator used by Doya and Yoshizawa (1992).

et al. (1998) and other work by Kopell) and has the aim of understanding the dynamics of biological neurons. The analysis presented in this thesis provides accurate, simple analysis which is geared towards understanding the behavior of oscillators coupled to systems, as well as providing practical information for tuning purposes.

2.6 Biological central pattern generator evidence

There is considerable evidence for oscillator or central pattern generator control of human and animal limbs. The most striking evidence comes from observing the behavior of the hind legs of spinalized cats, when placed on a treadmill. The action of the treadmill on the legs causes the legs to walk, even though there is no direct spinal control (see Rossignol and Dubuc (1994), Rossignol (1996) for reviews). The legs coordinate with each other, and produce a fairly normal looking stepping pattern. There is also evidence that humans have a similar behavior (see Barbeau and Rossignol (1994) for a review).

Another thread of research is the physiology of invertebrates, where pattern generators are thought to be responsible for locomotion of insects, lampreys etc.. Arshavsky et al. (1991) and Getting (1988) provide reviews of the literature in this field, and Pearson (1993) relates the invertebrate and vertebrate literature.

2.7 Human arm control

There is considerable evidence that humans exploit the dynamics of their bodies in the way that we perform tasks, move, and perceive objects.

For example, we organize our kinematics in many different ways to be appropriate for different tasks. When playing pool, our whole bodies are arranged to put the cue in line of sight of our eyes, and the mechanical control is directed from one isolated joint. When writing or performing a delicate task, we often brace our hands on a hard surface, isolating and removing excess degrees of freedom, and aligning those that are left with the axes of the task. Another example is using our skeleton rather than our muscles to carry loads. Hogan (1985b) described the effect of postural changes and musculature on the mechanical properties of arms.

There is evidence that human exploit the compliant spring-like dynamics of our limbs in a variety of ways. The compliance itself appears to be passive (Mussa-Ivaldi et al., 1985) and thus appropriate for interacting with objects in a stable manner. The spring-like properties of our muscles and tendons are exploited during running (Alexander, 1990). The resonant properties of our limbs are also exploited during ordinary movement, Herr (1993) and Hatsopoulos and Warren (1996) showing that the preferred frequency of swinging arms is the resonant frequency of the limb and its environment, determined by measuring the most “comfortable” swinging frequency when the arm dynamics had been augmented by either extra springs or added weights. In addition Herr (1993) provided experimental data that the comfortable frequency corresponded to a minimum of work. Even babies seem to be able to respond to the resonant properties of the environment, quickly finding the resonant frequency of bouncing supports (Goldfield et al., 1993). Bingham et al. (1989) showed that when throwing objects a sequence of postures is used which exploits the dynamical coupling between the limb segments. He also showed that the mass properties of the object are estimated in order to tune muscle stiffnesses for a powerful throw. The idea that when the arm performs a task it is a “smart” mechanism, adjusting its passive properties for the task has been suggested by Saltzman and Kelso (1987) and Bingham et al. (1991) amongst others.

There is some evidence from development (Thelen et al., 1992) and also from motor learning (Schneider et al., 1989) that the goal of learning is to correct the timing and strength of the muscle forces to complement the natural dynamics. Thus the interaction forces between the limb segments are harnessed to produce the movement, or in the words of Bernstein:

...the secret of co-ordination lies not only in not wasting superfluous force in extinguishing reactive phenomena but, on the contrary, in employing the latter in such a way as to employ active muscle forces only in the capacity of complementary force. (Bernstein, 1967, p. 109)

The interplay between explicit control and natural dynamics is a subject of debate. If the natural dynamics were the most important aspect of arm control, one might expect the same task to be carried out differently in different parts of the workspace because of the different mechanisms used. This does not seem to be the case, there being a remarkable similarity between arm motions in different parts of the workspace, and tasks performed at different scales. For example, signatures tend to look the same when written very small, very large or even when written with other body parts such as one's foot. Humans also tend to produce roughly straight movements over the workspace of the arm (Morasso, 1981, Atkeson and Hollerbach, 1985). Although some researchers argue that many of the qualities of arm motion can be explained either by complex muscle models (Gribble et al., 1998), or neural network models (Massone and Myres, 1996), the similarity of movements suggests that a fairly complex mechanism is controlling the arm to compensate for the dynamics of the arm itself.

A complex mechanism is required because the dynamics of human arms are considerable, their stiffness being low and the stiffness decreasing during movement (Zajac, 1989, Bennett et al., 1992). Experiments where the arm dynamics are perturbed (e.g. Gandolfo et al. (1996)) or where the perception of the arm is perturbed (e.g. Wolpert et al. (1995)) both show humans rapidly adjusting their movement towards straight line motion. This further suggests active control of the arm. Whether this control is an explicit controller or whether it is something more aligned with Bernstein ideas is also a subject of debate.

One experiment which illustrates the trade off between control and exploiting dynamics is the study on obstacle avoidance by Sabes and Jordan (1997). Sabes found that the closest distance between the object and the arm correlated well with the arm's inertial properties at that point. The arm was controlled to avoid the obstacle, but the "safe distance" was determined by the arm dynamic properties.

The dynamics of objects also appear to be important in perception. Many perceptual properties can be explained by humans sensing the inertia of objects (Turvey and Carello, 1995). For example, humans are good at estimating the lengths of objects they manipulate, properties that correlate well with the objects inertia. (Pagano and Turvey, 1995) suggested that this mechanism is used to determine limb orientation. Weights were added to the arm which caused subjects to point not with the axis of their arm but with the inertial axis of their arm and the added weights.

2.8 Conclusion

This chapter has situated the work in this thesis in the relevant literature. The approach taken exploits the natural dynamics of the robot and the task, and does so both in terms of the robot design, but also in the choice of control. The work extends previous work on oscillators by addressing a wide range of complex tasks, and by being implemented on a real robot. The emphasis on feedback rather than networks of oscillators also differentiates this work from most oscillator research. The analysis tools presented in this thesis also extend the available analysis techniques.

The chapter has also shown that using oscillators has a firm biological basis, and that exploiting natural dynamics is an approach which humans take when performing particular tasks, and may also be important from general motion.

Chapter 3

Single degree of freedom motions

This chapter introduces the oscillators which are used throughout this thesis. It describes their governing equations, and how they are coupled to the joints of the arm. The chapter then introduces an analysis technique using describing functions which is useful for understanding the non-linear behavior of the oscillators. This analysis technique is accurate, and can be used in a number of ways. The analysis makes clear the effect of parameters on the overall system behavior, which facilitates tuning of parameters. It also shows the inherent robustness of the oscillator properties to parameter and system changes.

The analysis is not restricted to the Matsuoka oscillator and can be equally applied to other oscillator types. The chapter shows its application to the well known Van der Pol oscillator. Neither is the analysis restricted to simple tasks, the chapter describing its application to the design of oscillator driven juggling. This is implemented on the robot.

The chapter concludes by stressing the coordination between the oscillator dynamics and the arm dynamics, and the sensitivity of the oscillators to the exact motion of the arms.

3.1 Introduction

The oscillator used throughout this thesis was originally analyzed by Matsuoka (1985, 1987). It consists of two simulated neurons in mutual inhibition as shown in figure 3-1. The oscillator model approximates the envelope of the firing rate of a real biological neuron with self-inhibition. The model has 4 state variables, governed by the following equations:

$$\tau_1 \dot{x}_1 = c - x_1 - \beta v_1 - \gamma [x_2]^+ - \Sigma_j h_j [g_j]^+ \quad (3.1)$$

$$\tau_2 \dot{v}_1 = [x_1]^+ - v_1 \quad (3.2)$$

$$\tau_1 \dot{x}_2 = c - x_2 - \beta v_2 - \gamma [x_1]^+ - \Sigma_j h_j [g_j]^- \quad (3.3)$$

$$\tau_2 \dot{v}_2 = [x_2]^+ - v_2 \quad (3.4)$$

$$y_i = [x_i]^+ = \max(x_i, 0) \quad (3.5)$$

$$y_{out} = [x_1]^+ - [x_2]^+ = y_1 - y_2 \quad (3.6)$$

Each simulated neuron is governed by two equations, neuron 1 by (3.1) and (3.2) and neuron 2 by (3.3) and (3.4). The variables x_1, x_2 represent the firing rate of the neurons, and the v_1, v_2 variables are internal states representing the self-inhibition. The outputs of the neurons are taken to be the positive parts of the firing rates, with $y_1 = [x_1]^+$ and $y_2 = [x_2]^+$. The output of the oscillator is taken to be the difference of these signals (3.6). The neurons inhibit one another, with the output of the other neuron appearing in the update for the firing rate, as in the term $-\gamma [x_2]^+$ in (3.1).

The parameters of the oscillator are β, γ which are constant,¹ the constant or tonic parameter c , which determines the amplitude of the oscillator output (defined as A_n), and the two time constants

¹The values used throughout the thesis were $\beta = 2, \gamma = 2$. Working parameter ranges are described by Matsuoka (1985)

τ_1, τ_2 which determine the frequency and shape of the output. For stable oscillations τ_1/τ_2 should be in the range 0.1–0.5.² Keeping the ratio τ_1/τ_2 constant makes the natural frequency of the oscillator ω_n (the frequency of the oscillator without an input) proportional to $1/\tau_1$, as shown in figure 3-2. Figure 3-3 shows a typical output from the oscillator.

Inputs are applied to the oscillator through the variables g_j , weighted by gains h_j . The inputs are arranged to always inhibit the neuron, applying the positive part $[g_j]^+$ to one neuron, and the negative part $[g_j]^- = \max(-g_j, 0)$ to the other. To remove offsets in g_j a high pass filter is used to remove the DC component.

When no input is applied to the oscillator, it oscillates at a natural frequency w_n determined by the time constants τ_1, τ_2 , with a fixed amplitude defined by the tonic c , as shown in figure 3-2. If an oscillatory input is applied, the oscillator entrains the input, producing an output at the same frequency as the input. This entrainment behavior is illustrated in figure 3-4 which shows the output of the oscillator as the size of the input signal is increased. The oscillator tends to entrain very quickly, within one cycle in the lower graph in figure 3-4. The oscillator can lock onto input frequencies over a wide range of frequencies and sizes of inputs as illustrated in Figure 3-5. The figure shows the minimum input required to frequency lock the oscillator as a function of frequency. The plot was obtained by varying the input magnitude and comparing the oscillator frequency (taken as the frequency with the maximum magnitude in a Fourier transform of the output), with the input frequency. The entrainment range is large, in this case $w_n = 7$ rad/s, and the range is 1.5 to 35 rad/s.

For most of the work in this thesis, the oscillator is connected to the joints of the robot arm. As described in chapter 1, the position of each joint is controlled using a simple proportional-derivative control law, making the commanded torque at the i th joint

$$u_{di} = k_i(\theta_{vi} - \theta_i) - b_i\dot{\theta}_i \quad (3.7)$$

where k_i is the stiffness of the joint, b_i the damping, θ_i the joint angle and θ_{vi} the equilibrium point. The oscillator is connected to the arm using the oscillator output to control the setpoint for the joint with an output gain h_o and offset θ_{pi}

$$\theta_{vi} = h_o y_{out} + \theta_{pi} \quad (3.8)$$

and connecting either the joint angle θ_i or the joint torque u_i to the input g of the oscillator. The offsets in these signals are removed using a high pass filter. Coupling the oscillator in this way results in a number of useful properties. These are described in the following sections.

3.2 Analysis methodology

The oscillator is a non-linear system, and its behavior when coupled to the mass-spring system of a robot joint is complex. The final motion depends on the interaction between the oscillator dynamics and the system dynamics. This section describes an analysis method that can be used to understand the nature of this interaction, and shed light on the oscillator behavior, stability, tuning and robustness. The analysis is based on expressing the oscillator and the driven system as frequency responses, and determining conditions for oscillation. This technique is known as describing function analysis (Gelb and Vander Velde, 1968, Slotine and Li, 1991, Khalil, 1996).

The coupled system is illustrated in figure 3-6, where the driven system or plant is assumed to be linear with transfer function $G(j\omega)$, where $j = \sqrt{-1}$ and ω is the frequency. The oscillator is non-linear and so does not have a frequency response. An approximation to the frequency response can be made by sampling over a range of frequencies and input amplitudes. Writing this linearized

²The ratio used throughout the thesis was $\tau_1/\tau_2 = 0.5$.




Figure 3-1: Schematic of the oscillator. The oscillator equations simulate two neurons in mutual inhibition as shown here. Black circles correspond to inhibitory connections, open to excitatory. The mutual inhibition is through the $\gamma[x_i]^+$ connections, and the βv_i connections correspond to self-inhibition. The input g_j is weighted by a gain h_j , and then split into positive and negative parts. The positive part inhibits neuron 1, and the negative part neuron 2. The output of each part neuron y_i is taken to be the positive part of the firing rate x_i , and the output of the oscillator as a whole is the difference of the two outputs.




Figure 3-2: SIM Oscillator behavior under changing tonic c and time constant τ_1 . The left hand figure shows the output frequency of the oscillator w_n plotted against $1/\tau_1$, and the right hand graph shows the amplitude of the oscillator output A_n plotted against the tonic excitation c . For this example $\tau_1/\tau_2 = 0.5, \beta = 2.5, \gamma = 2.5$.




Figure 3-3: SIM A sample output from the oscillator. The top graph shows the variation with time of the states of the oscillator x_1, x_2, v_1, v_2 during normal operation. The bottom graph shows the output of the oscillator $y_{out} = [x_1]^+ - [x_2]^+$. For this example $\tau_1 = 0.25, \tau_2 = 0.5, c = 1.5, \beta = 2.5, \gamma = 2.5$.




Figure 3-4: SIM Figure showing effect of increasing input signal. For small input (top graph), the oscillator is not entrained, and oscillates at its endogenous frequency w_n . In the middle graph, the input is larger, and the oscillator is almost entrained, but slips every couple of cycles. The lower graph shows the oscillator locked onto the frequency of the input. For this example $c = 1.0, \tau_1 = 0.1, \tau_2 = 0.2$.




Figure 3-5: SIM Minimum input for entrainment. The figure shows the minimum input signal required for entrainment of the oscillator. The oscillator natural amplitude and frequency are given by the horizontal and vertical lines respectively. For this example $c = 1.0, \tau_1 = 0.1, \tau_2 = 0.2$.

response at frequency ω and input amplitude A as $N(j\omega, A)$, then the condition for steady state oscillation is that the loop gain is unity:

$$N(j\omega, A)G(j\omega) = 1 \quad (3.9)$$

Calculating N and solving this equation for the frequency and amplitude of the limit cycle solution (ω_f, A_f) provides much insight into the capabilities and tuning of oscillator driven systems. The following section describes how to calculate N for a particular choice of oscillator, and how to solve this equation to calculate the final solution.

3.3 Oscillator describing function

The oscillator describing function $N(j\omega, A)$ expresses the gain and phase difference between input and output of the oscillator as a function of input amplitude and frequency. This can be calculated by applying an input $g_j = A \sin(\omega t)$, and measuring the output y_{out} . If the oscillator is entrained, the frequencies of input and output are the same and a Fourier transform can be used to calculate an approximation to the gain and phase of the oscillator. By applying inputs over a range of frequencies and amplitudes and removing points where the oscillator is not entrained, N can be calculated.³

Figure 3-7 shows N plotted as a Bode plot. The plots show the variation of gain and phase with frequency. The multiple lines correspond to different values of A , the arrow referring to the direction of increasing A . Changes in amplitude affect the gain (gain decreasing with increasing A), but not the phase, which remains roughly independent of A . For this oscillator, the gain is actually proportional to $1/A$, as the output amplitude is approximately constant. This is because the input is arranged to always inhibit the neurons. If the input is applied without the max operator, the output

³It is also possible to calculate this analytically, using some advanced techniques in describing function analysis, taken from Gelb and Vander Velde (1968).




Figure 3-6: System schematic. The figure shows an oscillator tightly coupled to a system. The oscillator output y_{out} drives the input of the system, whose output g is the oscillator input. The analysis of the system proceeds by looking at the frequency responses of each system and calculating the condition for oscillation. At a frequency ω , and input amplitude A , the linearized frequency response of the oscillator can be written as $N(j\omega, A)$, where $j = +\sqrt{-1}$. The frequency response of the system is $G(j\omega)$. There will only be stable oscillations when the loop gain is unity, or $N(j\omega, A)G(j\omega) = 1$.

amplitude becomes a function of A , and the range of possible gains is reduced (see Appendix C for more details).

A common situation is the oscillator driving a single link of the arm. In this case the plant or driven system can be modeled as a mass m actuated through a spring of stiffness k , and a damper b (where k and b are the stiffness and damping terms in the joint level control, equation (3.7)). Ignoring the actuator dynamics, the dynamics of the link are given by the equation

$$m\ddot{\theta} + b\dot{\theta} + k\theta = k\theta_v \quad (3.10)$$

where θ is the position of the link, and θ_v is the desired position. Since the link is connected in the usual way to the oscillator, with $y_{out} = \theta_v$, and $g_j = \theta$, the transfer function for the plant is between θ_v and θ :

$$G(j\omega) = k/(k - m\omega^2 + j\omega b) \quad (3.11)$$

The condition for limit cycles (equation (3.9)) can now be solved graphically by plotting $G(j\omega)$ and $1/N(j\omega, A)$ on the complex plane and finding points where the graphs intersect at the same frequency. This plot is shown in figure 3-8. There are many intersection points, but there is only one where the frequencies of both transfer functions are the same. Since the exact solution is unlikely to be at the points sampled when calculating N and G , a simple interpolation routine is used to find the limit cycle solution. The intersection point (marked with a $\bullet$), determines both the frequency and amplitude of the limit cycle solution. In this case $\omega_f = 8.9$ and $A_f = 0.52$.

3.4 Predictive accuracy

Describing function analysis is an approximate method, and relies on the assumption that the system is well described by the lowest frequency harmonic (Slotine and Li, 1991, Khalil, 1996). This requires firstly that the output of the oscillator not contain many higher harmonics, and secondly that the gain of the linear system be much lower at these harmonics, i.e. that the linear system be a low pass filter. The Fourier transform of the oscillator output is shown in figure 3-9. Most of the energy is located around the natural frequency of the oscillator. Since the mass-spring system has a low pass characteristic, it is appropriate to use the describing function analysis in this case.

The accuracy of the prediction can be found using simulation. Figure 3-10 shows the predicted and measured (from a simulation of the system) frequency and amplitude as the stiffness of the spring




Figure 3-7: **[SIM]** Oscillator Bode plot. The top graph shows the gain $|N(j\omega, A)|$ plotted against frequency, and the lower plot the phase. The multiple lines correspond to different values of input amplitude A as indicated by the numbers and arrows on the plot. The gain is inversely related to A and is roughly constant with frequency. The phase is less dependent on A , and reduces from around 180° to 60° as the frequency is increased. The natural frequency of the oscillator is indicated by the vertical line.




Figure 3-8: **[SIM]** Plot of $G(j\omega)$ and $1/N(j\omega, A)$ in the complex plane. $G(j\omega)$ is the thick line, with numbers indicating frequency. The lines for $1/N(j\omega, A)$ correspond to constant frequency, with amplitude A increasing along each line as indicated by the numbers and arrow. There is a limit cycle solution at the point where the two graphs intersect at the same frequency. The intersection point is indicated by a filled circle, at $\omega_f = 8.9$ and $A_f = 0.52$. This plot was produced for $\tau_1 = 0.1, \tau_2 = 0.2, c = 1, h_j = 1, k = 20, m = 0.4, b = 2$.




Figure 3-9: SIM Power density spectrum for the Matsuoka oscillator. The plot shows power plotted against frequency, expressed as a function of the oscillator natural frequency ω_n . The plot shows a clear peak at the natural frequency, and a much smaller one at $\omega = 3\omega_n$. The sinusoidal nature of the oscillator output confirms the accuracy of the describing function approximation.

is altered. Changing stiffness alters the shape of the $G(j\omega)$ plot, and so alters the final solution found by the oscillator. The figure shows that the prediction is accurate, with possible errors coming from the calculation of $N(j\omega, A)$, as well as the estimation of the frequency and amplitude of the mass motion.

3.5 Local stability

The plot of $G(j\omega)$ and $1/N(j\omega, A)$ can also be used to determine the stability of the final motion, using an extended version of the Nyquist Criterion. The criterion is taken from (Slotine and Li, 1991, p. 186):

Limit Cycle Criterion: *Each intersection of the curve $G(j\omega)$ and the curve $1/N(A)$ corresponds to a limit cycle. If points near the intersection and along the increasing- A side of the curve $1/N(A)$ are not encircled by the curve $G(j\omega)$, then the corresponding limit cycle is stable. Otherwise the limit cycle is unstable.*

In the case of the Matsuoka oscillator driving a mass, examining figure 3-8 shows that for low A , the points are encircled by $G(j\omega)$, but, as A increases, the points move into an un-encircled region. This corresponds to locally stable behavior. Since the analysis only considers the steady state solutions and not the transient behavior, it cannot predict global stability. Experimentally, the cycles are also observed to be stable. Figure 3-11 shows a phase plot for the motion of a robot arm link, for a number of different initial conditions. The system converges to the limit cycle very rapidly, within one cycle in this case.




Figure 3-10: SIM Predictive accuracy. The plot shows the prediction (line) and measured (*) frequency (top plot) and amplitude (lower plot) for an oscillator driving a mass, as the stiffness of the spring was varied, so varying the natural frequency of the system. The accuracy is good, with a constant offset error for amplitude. The error in frequency is partly due to noise in the frequency measurement of the simulated mass motion.




Figure 3-11: REAL Phase plot of robot motion. The plot shows velocity v. position for an oscillator driving a mass, recorded from one link of a robot arm. Shown are three traces overlaid starting from three different initial conditions. The three traces all converge to the limit cycle, indicative of the anecdotal observation that all initial conditions converge to the limit cycle.

3.6 Effect of parameter changes (a): Design

The oscillator has a number of parameters, time constants, input gains, tonic excitations etc. which affect the final motion of the coupled system. The describing function analysis sheds light on the effect of changes in those parameters, by looking at the effect on the plots $G(j\omega)$ and $N(j\omega, A)$. The effect of most parameters is to scale or shift N while preserving its shape. This means that design can be accomplished without requiring the laborious recalculation of N .

Frequency and amplitude parameters. Altering the natural frequency of the oscillator does not alter the shape of N , but simply scales the frequency of the plot. The shape of the plot is the same if the frequencies are measured with respect to the oscillator natural frequency ω_n . Since for the Matsuoka oscillator $\omega_n \propto 1/\tau_1$, changes in time constants can be easily accounted for.

Similarly, the amplitude of the oscillator is proportional to the parameter c . Changing c increases the amplitude of the oscillator, but does not change N , except that larger inputs are needed to entrain the system. The shape of the plot remains the same when the input is scaled by the oscillator natural amplitude A_n .

Input and output gains. In general, the oscillator input is multiplied by a gain h_j , and the output also has an output gain h_o , as shown in figure 3-12. The effect of these gains is to change $N(j\omega, A)$ to $h_j h_o N(j\omega, h_j g_j)$, which alters the magnitude but not the phase. Changing the gains thus changes the overall size of N , and can be used to move the plots over one another, or to change the limit cycle solution. Rather than recalculating N , the effect of the gains can be lumped with G , i.e. $G'(j\omega) = h_j h_o G(j\omega)$, which is easily recalculated.




Figure 3-12: If N has an input gain h_j , and an output gain h_o , the overall gain of the oscillator becomes $h_j h_o N(j\omega, A)$. Since N depends on A , i.e. amplitude after the gain h_j , it makes sense to combine the gains with the plant transfer function G , i.e. $G'(j\omega) = h_j h_o G(j\omega)$.

Input and output variables. Choosing a particular oscillator type fixes $N(j\omega, A)$, while choosing outputs and inputs fixes $G(j\omega)$. As shown above, changing oscillator parameters can scale the size of N , but not alter its phase. Thus G needs to be chosen so that the two plots can intersect (there are phases that match). It is then relatively easy to find gains to place the plots over one another. This aspect of tuning is also discussed in appendix B.

Taking for example the mass-spring system of equation (3.10). If instead of position θ , the velocity of the mass $\dot{\theta}$ is used as input, the frequency response becomes

$$G'(j\omega) = j\omega G(j\omega) = j\omega k / (k - m\omega^2 + j\omega b) \quad (3.12)$$

which is the same as before, only rotated 90° counterclockwise, and scaled by ω .

3.7 Effect of parameter changes (b): Robustness

The describing function analysis can also be used to examine the robustness of the oscillators. Robustness in this context means either that the coupled system has the same performance or behavior over a wide range of oscillator parameters, or that given a single set of parameters, a wide range of system properties give the same behavior. The oscillators are robust in both senses of the word. The oscillator parameters are generally well-behaved, with a smooth monotonic effect on the system behavior, decoupled from the effects of the other parameters.

Robust properties greatly ease the implementation of oscillator driven systems since very little tuning is required to achieve good results. This section evaluates each of the oscillator parameters in turn.

Input gains and tonics. The tonic parameter c affects the output amplitude of the oscillator and is effectively decoupled from all the other parameters. The output amplitude is independent of input size and frequency (as shown in figure 3-7). This independence is a consequence of the non-linearities on the input g in equations (3.1) and (3.3). Appendix C shows that without these non-linear terms the system output is a function of input size, and is less well-behaved. The effect of the tonic parameter is monotonic, as was shown in figure 3-2.

The input parameter h is also decoupled from the oscillator parameters which set frequency. It has a more interesting effect on the overall system properties. In practice, the parameter does not have much effect on the final motion. It needs to be large enough to cause entrainment, but other than that its value is unimportant. The reason for this can be seen in figure 3-13, which plots the describing function for the oscillator with two different values of input gain. The two different solutions are $\omega_f = 8.82, A_f = 0.38$ for $h = 0.7$, and $\omega_f = 8.98, A_f = 0.77$, for $h = 1.7$. In spite of the gain more than doubling, the frequency changes by less than 2%. Since A_f is the amplitude of input to the oscillator, the actual motion has amplitude A_f/h , which is 0.55 when $h = 0.7$, and 0.45 when $h = 1.7$, a change of about 15%. Looking at figure 3-13, the reason for this robustness is the approximately radial lines of the oscillator describing function, particularly in the middle frequencies. The oscillator is inherently robust to changes in input size, giving approximately the same performance for a wide variety of parameter values.

For a fixed set of gains, the range of phases, frequencies, and gains that the oscillator can produce is large, so that any part of the system phase plot which overlaps with the area covered by the oscillator will produce an oscillatory solution. This contributes to the robustness of the oscillator itself.

Time constants. The main effect of changing the time constants of the oscillator is to change the frequency range of the oscillator. This effect is monotonic, as shown clearly in figure 3-2. As mentioned above, the effect is decoupled from the input gain and tonic parameters.

Referring to figure 3-13, if the lines were radial, then the effect of changing the time constants would be a monotonic change in the frequency of the coupled system. The lines are not exactly radial, which gives rise to the resonance tuning behavior of the system. The oscillator can tune into the resonant frequency of a driven system for a wide range of oscillator parameters, and conversely, for a fixed set of oscillator parameters, can tune into the frequency of a range of systems.

Figure 3-14 shows the range of settings of input gain and time constant which have a final solution which is within 10% of the system natural frequency. The plot shows that the oscillator natural frequency can vary from 3.5–5.5 rad/s which is a percentage change of 40 %, and the gain can vary even more (from 0.2 to 2), while still finding the resonant frequency. This was for a simulated system with $k = 20, m = 0.4, b = 5.2$.

Taking parameters for the oscillator in the middle of these values, i.e. $w_n = 4.5, h = 1$, the robustness of the oscillator to changing systems can be measured. The stiffness can be varied from 14–25.4 Nm/rad with the final frequency remaining within 10% of the natural frequency. This is a




Figure 3-13: SIM Figure showing the effect of the gain on the describing function of the oscillator. The plot shows $1/N(j\omega, A)$ plotted on the complex plane for two values of input gain h , as marked. It also shows $G(j\omega)$ for a typical mass-spring system, with the limit cycles marked with a $\circ$ for $h = 0.7$, and a $\square$ for $h = 1.7$. Because the lines of $1/N$ are approximately radial (especially for the middle frequencies), the effect of doubling the gain on the coupled system behavior is small. For high and low frequencies, changing the gain will alter how the oscillator couples with another system but the effect of the gain is still minor, given the acute angles between the lines and the radial directions.




Figure 3-14: SIM Figure showing the range of values of the oscillator natural frequency, and oscillator input gain over which the final frequency of a coupled oscillator, mass-spring system was within 10% of the system natural frequency. The mass-spring system had values $k = 20, m = 0.4, b = 5.2$ corresponding to a natural frequency of 8.66 rad/s and damping ratio of 0.3.

30% change in stiffness, or approximately a 12% change in system resonant frequency. These results demonstrate clearly the robustness of the oscillator to parameter values and to changes in the system properties.

This resonance tuning behavior is the natural behavior of the coupled system. Since the oscillators entrain quickly, this means that the oscillator system is reactive to fast changes in the dynamics. An alternative method would be to identify the system and use a constant frequency command to excite the resonant frequency. While this method would probably give more accurate resonant tuning, it would not be responsive to changes in the dynamics without re-identifying the system.

The resonant frequency is also the most efficient speed to move the mass-spring system, with the minimum of real work required to produce the motion. The fact that the oscillator can find and drive at this frequency, while still adapting to the system properties is thus a useful behavior.

System changes The oscillator is also robust to other changes in the system properties. For example, changing the damping in the system alters the shape of $G(j\omega)$, but does not alter the frequency. Since the oscillator can produce a variety of gains and frequencies, it is automatically robust to these changes. For example, the damping ratio of the example system described above can be altered from 0.12 to 0.5 while still remaining entrained, with the entrained frequency varying from 7.56 to 6.59 rad/s where the resonant frequency is 7.07 rad/s, as shown in figure 3-15. This corresponds to a 7% change in frequency.

Entrainment The final aspect of the oscillator control which is practically useful is that it tends to entrain very quickly, usually within one cycle. This means that it is easy to tell whether a parameter is correct or not as the system settles quickly. This also improves the robustness, since the oscillator responds quickly to changes in the dynamics.




Figure 3-15: SIM Figure showing the effect of varying the damping ratio of the mass-spring system. The plot shows the describing function for the oscillator with lines for $G(j\omega)$ for different values of the damping ratio as marked. The damping ratio can be altered from 0.12 to 0.5 while remaining entrained, with only a 7% change in the output frequency away from the resonant frequency: frequency varies 7.56–6.69 rad/s, with the resonant frequency of 7.07 rad/s.

3.8 Design example: Resonance tuning

The describing function analysis is not only appropriate for analyzing the Matsuoka oscillator. This section shows how it can be used to analyze the well known Van der Pol oscillator, for the task of resonance tuning. There have been a number of papers concerning the resonance tuning property of neural oscillators (e.g. Hatsopoulos et al. (1992), Hatsopoulos and Warren (1996), Epstein and Kopell (1999)), without any full explanation of why this behavior is observed or predictions of its limits. The describing function analysis provides answers to both of those questions.

The Van der Pol (VDP) oscillator (Strogatz, 1994, p. 198), is described by the following equations:

$$\tau \dot{u} = wv + f(u) + h_j g_j \quad (3.13)$$

$$\tau \dot{v} = -\epsilon u \quad (3.14)$$

$$f(u) = u - u^3/3 \quad (3.15)$$

$$y_{out} = h_o u \quad (3.16)$$

where u and v are state variables, g_j is an input, h_j a gain on the input, h_o is a gain on the output, τ determines the speed of the oscillator, and ϵ is a parameter.⁴

The describing function for this oscillator is illustrated in 3-16, which shows a rather different pattern from the Matsuoka oscillator. The same trends are evident; the gain reduces with increasing A , although the output amplitude alters as a function of A , making the range of gains less uniform. The gain as before is roughly constant with frequency. The phase is roughly independent of amplitude, in the range $\pm 30^\circ$.

⁴Reasonable values are $0.1 \leq \epsilon \leq 0.1, 1 \leq h_j \leq 30, 0.1 \leq h_o \leq 1$.




Figure 3-16: SIM Van der Pol Bode plot. The top graph shows the gain $|N(j\omega, A)|$ plotted against frequency, and the lower plot the phase. The multiple lines correspond to different values of input amplitude A as indicated by the numbers and arrows on the plot. The gain is inversely related to A and is roughly constant with frequency. Unlike the Matsuoka oscillator, the output amplitude of the VDP oscillator is a function of the input amplitude. This results in the lines in the gain plot not being equally spaced (compare to figure 3-7). The phase is less dependent on A , and is grouped around zero degrees. The natural frequency of the oscillator is indicated by the vertical line.




Figure 3-17: SIM Plot of $G(j\omega)$ and $1/N(j\omega, A)$ in the complex plane for the VDP oscillator. $G(j\omega)$ is the thick line, with numbers indicating frequency. The lines for $1/N(j\omega, A)$ correspond to constant frequency, with amplitude A increasing along each line as indicated by the numbers and arrow. There is a limit cycle solution at the point where the two graphs intersect at the same frequency. The intersection point is indicated by a filled circle, at $\omega = 7.0$ and $A = 7.4$. This plot was produced for $\tau_1 = 0.03, \epsilon = 0.1, h_o = 0.2, h_j = 1, k = 20, m = 0.4, b = 2$.

If this oscillator is connected to a mass-spring system as before, because the Van der Pol oscillator can only produce phases in the range $\pm 30^\circ$, the intersections would all be at low frequencies. However, if the velocity rather than the position is used as input, those phases correspond to points near the resonant frequency of the system. This is illustrated in figure 3-17, which shows $G(j\omega)$ for a mass-spring system with input θ_v , and output velocity $\dot{\theta}$, plotted together with $1/N(j\omega, A)$. The input and output gains are calculated to position the two plots over one another, here $h_j = 1, h_o = 0.2$. As before, there are limit cycle solutions where the graphs intersect at the same frequency, here $\omega = 7.0, A = 7.4$.

Because of the shape of the phase plot of the VDP, the oscillator can tune into the resonant frequency of the plant motion over a range of system frequencies. The Matsuoka oscillator also has this effect, but because the lines of its describing function are not as parallel, or as aligned with the resonant phase difference as the VDP characteristic, the resonance tuning behavior is not as strong. Using the describing function analysis, the resonance tuning effect of the VDP can be calculated, and measured in simulation. Figure 3-18 shows the result of this prediction. The plot shows the oscillator tuning into the resonant frequency of the system over a large range of frequencies. The accuracy of the describing function prediction is also shown in the graph.

To confirm this result, the VDP oscillator was implemented on the robot, and a similar experiment was conducted, varying the stiffness of a robot joint, and measuring the final frequency. The results from this are plotted in figure 3-19. The results show that over a range of frequencies corresponding to a four-fold increase in stiffness, the VDP oscillator drove the system at its resonant frequency.

This example illustrated not only the performance of the VDP oscillator system in this task, but also the use of describing function analysis to design and predict final system performance.

The VDP oscillator is in general less robust than the Matsuoka oscillator. This is because the




Figure 3-18: SIM Prediction of VDP oscillator driving a mass, using velocity as the oscillator input. The upper plot shows the accuracy of the prediction, and the resonance tuning behavior of the oscillator. The oscillator drives the mass at the resonant frequency of the mass-spring system, over a range of different conditions (here the frequency of the mass-spring system is varying from 5 to 9 rad/s). The lower plot shows the accuracy of the prediction of motion amplitude. For these plots $\tau_1 = 0.03$, $\epsilon = 0.1$, $h_o = 0.1$, $h_j = 2$, $k = 20$, $m = 0.4$, $b = 2$.

output amplitude is dependent on the input amplitude, making tuning difficult, and because the range of output phases is much less.

3.9 Design example: Juggling

This section describes the application of describing function analysis to a more complex task, that of juggling a ball on a paddle. The setup is illustrated in figure 3-20. This task has been addressed by a number of researchers, (Schaal and Atkeson, 1993, Rizzi and Koditschek, 1994), including one using neural oscillators (Miyakoshi et al., 1994). Using the method presented in this chapter, not only can the parameters be easily tuned without intensive simulation, but the robustness of the system can be calculated and compared to other approaches.

Schaal and Atkeson (1993) showed that when juggling a ball on a paddle, stable motion is achieved when there is a particular constant phase difference between the paddle and ball motion. This motivates the approach of writing the juggling problem as a describing function, where stable juggling is defined in terms of a phase between input (paddle motion) and output (ball motion), as well as a gain (the amplitude of the ball). Although this is a rough approximation to juggling, it allows analysis of the situation. The setup is shown in figure 3-21, where the paddle is assumed to move sinusoidally, i.e. $y(t) = B \sin(\omega t)$, and the ball bounces with height x . Since the juggling is stable, the velocity of the ball before and after impact is the same, so the impact equation is

$$(\dot{x} - \dot{y}) = -\alpha(-\dot{x} - \dot{y}) \quad (3.17)$$

where α is the coefficient of restitution. Calculating the time between bounces (which is $2\pi/\omega$), and




Figure 3-19: **REAL** Plot shows resonance tuning action of a VDP oscillator actuating a joint on the real robot. The stars mark measured frequency, plotted against the natural frequency of the system (also measured), as the stiffness of the system was altered. The results fit with the predictions of the describing function analysis. Because the natural frequency is a function of the square root of the system stiffness, the actual stiffness varied from 9 to 42.




Figure 3-20: Picture of robot in juggling task. The paddle is a table tennis bat attached to the robot's arm. The ball is restrained to travel in one dimension by being mounted on a rotating boom. The angle of the boom is measured using a potentiometer at the pivot point, and this feedback signal is used by the oscillator to coordinate the arm motion with the ball motion.




Figure 3-21: Schematic of the juggling action. The thick line shows the motion of the paddle $y(t) = B \sin(\omega t)$, and the thin line the ball trajectory $x(t)$. The ball impacts the paddle at a phase ϕ of the sine wave. The analysis expresses the juggling as a describing function, writing a stable solution in terms of the phase difference between ball and paddle trajectories ($\phi + \pi/2$), and the amplitude of the ball motion $\bar{x}$.




Figure 3-22: Oscillator driven juggling. The amplitude of the ball motion is $\bar{x}$, the input to the oscillator is $A = h_j \bar{x}$, and the output of the oscillator drives the paddle ($G(j\omega)$) with amplitude B . The ball motion can then be calculated using N_j . As before, there can be only a limit cycle when the loop gain is unity.

substituting $\dot{y} = B\omega \cos(\phi)$ at impact gives an equation for the phase at impact:

$$\phi = \arccos \left[\frac{\pi g}{B\omega^2} \left(\frac{1 - \alpha}{1 + \alpha} \right) \right] \quad (3.18)$$

Calculating the maximum height of the ball (at the time halfway between bounces) gives an expression for the ball amplitude:

$$\bar{x} = g\pi^2 / 4\omega^2 \quad (3.19)$$

This makes the describing function for the juggling $N_j(j\omega, B, \alpha) = \bar{x} \exp(-j(\phi + \pi/2))$. The paddle is driven by an oscillator with the same control system as before, i.e. through a spring, so that the oscillator output drives the desired position of the paddle, and the input is the ball trajectory x , as illustrated in figure 3-22. Since the juggle describing function is a non-linear function of B , which is itself generated by the non-linear oscillator/paddle, the functions cannot be separated and plotted as before. A loop gain, and a condition for stable oscillator driven juggling can still be calculated:

$$\bar{x} = N_j(j\omega, [h_j \bar{x} N(j\omega, h_j \bar{x}) G(j\omega)], \alpha) \quad (3.20)$$

This can be solved numerically for the values of A , and ω which give stable juggling given the other parameters (h_j, A_n, k, m etc.).




Figure 3-23: SIM Robustness to changes in α . The plot shows ranges of α for which stable juggling can be achieved by either an oscillator with natural frequency ω_n , ($\bullet$), or a constant frequency sine wave $B \sin(\omega_n t)$, ($\circ$), as the frequency ω_n is varied. Because the oscillator can alter its frequency, it is more robust to changes in α than the constant frequency solution, and thus more suitable for the task.

Solving this equation is a powerful tuning tool. Using the equation, it was found that using ball velocity $\dot{x}$ rather than the ball trajectory x as input gave more robust solutions. This would have been very difficult to discover without this analysis.

Solving this equation is also powerful for measuring robustness and comparing solutions. The robustness of the oscillator solution can be measured by fixing the oscillator parameters and calculating whether solutions exist for different values of restitution coefficient α . This can be compared to a similar calculation for a constant sine wave $y(t) = B \sin(\omega t)$. For the sine wave, the limit on α is given by real solutions of equation (3.18). The maximum value of α is 1, and the minimum is a function of frequency:

$$\alpha_{min} = \min(0, (1 - B\omega^2/g\pi)/(1 + B\omega^2/g\pi)) \quad (3.21)$$

Figure 3-23 shows the comparison between oscillators with natural frequency ω_n against sine waves of frequency ω_n , over a range of different ω_n 's. The oscillators use velocity as input. The lines on the graph show the range of α 's over which the different methods could stably juggle. Because the oscillator can produce outputs over a range of frequencies, it can juggle with a wider range of α 's than the constant frequency solution. This result shows that oscillators are a better solution than a constant sine wave for this task.

Motivated by these results, juggling was implemented using a paddle attached to the arm of the robot. The setup is shown in figure 3-20. The ball was restricted to move in one dimension by a lightweight boom whose angle was measured using a potentiometer. As indicated by the analysis, ball velocity was chosen as the input signal to the oscillator. Figure 3-24 shows a short transient from the juggling performance. The juggling is much more variable than a simulated version, since there is a great deal of noise not only in the sensor signal, but also in the way the robot hits the tethered ball, effectively altering α for each impact. The change in speed of the oscillator driven paddle is visible in the plot. The oscillator responds on a hit by hit basis, remaining coordinated




Figure 3-24: REAL Plot showing paddle position (solid) and ball trajectory (dashed) against time for 2 traces of the robot juggling. Because of the way the arm hits the ball, there is significant noise in the system. The oscillator responds to the changing amplitude of the ball by adjusting its speed slightly, which is noticeable in the middle of the trace.

with the ball motion.

3.10 Adaptation

As well as responding to the type of system, and adjusting its frequency with respect to the system properties, the oscillator synchronizes itself with the mass motion. This was seen in the juggling example, the oscillator coordinating the motion of the paddle with the height of the ball on a hit by hit basis. The same action occurs when driving a single robot joint. This is because the oscillator is tightly coupled, and the oscillator output is referenced relative to its input, which is the motion of the joint. This means that the oscillator is sensitive to changes in the motion of the joint. For example if the oscillator is moving a certain link back and forth, and something or someone moves the link, the oscillator entrains with the imposed motion, and remains synchronized. This is illustrated in figure 3-25. The oscillator is sensitive to the exact motion of the joints, and not just the general system properties. This behavior and sensitivity is important and useful for providing coordination between the joints through the natural dynamics. This topic is examined in more detail in chapter 4.

3.11 Conclusion

This chapter has introduced the oscillators, and described their behavior coupled to single degree of freedom mechanical systems. Analysis using describing functions was introduced, and was shown to be powerful for analyzing the motion, predicting the behavior of the coupled system, and a practical tool for tuning, evaluating the robustness of the parameters, and evaluating the oscillator approach.

Using this analysis on oscillators reveals some general principles which apply to the oscillator solutions.




Figure 3-25: REAL Effect of external disturbances on the oscillator entrainment. The plots show the oscillator outputs (solid) and the joint angle (dash dot) for a single joint of the robot. Approximately halfway through the trace, the arm was moved by the author at a higher frequency and larger amplitude. When feedback is used (top graph), the oscillator responds to the extra motion, and remains entrained, its output tracking the imposed motion. In the lower trace when the feedback is off, the oscillator ignores the imposed motion, which is thus more jerky, since the oscillator and the disturbance are not coordinated.

- Robustness in parameter settings. Due to the wide range of gains that the oscillator can produce, the oscillator will entrain with a wide range of different systems. The behavior of the oscillator also makes the parameters of the oscillator robust because changing them either has a minor, or a predictable effect. The effect of varying parameters is smooth, monotonic, and they are decoupled from one another. This has the consequence that the system is easy to use in practice, it being easy to choose working parameters.
- Sensitivity to system changes. Again due to the shape of N , oscillator systems are sensitive to changes in the resonant frequencies of the mechanical systems, driving them at frequencies close to the resonant frequency. This occurs over a wide range of oscillator parameters, and a range of system frequencies. The oscillator behavior is automatic and is reactive, altering quickly when the system dynamics change. Since the resonant frequency is an efficient speed to move a mass-spring system, the oscillator behavior is finding the best speed, dependent on the natural dynamics.
- Stability of final solutions. Due to the shape of the plots, the oscillators produce stable motions actuating common mass-spring mechanical systems. This together with their observed fast entrainment to limit cycle solutions makes them well suited for practical applications.
- Synchronization. The oscillator produces a command to the joints which is defined relative to the joint motion. The oscillator is thus sensitive and responds to the joint motion, driving slower or faster when the joint is forced by the system, be it another joint of the arm or a human grabbing and moving the arm. This adaptation is crucial to the oscillator behavior when driving multiple degrees of freedom, as described in the following chapter.

Chapter 4

Coupling through natural dynamics

This chapter examines the behavior of the oscillators while driving multiple degrees of freedom of the arm, as opposed to the single degree of freedom motions examined in the previous chapter. The oscillators are connected to drive each joint of the arm independently, as shown in figure 4-1. There is a tightly coupled oscillator at each joint, with no connections between them. The oscillators use mechanical coupling through the physical arm to coordinate with one another and the task.




Figure 4-1: Figure showing configuration of oscillators and arm throughout this chapter. The oscillators are tightly coupled to each joint, with their output driving the setpoint θ_{vi} , and their input being either the joint angle θ_i , or the joint torque u_i . There are no software connections between the oscillators; they use mechanical coupling through the physical arm detected using the feedback to coordinate with one another and the task.

The versatility of this configuration is demonstrated using a number of examples, including crank turning using both one and two arms, and with the arms in both redundant and non-redundant configurations. Other position constrained tasks such as pumping a bicycle pump have also been demonstrated. The configuration can also exploit the dynamics of objects for coordination of the joints; the chapter includes the example of two arms coordinated through the dynamics of a Slinky toy, which is passed from hand to hand.

The analysis tools introduced in chapter 3 are extended to analyze these multi-degree of freedom motions. The analysis indicates again the sensitivity of the oscillators, showing that their behavior is to drive the resonant mode of the mechanical system. This is appropriate for a range of tasks, for example, attaching a crank to the arm creates a low frequency resonant mode of the system, which the oscillators automatically find and drive.

The crank turning example is analyzed in the context of these results, showing also the robustness of the oscillator solutions, with many different parameter settings giving crank turning, and a single set of oscillator parameters being robust to large changes in the arm dynamics.



Figure 4-2: REAL Crank turning with one arm. The picture shows the robot using its left arm to turn the crank. Oscillators are used at the shoulder and elbow in this configuration. The oscillators are initially uncoordinated, but use feedback from the motion of the joints to adjust their outputs, so turning the crank.

The chapter as a whole points out the self-organizing property of the oscillators, performing complex mechanical tasks which require coordination between the joints in a simple and robust manner.

4.1 Examples

The first example is crank turning. The oscillators have been used to turn cranks in a variety of different configurations, the most simple of which is shown in figure 4-2. Here two oscillators driving the shoulder and elbow joints are used to turn the crank in the plane. They use feedback of the joint angle to modify their outputs. As described above, there is no explicit connection between the oscillators at the two joints, their only connection being through the mechanics of the arm, and their feedback.

When the feedback is off, the oscillators oscillate at their natural frequencies, producing uncoordinated rhythmic drives to the arm joints. When the feedback is switched on, the oscillators respond to the dynamics of the situation, become coordinated with one another, and turn the crank smoothly, as shown in the transients in figure 4-3. Because the oscillators entrain rapidly, the crank turning settles down rapidly to a steady motion. The crank turning behavior is robust to changes in most of the oscillator parameters (gains, time constants etc), and is sensitive more to the posture of the arm, and the sizes of the oscillator amplitudes (tonic). The posture and the amplitudes determine whether the arm goes all the way round, and whether the final motion is smooth and energy efficient, or more violent. Because the arm is compliant, the system is still fairly robust to these values, with a variety of postures and amplitudes giving crank turning motion. This is considered




Figure 4-3: REAL Figure showing the effect of feedback for crank turning using two degrees of freedom in the configuration shown in figure 4-2. The lines show the angles of the joints. When the feedback is on, the oscillators are coordinated with one another, and the crank is turned (the crank angle is the dash-dot line, which wraps around at 180 degrees). Without feedback the oscillators revert to their natural frequencies, which does not result in smooth crank turning. The glitch in the crank trace is caused by the angle sensor wrapping round, not by a jerk in the motion.




Figure 4-4: REAL Crank turning with more than one arm. Since the oscillators are independent, the crank turning can be easily extended to use more than one arm. Here two arms are used to turn the crank, using four oscillators with two on each arm. As before, the oscillators are synchronized through mechanical coupling, there being no explicit connections between the oscillators on a single arm, or between the two arms.

in detail in section 4.7.

Since the control of the joints is independent, it makes no difference whether one or two arms are used. Figure 4-4 shows two arms turning the crank where each arm is in the same configuration as before. Four oscillators are used to achieve the motion, driving both shoulders and both elbows.

The crank is connected to the arm, and thus provides a strong constraint on how the arm can move. This is particularly the case in the configurations shown in figures 4-2 and 4-4 because the arm is not in a redundant configuration.¹ The motion of the crank determines how the shoulder and elbow joints can move. Because the oscillators use feedback from the joint angle, and because they provide an output which is a certain phase difference from their input signal (as described in chapter 3), they respond to the constraint by driving the arm around it.

The oscillators can also find crank turning solutions when the arm is redundant, i.e. configurations where the crank does not completely specify how the joints should move. A redundant configuration is shown in figure 4-5, where between four and six degrees of freedom can be used to turn the crank. The turning in this redundant case has the same properties as the non-redundant configurations: coordination through mechanical coupling, quick entrainment (see figure 4-6), and robustness to parameter changes. The oscillators have also been used for other constrained tasks such as pumping a bicycle pump.

Since the oscillator outputs are approximately sinusoidal, and the joint angle motions induced by the constraint are not sinusoidal (as in figure 4-6), there is a considerable tracking error. If the arm were stiff, this would be a problem and might cause the arm to jam. However, because the arm is compliant, the error is absorbed in the arm compliance, and the final motion is smooth. The compliance of the arm thus gives robustness. In addition, the compliance together with the position constraint expands the repertoire of the oscillators. The system exploits the crank itself to create a motion which is not solely created by rhythmic commands at the joints.

There have been many robotic approaches to crank turning, for example using impedance control

¹There are actually three joints in the plane of the arm, the shoulder, elbow and wrist, which makes this configuration redundant. However, the wrist joint was not actuated and maintained a constant stiffness, so to all intents and purposes the configuration is not redundant.




Figure 4-5: **REAL** Crank turning using a redundant arm. The oscillators can be used to turn cranks even when the arm is redundant, as shown here. The arm uses 4 oscillators, driving 4 degrees of freedom, two in the shoulder, and two in the elbow. The transients from this system are shown in figure 4-6.

(Hogan, 1985a), or hybrid force/position control (Raibert and Craig, 1981). These approaches use knowledge about the arm and crank kinematics to both move the arm to the crank, and to coordinate multiple degrees of freedom to turn the crank. They use this kinematic knowledge to deal with cranks of different sizes and different locations. The oscillator solution requires no calibration or kinematic transformations, but is only solving half of the problem. The arm starts connected to the crank, the posture and oscillator amplitudes set for the particular crank size and position, and the oscillators only provide the coordination for the motion. Even for creating the motion, the oscillator solution is more simple, with uncalibrated identical local controllers interacting through the arm mechanics, compared to a single complex calibrated kinematic controller.

The reason why the local control works is that the arm is compliant, making the interaction with the crank robust and giving some dynamics to exploit, and that the oscillators have the right properties to be local controllers for these tasks. These properties are examined later in this chapter.

Even in the redundant case, the crank provides a strong constraint on how the arm can move. However, the oscillators do not need such a strong constraint in order to coordinate multiple degrees of freedom through mechanical coupling. A different task accomplished in this manner is passing a Slinky toy from hand to hand, as illustrated in figure 4-7. An independent oscillator is used to control each arm, only coupled by mechanical interactions with the Slinky toy itself. The oscillators quickly converge on an out of phase motion of the Slinky, as shown in figure 4-8. The coupling forces from the Slinky are small compared to the crank forces, but are still large enough to entrain the oscillators.

In all these examples, the oscillators are using the mechanical coupling through a physical structure to become coordinated with one another. This allows them to perform complicated tasks in a very simple manner. In the following section a simple model of this coupling is developed. The model shows that the behavior of the oscillators is to find the resonant mode of the mechanical system. Connecting the arm to the crank, or connecting the arms with a Slinky can thus be seen as altering the natural dynamics of the system so that the resonant mode is the movement required to successfully complete the task.




Figure 4-6: **REAL** Transients of crank turning with a redundant arm (the configuration shown in figure 4-5). Three degrees of freedom are shown here, for two shoulder joints and one elbow joint. The top two graphs show the joint angles, and the motion of the crank as the feedback is turned on, which occurs at the vertical line. The system has a short transient before finding the stable motion. The lower two plots show the same result, only this time the feedback is turned off at the vertical line. The system quickly falls out of coordination, and the crank stops being turned smoothly. The glitch in the crank trace is caused by the angle sensor wrapping round, not by a jerk in the motion.




Figure 4-7: Picture of Cog passing the Slinky toy from hand to hand. The two elbow joints are used to move the hands up and down, where the coordination between the hands is given by the interaction between the oscillator dynamics and the coupled arm-slinky system.




Figure 4-8: REAL Two examples of Slinky operation. Both plots show the outputs from the oscillators as the torque feedback (dash-dot) is turned on and off. When the traces are in phase, the Slinky is moving in anti-phase. When the feedback is on, the two arms are coordinated and the outputs are synchronized, but when off, the oscillators are no longer synchronized. The only connection between the oscillators is through the physical structure of the Slinky.

4.2 Coupling model

The action of the mechanical coupling through the arm is to constrain the motion of each joint dependent on the motion of the other joints. Since the arm is compliant and in most configurations redundant, the coupling is not very stiff, i.e. there is some slop in the system. A simple approximation to this coupling is shown in figure 4-9. This shows two robot links represented as masses, with the joint level control appearing as a spring and damper connected to each mass. The coupling is included as a spring coupling the two masses, the stiffness of which can be varied to model the strength of the coupling. Using describing function analysis to represent the oscillators driving each mass extends the results in chapter 3 to this higher dimensional system.

In detail, the coupling model in figure 4-9 consists of two masses m_1, m_2 driven through springs k_1, k_2 and dampers c_1, c_2 by oscillators. The motion of each mass is coupled to the other by the coupling spring k_T .² The motion or angle of each mass is θ_1, θ_2 , and the oscillator outputs are θ_{v1}, θ_{v2} .




Figure 4-9: A simple model of coupling through the natural dynamics. The model consists of two masses driven by oscillators, connected by a coupling spring k_T .

The equations of motion for this system are

$$\begin{aligned} m_1 \ddot{\theta}_1 + c_1 \dot{\theta}_1 + (k_1 + k_T) \theta_1 - k_T \theta_2 &= k_1 \theta_{v1} \\ m_2 \ddot{\theta}_2 + c_2 \dot{\theta}_2 - k_T \theta_1 + (k_2 + k_T) \theta_2 &= k_2 \theta_{v2} \end{aligned} \quad (4.1)$$

This is a resonant mass-spring system, which has a free-vibration behavior when there is no driving input (i.e. $\theta_{v1} = \theta_{v2} = 0$). There are two resonant modes, at two different resonant frequencies. These can be determined by assuming solutions of the form $\Theta = Ae^{j\omega t}$, and solving the eigenvalue problem

$$\begin{pmatrix} (k_1 + k_T)/m_1 & -k_T/m_1 \\ -k_T/m_2 & (k_2 + k_T)/m_2 \end{pmatrix} \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} = \omega^2 \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} \quad (4.2)$$

for values of ω and ratios Θ_1/Θ_2 .

When this system is driven by an oscillator, the behavior of the oscillator can be modeled using the describing function analysis developed in chapter 3. The oscillator behavior is expressed as a transfer function $N_i(j\omega, A)$, with a gain g_i and phase ϕ_i , evaluated over a range of frequencies and input amplitudes, i.e.

$$g_i e^{j\phi_i} = N_i(j\omega, A) \quad (4.3)$$

²In general the coupling should include a damper. Including the damper complicates the equations without changing their fundamental result so it was ignored.

Because the oscillator provides a driving force which is a function of the joint angles Θ , the effect of the oscillator is to turn the system from being a driven resonant system (4.1), to a freely vibrating system:

$$\begin{pmatrix} (k_1 + k_T + j\omega c_1 - k_1 g_1 e^{j\phi_1})/m_1 & -k_T/m_1 \\ -k_T/m_2 & (k_2 + k_T + j\omega c_2 - k_2 g_2 e^{j\phi_2})/m_2 \end{pmatrix} \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} = \omega^2 \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} \quad (4.4)$$

If this system has a steady state vibrating solution, then the solutions for ω must be real, since complex values of ω correspond to solutions which either decay or grow. This implies that the oscillator must cancel out the damping in the system, or in this case that the imaginary parts on the diagonals must be zero (Strang, 1993).

$$\omega c_1 - k_1 \text{Im}[g_1 e^{j\phi_1}] = \omega c_2 - k_2 \text{Im}[g_2 e^{j\phi_2}] = 0 \quad (4.5)$$

If the two oscillator driven systems are assumed to have the same damping factor ζ then this equation can be simplified. Writing $\omega_{ni} = \sqrt{k_i/m_i}$ and $2\zeta\omega_{ni} = c_i/m_i$ for each mass, the equation becomes

$$2\zeta\omega = \omega_{n1}g_1 \sin \phi_1 = \omega_{n2}g_2 \sin \phi_2 \quad (4.6)$$

This defines an important relation between the two oscillators, as well as a relation between the gain and phase of the individual oscillators and frequency. Because g_i and ϕ_i are both functions of frequency this equation can be solved numerically, fixing the frequency, and finding the value of A where

$$2\zeta\omega = \omega_{ni} \text{Im}[N_i(j\omega, A)] \quad (4.7)$$

Figure 4-10 shows the solution to this equation for a typical choice of values for the masses and springs. The solution for each oscillator is indicated by the squares and stars on the plot. There are a range of solutions which vary in frequency, the extent of the solutions indicating the range of frequencies over which the oscillators can cancel out the damping, and produce steady state solutions.

The actual frequency of the motion depends on the solution for the eigenvectors of the complete system. When the imaginary parts have been removed, the problem reduces to

$$\begin{pmatrix} (k_1 + k_T - k_1 \text{Re}[g_1 e^{j\phi_1}])/m_1 & -k_T/m_1 \\ -k_T/m_2 & (k_2 + k_T - k_2 \text{Re}[g_2 e^{j\phi_2}])/m_2 \end{pmatrix} \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} = \omega^2 \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} \quad (4.8)$$

The effect of the oscillator is through the real part of its transfer function, $\text{Re}[g_i e^{j\phi_i}] = g_i \cos \phi_i$. If the phase $\phi_i = 90^\circ$, then this term is zero, and the equations reduce to the underlying mechanical system. This would imply that the oscillator driving this system simply removed the damping, but did not interfere with the shape of the motion. Unfortunately the phases are not exactly 90° , as show in figure 4-10. They lie in the range $90^\circ - 160^\circ$, which means that they affect the system. However, as the analysis below will show, they affect the frequency of the oscillation far more than the final mode shape.

The expression for the eigenfrequencies is obtained by solving (4.8) and is complicated:

$$\omega^2 = \frac{1}{2} \left[G_1 + G_2 + k_T \left(\frac{1}{m_1} + \frac{1}{m_2} \right) \pm \sqrt{(G_1 - G_2)^2 + 2k_T \left(\frac{1}{m_1} - \frac{1}{m_2} \right) (G_1 - G_2) + k_T^2 \left(\frac{1}{m_1} + \frac{1}{m_2} \right)^2} \right] \quad (4.9)$$

where $G_i = (k_i/m_i)(1 - g_i \cos \phi_i)$. The values of g_i and ϕ_i are taken from the solutions to the damping constraint and are both functions of frequency. The equation can be solved numerically, with the




Figure 4-10: SIM Figure showing solutions to the damping constraint (4.6). The plot shows $N(j\omega, A)$ in the complex plane, where the lines correspond to constant frequency, with increasing amplitude A as shown by the arrow. The squares and stars show the points where the damping in the system is zero (* - mass 1, $\square$ - mass 2). The numbers show the frequencies of these solutions. The range of solutions indicates the range of frequencies over which the oscillators can drive the system. The solutions at the same frequency are close together, with similar amplitudes and phases. The phases vary from about 90° for $\omega = 10$ to near to 160° for $\omega = 6.8$. For this example, $k_1 = 30, m_1 = 1, c_1 = 4.65, k_2 = 25, m_2 = 1, c_2 = 3, k_T = 30$.

graphical interpretation of the solution shown in figure 4-11. This shows ω^2 plotted together with the right hand side of (4.9) evaluated at the damping solutions. In this case there are two possible solutions where the lines intersect, at $\omega = 7.4, 9.7$. The two solutions correspond to two different modes of the system motion, which can be calculated by evaluating the eigenvectors at those solution frequencies.




Figure 4-11: SIM Figure showing solutions for the final system frequency. The plot shows ω^2 plotted together with the expression for the eigenvalues (equation (4.9)) evaluated at the oscillator solutions to the damping constraint. There are steady state solutions where the lines intersect, in this case there are two solutions at $\omega = 7.4, 9.7$. The plot shows that the oscillators produce at most as many solutions as there are modes of the underlying system, since there are only two possible intersections here.

The expression for the eigenvectors (4.9) carries little intuition, but can be simplified by considering the relationship between the damping constraint and the oscillator properties. The damping constraint (4.6) implies that the imaginary part of N multiplied by the local resonant frequency ω_{ni} is constant. The oscillators behavior is such that at constant frequency, the phase of N does not change greatly with input amplitude (see chapter 3, figure 3-7). This is also evident in figure 4-10, where the phases for the two solutions are approximately equal. Under this assumption,

$$\begin{aligned} G_1 - G_2 &= (\omega_{n1}^2 - \omega_{n2}^2) - (\omega_{n1}^2 g_1 \cos \phi_1 - \omega_{n2}^2 g_2 \cos \phi_2) \\ &\approx (\omega_{n1}^2 - \omega_{n2}^2) \end{aligned} \quad (4.10)$$

Assuming that this can be neglected compared to the other terms in the eigenvalue expression, (4.9) can be simplified, giving the two values of frequency ω :

$$\omega^2 \approx \begin{cases} (G_1 + G_2)/2 + k_T(\frac{1}{m_1} + \frac{1}{m_2}) \\ (G_1 + G_2)/2 \end{cases} \quad (4.11)$$

From which the eigenmodes of the system can be calculated. The ratio Θ_1/Θ_2 also has two

values:

$$\Theta_1/\Theta_2 \approx \left\{ \begin{array}{l} \frac{2k_T/m_1}{k_1/m_1 - k_2/m_2 - 2k_T/m_1} \\ \frac{2k_T/m_1}{k_1/m_1 - k_2/m_2 + 2k_T/m_2} \end{array} \right. \quad (4.12)$$

which is approximately the same as the resonant modes of the underlying mechanical system. Figure 4-12 shows the comparison between the modes of the mechanical system, and the modes found by the oscillator in this case.

Interestingly and importantly, the oscillator only finds one mode of the system. This is unlike the linear case, where the final solution is normally a superposition of the eigenmodes. One important constraint on the oscillator behavior is that it can only entrain over a limited range of frequencies. If the mode of the original system has a natural frequency which does not lie in this range, then the oscillator cannot have that mode as a periodic solution i.e. the lines in figure 4-11 will not intersect.

There are cases where there are two possible stable solutions, as in the example above. In that case the oscillator finds one or the other depending on initial conditions. Figure 4-14 shows two time traces of the transients finding either of the two modes for the example above. The reason for finding one solution is that the two modes correspond to different entrained frequencies for the oscillator. The oscillator can only output one frequency, (see chapter 3, figure 3-9), and thus cannot drive both modes simultaneously. This characteristic appears to be peculiar to the Matsuoka oscillator, as the Van der Pol oscillator (introduced in section 3.8) can drive two modes at once, as shown in figure 4-15. Driving two modes is undesirable for most applications.




Figure 4-12: SIM Plot showing eigenmodes of original system (thick arrow), and oscillator driven system (thin arrow). The mode is plotted as a vector in Θ_1, Θ_2 space. The arrows are close together, indicating that the oscillator solution is close to the resonant mode of the underlying mechanical system.

This behavior is robust, both to changes in oscillator and system parameters as shown in figure 4-13. It is also accurate. The oscillator drives relative to the motion of the system and accurately tracks the resonant mode. If the resonant mode of the system is the desired motion (as in the crank turning case described in the following section), then the oscillator behavior is automatically synchronized with the correct motion for the task.




Figure 4-13: SIM Plot showing robustness of oscillator to finding mode of system, under changing oscillator and system parameters. The y axis of the plots shows the angle of the mode in Θ_1, Θ_2 space, i.e. $\arctan(\Theta_1/\Theta_2)$. The left hand plot shows the effect of varying the natural frequency of the oscillator, with the lines indicating the underlying mechanical system modes. The accuracy of the oscillator is remarkable, given that the oscillator frequency doubles over the plot range. The right hand plot shows the effect of altering the stiffness of k_1 , while keeping the other stiffness constant at $k_2 = 25$ Nm/rad. When k_1 is close to 25 Nm/rad the error in the modes is very low. As the stiffness gets bigger or smaller, the approximation $k_1/m_1 \approx k_2/m_2$ becomes less accurate, and the error gets bigger. However this is a minor effect. The robustness of the oscillator system is remarkable, given the change in stiffness of six times ($k_1 = 10 \leftrightarrow 60$ Nm/rad).




Figure 4-14: SIM Plot showing the motion of mass 1 (solid) and mass 2 (dashed) of the oscillator driven system under two different initial conditions. The oscillator converges on one of the two modes of the system, with no mixing or superposition.




Figure 4-15: SIM Plot showing transient of the two mass system when a Van der Pol oscillator is used. This oscillator was described in section 3.8. The plot shows the motion of mass 1 (solid) and mass 2 (dashed) over time. Unlike the Matsuoka oscillator, this oscillator can give final periodic solutions which consist of a mixture of frequencies, as shown in the traces.

4.3 Local stability

The local stability of the oscillator limit cycle motion can be easily determined by examining the effect of amplitude changes on the damping in the system. At steady state, the oscillators exactly cancel the damping in the system, giving a constant amplitude oscillation. If an increase in amplitude results in positive (dissipative) damping, and a decrease in amplitude results in negative (excitatory) damping, then the limit cycle is stable.

The damping term is given by the diagonal elements of the system matrix:

$$c_{tot} = 2\zeta\omega - \omega_{ni}\text{Im}[g_ie^{j\phi_i}] \quad (4.13)$$

Figure 4-16 shows $g_ie^{j\phi_i}$ plotted together with the line $j2\zeta\omega/\omega_{ni}$, for an example frequency $\omega = 7.8$. The steady state solution is marked with a box. If the amplitude of the oscillation increases, $\text{Im}[g_ie^{j\phi_i}] < 2\zeta\omega/\omega_{ni}$, so $c_{tot} > 0$ which corresponds to positive damping, reducing the size of the oscillation. A decrease in amplitude has the opposite effect, confirming that the steady state oscillation is locally stable.

4.4 Summary of conclusions from the model

To conclude, the periodic solutions which exist for a set of oscillators driving a spring-mass system have the following characteristics

- The solutions are stable steady state solutions, with the oscillator canceling out the damping in the original system.
- The modes of the system are close to the eigenmodes of the original unactuated system, over a wide range of oscillator and system properties.




Figure 4-16: SIM Figure showing $N(j\omega, A)$ and the line $j2\zeta\omega/\omega_{ni}$ to indicate the stability of the oscillator driven system. The damping term is given by $c_{tot} = 2\zeta\omega/\omega_{ni} - \text{Im}[g_i e^{j\phi_i}]$, which is zero at the steady state solution marked by a $\square$. An increase in the amplitude of the motion moves the system to the point marked by the $\circ$, for which point the damping is positive because $2\zeta\omega/\omega_{ni} > \text{Im}[g_i e^{j\phi_i}]$. The positive damping will make the amplitude decrease. Similarly, a perturbation to lower amplitude results in negative damping, causing an increase in amplitude. The periodic solution is thus stable.

- The frequency of the final solution is not the eigenfrequency of the original mode, but is given by the interaction of the oscillator and system as defined by equations (4.6), (4.9).

The oscillator behavior of finding the resonant mode is useful because in some tasks such as crank turning the mode is exactly the desired motion for the task. The oscillators thus automatically find the correct coordination with the task. This complex property emerges from the interaction of the oscillators with the dynamics of the arm. No calibration, very little tuning, and no kinematic calculations are required to drive the arm along its resonant mode.

The following sections will show that the robot performance during both the crank turning and the Slinky toy tasks fit well with the simple model, providing evidence that exploiting the resonant mode is useful in practice.

4.5 Crank turning



Figure 4-17: [Picture of Cog turning a crank] Picture of Cog turning a crank. The robot is using two shoulder joints and two elbow joints to create the motion of the crank.

The crank turning behavior is a good example of the oscillators coordinated through the natural dynamics. The task is illustrated in figure 4-17. The crank turning can be understood in the context of the analysis of the previous chapter as a resonance of the springy arm around the constraint of the crank. The “springs” between the joints in this case are highly non-linear, but can be approximated by considering the solution of one mode of the system. This section compares data from the robot crank turning to the simple model developed in the previous section.

The full dynamics of crank turning are complicated and non-linear, but have the same general form as the full arm dynamics:

$$M(\Theta)\ddot{\Theta} + C(\Theta)\dot{\Theta} + K(\Theta)\Theta = K'\Theta_v \quad (4.14)$$

When this system is undergoing steady state oscillation it can be approximated a linear system, written as a set of eigenvalues and vectors. Transforming the variables $\Theta = Uq$, where U is the matrix of eigenmodes, and writing the effect of the oscillator as a diagonal matrix $G \exp(j\Phi)$, this equation can be written

$$U^T M U \ddot{q} + U^T C U \dot{q} + U^T K U q = U^T K' G e^{j\Phi} U q \quad (4.15)$$

Which looking at one mode u_i reduces to the single equation:

$$u_i^T M u_i \ddot{q}_i + u_i^T C u_i \dot{q}_i + u_i^T K u_i q_i = u_i^T K' G e^{j\Phi} u_i q_i \quad (4.16)$$

or approximately

$$\tilde{m} \ddot{q}_i + \tilde{c} \dot{q}_i + \tilde{k} q_i = u_i^T K' G e^{j\Phi} u_i q_i \quad (4.17)$$

The accuracy of this as a model of the crank turning problem can then be assessed by comparing real robot data to the predictions of the model.

The crank turning performance was measured for the robot in a configuration similar to that shown in figure 4-17, using four oscillators to actuate both shoulder, and both elbow joints. The oscillator time constants and the arm stiffness were varied to examine the behavior over a range of conditions. The frequency of the motion was calculated using a zero-crossing detector, and the amplitude and phase of the various joint motions and oscillator outputs was measured using a single frequency Fourier transform. The mode of the system u_i was directly calculated from the amplitudes and phases of the joint motions, and the gain of the oscillator directly measured by comparing joint motions and oscillator outputs. The stiffness and damping at each joint was also measured.




Figure 4-18: REAL Plot of $\text{Im}[u_i^T K' G e^{j\Phi} u_i]$ against ω for the crank turning. The crank turning involves four joints of the arm, for which the mode u_i is calculated from measured amplitude and phase data. The three lines correspond to different values of the arm stiffness ($\circ$ low, $\square$ medium and $\diamond$ high stiffness). The theory predicts that these points should lie on straight lines, the slopes of which are detailed in table 4.1.

The model of the crank turning predicts that the effect of the oscillator is to cancel out the damping in the system, or

$$\tilde{\omega} \approx \text{Im}[u_i^T K' G e^{j\Phi} u_i] \quad (4.18)$$

Figure 4-18 shows a plot of the imaginary part of the oscillator behavior versus ω . The data points were taken by measuring the system behavior as the natural frequency of the oscillatory was altered at three different values of arm stiffness (scaling the stiffness of all the arm joints in the ratio 1 : 2 : 3). The real damping at each joint was kept constant.

The theory predicts that this should be a straight line with slope proportional to $\tilde{c}$. The lines are straight, which is itself a good result given the simplicity of the model and the complexity of the arm motion. The slopes of the lines (listed in table 4.1) increase with stiffness, indicating that there is some coupling between the stiffness and damping through the arm dynamics. The abscissa are roughly constant which is as predicted.




Figure 4-19: REAL Plot of $\text{Re}[u_i^T K' G e^{j\Phi} u_i]$ against ω^2 . The theory predicts that these lines should be straight, with approximately constant slope. The lines have gradients $-0.28, -0.24, -0.35$ for increasing stiffness which are indeed roughly constant. The line for high stiffness is the most different from constant, and also the most noisy. The abscissa are at 9.08, 15.52, 28.64 so roughly proportional to stiffness, as predicted by the model.

The theory also predicts that the real part of the oscillator should be inversely related to ω^2 :

$$\tilde{k} - \tilde{m}\omega^2 \approx \text{Re}[u_i^T K' G e^{j\Phi} u_i] \quad (4.19)$$

Figure 4-19 shows the plot for this situation. Again the straight lines are striking. The mass matrix of the system $\tilde{m}$ should be independent of the stiffness. The data supports this, with slopes which are roughly constant (see table 4.1). The abscissa reflect the effect of k in the equation above, being roughly proportional to the stiffness.

The accurate fit of the crank turning data with the model shows that the model is a good description of the overall system. It also shows that the crank turning can be thought of as a resonant mode of the arm-crank system, which the oscillators are tuning into and exciting. This perhaps explains why the crank turning generalized so easily to multiple arms, the crank turning

kay	$\text{Im}[u_i^T K' G e^{j\Phi} u_i]$				$\text{Re}[u_i^T K' G e^{j\Phi} u_i]$			
	Slope		Abscissa		Slope		Abscissa	
	Value	Norm	Value	Norm	Value	Norm	Value	Norm
1	-0.34	1	15.13	1	-0.28	1	9.08	1
2	0.56	-1.62	17.63	1.17	-0.24	0.86	15.52	1.71
3	1.66	-4.85	19.67	1.30	-0.35	1.22	28.64	3.15

Table 4.1: **REAL** Slopes and abscissa of the line fits in figures 4-18, 4-19. The raw values are shown in normal type, the normalized values in bold.

being the resonant mode of a larger system. It also suggests that the coupling through the Slinky toy also enforces a resonant mode of the two arm system.

The mode-finding property suggests that the oscillators would be suitable for other tasks where the resonance of the mass-spring system corresponds with the task. For example, other constrained tasks such as pumping a bicycle pump fall into this category. Methods to extend the oscillators to cases where the resonant mode of the system is not aligned with the task are described in chapter 5.

The oscillator behavior is useful because they automatically find the mode of the system without any extra computation, and they respond to the dynamics of the arm itself. The final solution is robust, because the oscillators are augmenting a natural motion of the arm.

The oscillator control is in fact more robust at high speeds, and does not work well at low speeds, which is the opposite of traditional robot control. For example, consider turning a crank using a stiff arm, with hybrid force/position control (Raibert and Craig, 1981). The arm would control force along the direction of the crank (and so exploit the physical crank to constrain the motion), and control position around the crank to move it. As the crank is turned faster and faster, the performance of both the position and force control will deteriorate as the dynamics of the arm become significant, requiring extra dynamical models. While intersegmental dynamics might change the exact structure of the crank-arm resonant mode, changing frequency will not change the fundamental structure of the dynamics. There will still be a resonant mode that the oscillators can find, and in any case the oscillators are responsive to the exact structure of the mode. By exploiting rather than canceling the arm dynamics, the oscillator system works well as the frequency increases.

4.6 Self-organization to find resonant mode

The oscillator system was always observed to find a rhythmic solution in the crank turning case. This corresponded to turning the crank all the way round over a wide set of parameters (see section 4.7), or to turning the crank part of the way round if the parameters were not set correctly. The coupling model predicts the existence of limit cycles and their local stability, but does not guarantee that the oscillators will find the solution. Unfortunately it is difficult to prove that the system will converge to the limit cycle (this is discussed further in chapter 6).

An intuitive explanation for the oscillator self-organizing behavior can be given in terms of the entrained state of the system. The oscillators start with some random phasing between one another, and by interacting with the crank settle on a solution where the phases differences between the oscillators are correct for the crank to be turned. The phases between the joints are varied by the oscillators falling in and out of entrainment with the arm.

At the steady state solution, all the joints need to move at the same frequency, with all the oscillators entrained. For the oscillator to be entrained, there needs to be a steady input, which would be given by the joint moving rhythmically. This means that there will only be a globally stable solution when all the oscillators have phases relative to one another so that they all get steady rhythmic signals from their joint angles. This situation corresponds to the crank being turned. If one of the oscillators does not get a steady input, it will become unentrained and its frequency will




Figure 4-20: REAL Transient of oscillator driven crank turning. The three plots show equilibrium points, joint angles and a schematic of the crank motion plotted against time. The feedback to the oscillators is controlled by the step trace, and the crank angle is indicated by the dashed line. When the feedback is turned on the oscillators at the different joints entrain to different frequencies causing the phases between the joints to change rapidly. When the stable turning phase is reached, the oscillators lock onto that phase, all the oscillators oscillate at the same frequency and the crank is turned stably.

change. The different frequency has the effect of rapidly varying the phases between the joints. When the phases between the joints becomes correct, the oscillators can stay entrained, and lock onto the crank motion. The process of falling into and out of entrainment thus allows the oscillators to change the phases between the joints, so finding the solution predicted by the describing function analysis.

Figure 4-20 shows a particularly long transient for the crank turning. When the feedback is switched on, the oscillators change their speed, which causes the phases between the different joints to vary. The elbow (dashed line) speeds up and then slows down when the system converges on the limit cycle motion.

4.7 Robustness

The oscillator solutions for crank turning are remarkably robust to choices of parameters. Crank turning can be achieved for a wide variety of different parameters, and a single set of parameters is appropriate for a wide range of system dynamic properties. The oscillator parameters can be divided into those which are extremely robust (values can be varied by factors of 3 or so without affecting performance), and those which are more sensitive (values can be varied by 15-20% without affecting performance).

The most robust parameters are the arm stiffness and inertia, oscillator time constants and oscillator feedback gains. Figure 4-21 shows the range of different arm stiffnesses and time constants which were tested on the robot. Both stiffness and time constants can be changed by factors of 3 without affecting performance. The feedback gain could be varied from values of 80 to 250 about a nominal value of 168. Figure 4-22 shows some results for two arm crank turning (the configuration shown in figure 4-4). The data shows the average turning speed and was taken by varying each parameter in turn while keeping the other parameters constant. The graphs show the robustness of the crank turning to these changes. In yet another experiment, 1 Kg masses were attached to both the upper and lower parts of the arm during crank turning, as shown in figure 4-23. These masses effectively doubled the inertia of each link so changing the arm natural frequency by approximately $\sqrt{2}$. Unfortunately no data was taken, but anecdotally the only effect was to reduce the speed of the crank turning slightly and not effect the coordination of the task.

Intuitively, the reason for this robustness is the entrainment properties of the oscillator. The final frequency of the crank depends on the arm stiffness and inertia as well as the oscillator frequency. Because the oscillator can entrain over a range of frequencies, the final motion can occur over a range of time constants and arm dynamic properties. The robustness to input gain is due to the oscillator properties; the oscillator output size and phase is for the most part independent of input size (see chapter 3).

The parameters which are more sensitive are the ones more tightly related to the resonant mode shape. These are the posture of the arm (about which it oscillates) and the amplitude of the joint motions (oscillator tonic parameter). These parameters have to be roughly correct: the crank turning will not work if all the joint amplitudes are zero, and the Slinky will not work if the hands are upside down. However, due to the compliance of the arm there is a significant range for these parameters which will perform the task.

Figure 4-24 shows the range of postures which gave reasonable turning for a two degree of freedom simulated arm (configuration similar to that shown in figure 4-2). The data was obtained by exhaustively testing postures near to a tuned solution. The posture of the shoulder and elbow joints can be independently varied by about 20% without affecting the crank turning. Figure 4-25 shows data from the real robot, again from the two arm crank configuration. The graphs plot the average turning speed as a function of parameter setting, with all other parameters kept constant. The posture of the shoulder has the most effect, as might be expected given its higher stiffness and it being the most powerful joint. Figure 4-26 shows the effect of varying the amplitudes at the




Figure 4-21: REAL Plot showing robustness of the oscillator driven crank turning to changes in arm stiffness and oscillator time constants. The lines show the extent of crank turning solutions for various experiments, and thus are samples of the underlying range of parameters possible. The ranges of working parameters was found by altering both the arm stiffness (by changing each joints stiffness by the ratio indicated from a nominal stiffness), and the oscillator natural frequency (changing τ_1 for all the oscillators). All these lines were taken with the same input scale. The range of robustness is clear, with both parameters changing by factors of 3, while still performing crank turning.




Figure 4-22: **REAL** The graphs show the average crank turning speed (in revs per minute) over four seconds starting from rest for two arm crank turning (the configuration shown in figure 4-4). Each graph shows the effect of varying one parameter while keeping all the others at their default values (indicated by vertical lines). The left hand column refers to changes in the left arm, and the right hand column to right arm changes. The inputs gains (marked pos-scale on figure titles) can be varied greatly without affecting performance. The time constants (marked tau1) also show some variation, requiring a minimum value for crank turning to work. The data shows the effect of varying the time constant at one oscillator while keeping all the other parameters constant, which is why the data is less robust than that in figure 4-21 where all the time constants were varied together. The ability of the oscillator to entrain over a range of frequencies is illustrated by insensitivity to time constants especially at higher values. The data was taken over a short period of time so part of the noise in the data is due to slow transients, rather than necessarily being a consequence of poor turning solutions.



Figure 4-23: **REAL** To demonstrate robustness of the oscillator system, it was used with two 1 Kg masses attached to the upper and lower parts of the arm. These masses effectively double the inertia of the arm links. Their effect on the oscillator control was to make the system turn the crank slower, but did not disrupt the crank motion. In this picture all six degrees of freedom of the arm were used to turn the crank.




Figure 4-24: **SIM** The graph shows the range of postures which give good crank turning for a two degree of freedom simulated arm. The turning was measured by simulating the system and recording those postures which gave 70% of the maximum average speed over a four second transient. The plot gives a good indication of the variation in posture that can be tolerated while still giving good turning. The solid line shows the posture with the highest speed.




Figure 4-25: REAL The graphs show the average crank turning speed (in revs per minute) over four seconds starting from rest for two arm crank turning (the configuration shown in figure 4-4). Each graph shows the effect of varying the arm posture while keeping all the others at their default values (indicated by vertical lines). The postures are expressed in degrees. The left hand graphs are for changes to the left arm, and vice versa. The graphs indicate that the range of postures for each joint for successful crank turning is about $10^{\circ} - 20^{\circ}$. The data was taken over a short period of time so part of the noise in the data is due to slow transients, rather than necessarily being a consequence of poor turning solutions.




Figure 4-26: REAL The graphs show the average crank turning speed (in revs per minute) over four seconds starting from rest for two arm crank turning (the configuration shown in figure 4-4). Each graph shows the effect of varying the oscillator output amplitudes while keeping all the others at their default values (indicated by vertical lines). The left graphs refer to changes in the left arm and vice versa. The only parameter that appears to be sensitive is the shoulder-a amplitude (tonic) of the left arm. Once this parameter is large enough, the crank turning is possible for a range of parameters. The system works with large values of these parameters by turning the crank more violently, the oscillators still finding the correct coordination.

joints. The shoulder parameter is most dominant with a minimum value for crank turning. The other amplitudes do not have a strong effect. The amplitudes can be increased without affecting the crank turning because although the crank is turned more violently, the oscillators can still tune into the correct coordination to produce the turning motion.

While these parameters are more sensitive than the other oscillator parameters, they are still tolerant to significant changes. This makes tuning parameters simple. While the general form of the arm posture and the joint amplitudes are important in defining the shape of the resonant mode and thus the final driven behavior, certainly in the crank turning case their exact values are not important. This is primarily because the compliance of the arm allows errors in position to be absorbed without causing excessive internal forces or jamming.

4.8 Quality of crank turning solution

If the oscillators are finding the resonant mode of the underlying mechanical system, then that should be reflected in the energy required to produce the motion. At resonance, the real work required to move the mass should be minimized, with a considerable amount of energy stored in the springs of the arm. To investigate this, the energy used by the oscillator solution was compared to a similar solution using sine waves to actuate the joints, where the amplitudes and phases between the sine waves were chosen to be similar to the oscillator solution.

The instantaneous work done $P(t)$ by the oscillator is given by the product of the torque at the joint $\tau(t)$ and the velocity of the set-point $\dot{\theta}_v(t)$.

$$P(t) = \tau(t)\dot{\theta}_v(t) \quad (4.20)$$

or

$$P(j\omega) = \tau(j\omega)\dot{\theta}_v^*(j\omega) \quad (4.21)$$

where $*$ is the complex conjugate. The work done $P(j\omega)$ is a complex number whose real part refers to the “active power” or work required to overcome dissipation in the system, and whose imaginary part refers to the “reactive power” or work stored and released in the springs (Bird, 1997).

The power was calculated for the oscillator and sine wave solutions, as illustrated in figure 4-27. The different graphs show the real and imaginary parts of the power for the different joints. The real parts are roughly the same, with the oscillator having a slightly higher value. For three out of the four joints, the oscillator solution has a higher reactive power. The only joint for which this is not the case is the shoulder joint, which has a small contribution to the total energy.

The ratio between energy stored and energy required is illustrated in 4-28, which shows that for three out of the four joints, the ratio is much higher for the oscillator driven system. This indicates that the oscillator driven solution stores energy in the compliance of the arm, so exploiting the dynamics of the arm to perform the task. The sine wave solution is close to the oscillator solution, but since it does not drive the resonant mode of the system, it does not produce a motion which exploits the natural dynamics of the arm as much as the oscillator solution.

4.9 Slinky toy

The second example of coordination through mechanical coupling is the task of passing a Slinky toy from hand to hand. Pictures and data for this task were included at the beginning of this chapter.

Data was collected from the operation of the Slinky in much the same way as for the crank turning. The results are included in appendix D. The main results are the same, with the data revealing the nature of the mechanical coupling through the Slinky. At low frequencies the effect of the Slinky mass provides coordination, while at high frequencies the spring-like properties of the




Figure 4-27: **REAL** Plot showing the real and imaginary parts of the power plotted against frequency for the different joints of the arm during crank turning. The imaginary part of the power (left hand graphs) refers to energy stored in the system, which is generally higher for the oscillator (*) than for the sine wave solution ($\diamond$). The only exception is for the shoulder-a joint, whose motion is very small, and the energy difference insignificant. The real part (right hand graphs) correspond to work done to move the arm, which is approximately the same for both control methods.




Figure 4-28: REAL Plot showing the ratio between energy stored and energy used to turn the crank, for the oscillator (*), and the sine wave ($\diamond$). The ratio is higher for the oscillator in all cases except for the first shoulder joint. This shows that the oscillator solution is exploiting the springy dynamics of the arm more than the sine wave solution.

Slinky have more effect. In either case the resonant mode corresponds to anti-phase motion, which is the desired motion for the Slinky.

The Slinky example is interesting because the forces on the hands due to the toy are much smaller than the forces due to the crank constraint, but the oscillators can still use the coupling to coordinate with one another.

4.10 Conclusion

This chapter has considered the behavior of the oscillators when driving multiple degrees of freedom of the arm, where the coordination between the oscillators comes from natural mechanical coupling, rather than explicit connections. The chapter presented an extension of the describing function analysis results in chapter 3 to multiple degree of freedom systems, showing that the oscillator driven limit cycles correspond to the resonant modes of the underlying mechanical system. The automatic behavior of the oscillator to find and drive systems in their resonant modes was shown to be useful in a variety of tasks, particularly for crank turning.

The oscillators are exploiting the natural dynamics of the arm in a number of ways to perform the task simply and robustly. Firstly the oscillators use the natural dynamics to couple the various joints, which removes the need for explicit connections. This makes the system more sensitive to the arm dynamics, and also reduces the dimensionality of the system, since there are less parameters to set. There is no need to specify the coordination between the links. The oscillators are also driving the arm in a natural motion, the resonant mode, which makes the overall system robust. In addition, the system exploits the constraints in the environment to constrain the movement of the arm, and so generate the movement.

The oscillator properties, together with the compliance of the arm, make the solutions robust to parameter and system changes (where parameters and stiffnesses can be multiplied by factors of 3 with no change to the system). This makes them appealing from a practical standpoint. The system quickly entrains to the task, and reacts appropriately to perturbations and changes to the system dynamics.

The oscillator solutions have also been shown to be versatile, since by actuating each joint independently the system scales easily to multiple joints, and multiple arms. Since the oscillators are simple and require virtually no tuning, scaling the system is easy. The number of tasks that can be cast as resonances is also large, especially given the sensitivity of the oscillator to both large and small coupling forces. Position constrained tasks, and ones where the manipulated object has some coupling dynamics are thus possible using this approach, as well as any other task where the resonant mode of the mechanical system is aligned with the desired motion. The method has also been extended to deal with motions which are not exactly determined by the resonant mode, these methods are considered in the following chapter.

Chapter 5

Coupling through explicit connections

5.1 Introduction

This chapter addresses the problem of extending the behavior of the oscillators to tasks which are not dictated by the natural dynamics. The problem here is constraining the oscillators in such a way that their robustness and self-organizational properties are preserved. The chapter highlights the difference between exploiting the self-organizational properties of the oscillators which are coordinated with the arm motion, and specifying constraints and relationships between the oscillators, which coordinate the oscillator outputs, not their inputs.

Three methods of constraining and modifying the oscillator behavior are considered in this chapter, illustrated schematically in figure 5-1. The first consists of adding connections between the oscillators, constraining the phase difference between the motion of the various joints. Adding connections immediately creates a tension between the oscillator entrainment with the mechanical system, and the oscillator entrainment within the network. The chapter will show that this tension is difficult to resolve, making the overall system sensitive to parameter changes and requiring tuning.

The second method uses a single oscillator to drive multiple joints. This is a generalization of the single degree of freedom systems (described in chapter 3) to multiple degrees of freedom. This method avoids some of the problems of using connections, but is still limited in its applicability. The method is illustrated by an implementation of robot sawing.

The third method relies on exploiting the resonant-mode finding behavior of the uncoupled oscillators, as described in chapter 4. The method uses extra artificially generated forces to manipulate the dynamics of the arm such that the final motion is a resonant mode of the system. The uncoupled oscillators can then find and drive the system in that mode.

5.2 Case (a): Networks of oscillators

Connecting oscillators into networks gives more complex behavior than a single oscillator because the phase relationships between the oscillators can be altered by changing the type and the strength of connections between the oscillators. Most applications of oscillators (for example in undulatory locomotion (Cohen et al., 1982), or legged locomotion (Taga et al., 1991, Kimura et al., 1998)) use oscillators connected into networks. An investigation into the stability and outputs from a variety of oscillator networks can be found in Matsuoka (1985).

Perhaps the simplest way to connect oscillators into a network is illustrated in figure 5-2, which shows two oscillators with inhibitory external connections. The connections make the input to the




Figure 5-1: Three methods for coupling oscillators. **(a)** shows connections between the oscillators which enforce phase differences between the joints. **(b)** shows a method using a single oscillator to drive a number of joints as a single unit. **(c)** shows the case where the natural dynamics of the arm have been artificially altered by adding extra forces either from a potential field, or from virtual springs as shown here. The extra forces modify the resonant mode of the system, and so modify the final steady state solution of the oscillators. The oscillators use mechanical coupling through the augmented dynamics of the arm to tune into the resonant mode, as described in chapter 4.




Figure 5-2: Simple oscillator network. The figure shows two oscillators with the usual mutual inhibition weights w and extra connections. The black circles on the end of the links correspond to inhibitory connections. The weights between the a neurons have strength $h_c \gamma$, and the cross coupled weights have strength $h_c (1 - \gamma)$, where $0 \leq \gamma \leq 1$. If $\gamma = 1$, neuron $2a$ will fire out of phase with neuron $1a$ and so in phase with neuron $1b$. The outputs of the neurons be thus be out of phase, i.e. $y_{out1} = -y_{out2}$. Setting $\gamma = 0$ will have the opposite effect, making the outputs oscillate in phase. Setting γ to some intermediate value will result in some intermediate phase between the outputs.

neurons for oscillator 1 :

$$\text{Input}_{1a} = \Sigma_j h_j [g_j]^+ + h_c \gamma [x_{2a}]^+ + h_c (1 - \gamma) [x_{2b}]^+ \quad (5.1)$$

$$\text{Input}_{1b} = \Sigma_j h_j [g_j]^- + h_c \gamma [x_{2b}]^+ + h_c (1 - \gamma) [x_{2a}]^+ \quad (5.2)$$

where the first term is the usual input from external systems, $[p]^+ = \max(p, 0)$, and the parameters h_c and γ ($0 \leq \gamma \leq 1$) control the effect of the oscillators on one another. The inputs for oscillator 2 are similar. Intuitively, these connections put the neurons in mutual inhibition, causing them to oscillate out of phase. For example, if $\gamma = 1$, the connections force $1a$ to be in phase with $2b$ and $1b$ in phase with $2a$, making the outputs out of phase (see figures 5-2, 5-3).

For values of γ different from 0 or 1, the intuition is not so clear. If $\gamma = 0.5$, the input to each neuron is the sum of the outputs from each neuron of the other oscillator. Since these are out of phase due to the mutual inhibition connections internal to that oscillator (the connections marked w in figure 5-2), the actual input will be rather complicated. This results in a final phase which is dependent on initial conditions, as illustrated in figure 5-3.

Although it is straightforward to achieve phase differences of $\pm\pi$, other phase differences are difficult to achieve. The values of the gains h_c and γ need to be tuned to provide the correct behavior, and even then it is not robust to initial conditions. The behavior of this network when connected to two different mechanical systems is considered in the next section.

5.2.1 Connections and feedback

When the network shown in figure 5-2 is connected to two different mass spring systems as shown in figure 5-4, the behavior of the overall system is complicated. There is a tension between the entrainment of the oscillators with the physical systems, and the entrainment within the network. This tension makes the system sensitive to parameter changes such as the strength of the coupling and the physical properties of the mass-spring systems, as well as sensitive to the initial state of the system. The overall system is thus difficult to tune.




Figure 5-3: SIM Plot showing how the relative phase between the oscillator outputs varies as a function of γ . The plot was produced by simulating the system from different initial conditions for each value of γ and calculating the phase using a Fourier transform. The oscillator outputs are in phase for small γ , and out of phase for high γ . For γ 's in the range $0.3 < \gamma < 0.7$ there is significant sensitivity to initial conditions, with the oscillator producing a variety of different phases.




Figure 5-4: Figure showing two oscillators with external connections with parameters h_c, γ , driving two different mass-spring systems. Each oscillator is tightly coupled to a mass-spring system, with output y_{outi} , and input g_i . If the oscillators were not connected together, they would entrain with the mass-spring dynamics to reach a steady-state frequency, which would not be the same for each oscillator. If the oscillators were not connected to the mass-spring systems, then they form a coupled oscillator system which would entrain to a third different frequency. When they are all coupled together the system is forced to have one frequency, which means that some part of the system is not fully entrained.




Figure 5-5: SIM Coupled system transients. The plots show the motion of the the set-points for the two masses (solid lines) and the actual mass motion (dashed), for two values of input gain h_c and $\gamma = 0$. The thick lines refer to one mass-spring system, and the thin lines to the other. This value of γ should result in the oscillators producing outputs in phase, but this is only realized for high values of the coupling strength h_c , as shown in the lower graph. The two solutions also have different frequencies, being slower when $h_c = 1$. The plot thus indicates the sensitivity of the system behavior to the values of the parameters. In addition, since the two mass-spring systems have different properties, even when the set-points are approximately in phase (lower graph), the mass motions have a phase difference between them. This means that some tuning would be required to make them move in phase.

Each part of the circuit shown in figure 5-4 would oscillate at a different frequency if isolated from the rest of the circuit. The oscillators driving the mass-spring systems would entrain to different frequencies (as described in chapter 3) and similarly the two oscillators entrain with one another through the connections at a third frequency. When the system is coupled together, each part is forced to oscillate at the same frequency, which implies that some part of the system is not fully entrained. Either the connections between the oscillators are dominant, the phase between the oscillators being set by the connection strengths and the effect of the masses largely ignored, or the mass-spring systems dominate. The strengths of the various gains and connections determine which signal dominates, which makes the overall system sensitive to parameter changes. This is in contrast to the oscillator systems analyzed in chapters 3 and 4, where changes in parameters might alter the speed of the motion, but do not radically alter the shape of it.

Figure 5-5 shows this effect. The figure shows two different responses from the coupled system for two different values of the connection strength h_c . When h_c is low, the mass motion dominates the system, and there is an appreciable phase difference between the oscillator outputs. On the other hand if h_c is large, the oscillator outputs are forced to be in phase.

The system behavior depends on the values of the gains in an inherently non-linear manner. This can be seen in the graph for $h_c = 1$ in figure 5-5. Although the oscillator outputs are in phase, the actual signals are distorted and different from one another. The distortion arises from the effect of the distorted input signals (shown in figure 5-6) on the oscillator dynamics. This behavior is more complicated than the oscillator response to a sinusoidal input (used for analysis in the previous




Figure 5-6: SIM Oscillator inputs. The graphs show the inputs to the individual neurons during the motion in the lower graph of figure 5-5. The top graph refers to oscillator 1, lower to oscillator 2. The inputs are not sine waves, particularly for oscillator 2 where the combination of inputs gives a very distorted signal. The distortion is created by adding the mass motion signal and the output from the other oscillator, and is one of the reasons for the sensitivity of oscillator systems coupled in this way.

chapters), and is correspondingly more difficult to predict. The input signal is itself distorted because it is a sum of approximately sinusoidal signals with different amplitudes and phases.

Even if the system could be tuned to produce outputs of a desired phase difference, that would still not be a good solution. As shown in the lower graph of figure 5-5, if the two driven systems have different dynamics, even driving them in phase will not result in in-phase motion. This is because the connections coordinate the oscillator outputs, not the oscillator inputs. Careful parameter tuning would be required to obtain the correct coordination between the masses. This contrasts with the systems analyzed in chapters 3 and 4. In those systems, the natural dynamics was used to coordinate the oscillator inputs, with the oscillator outputs driving relative to those inputs. This not only immediately gave the correct coordination with the task, but also had greater robustness to parameter and system changes than using explicit connections.

Some of these difficulties might be alleviated by ensuring that both systems have approximately equal resonant properties, and only requiring phase differences which are multiples of π . This is the case if oscillators are used for legged locomotion, since each leg has approximately equal properties, and for successful walking they must move out of phase. Nearly all of the successful implementations of oscillators with connections have been for legged locomotion (see chapter 2). For arm control, where different limbs segments have different resonant properties, and where few tasks can be easily specified in terms of phase differences between joints, connections in this way are not so appropriate. Choosing a different way of combining the inputs might also help, but since the oscillator behavior is so complicated, there is not an obvious combination method which does not introduce some undesirable non-linear effects into the system.




Figure 5-7: Master-slave network. A single oscillator is used to drive multiple joints of the arm using slave oscillators at the joints with explicit but unidirectional connections. Input is applied to the master, which specifies the sizes and phases for all the joints. For unidirectional connections, the effect of γ is also unreliable, with a similar behavior to that shown in figure 5-3. To overcome this, the slave oscillators can be conveniently replaced by constant gains as shown for the third joint in this picture. The magnitude of h_{o3} sets the amplitude for the third joint, with the phase (in-phase or anti-phase) set by the sign of h_{o3} .

5.3 Case (b): Master-slave network

One alternative which uses oscillators to control multiple joints of the arm is shown in figure 5-7. This uses a master oscillator to drive slave oscillators at the joints, where the connections from the master to the slaves are unidirectional. This configuration effectively makes a number of degrees of freedom operate together, under the control of one oscillator. Input to the master oscillator thus effects the motion of all the joints as a unit.

Given the unreliability of the connection strength γ (see figure 5-3), the slave oscillators can be replaced by simple gains. The magnitude of the gain determines the size of the joint motion, and the sign gives the phasing with respect to the master oscillator. The set-point for the i th joint thus becomes:

$$\theta_{vi} = h_{oi}y_o + \theta_{pi} \quad (5.3)$$

where h_{oi} is a gain multiplying the output of the reference oscillator y_o . If this gain is positive θ_{vi} oscillates in phase with y_o , and vice versa. θ_{pi} is the posture about which the oscillation occurs.

The input to the master oscillator should reflect the overall motion of the arm, as activated by the oscillator. Since the oscillator drives a number of joints as a unit, the input can be taken as a weighted average of all the joint motions. A sensible weighting factor is the size of the command to each joint, which is the gain h_{oi} . This makes the input to the master oscillator:

$$g = h_{in} \sum_i \frac{\theta_i}{h_{oi}} \quad (5.4)$$

where h_{in} is the normalization factor. Applying the input in this way gives a multiple degree of freedom generalization of the single degree of freedom system analyzed in chapter 3. By analogy, one would expect this system to have similar properties, giving a behavior similar to “resonance-tuning” as well as robustness to system properties and parameter values. In particular, one would expect this solutions to be more robust than the network considered in section 5.2.1. Although the approach is the same in that the oscillator outputs are coupled, the oscillator control ensures that these coupled outputs are synchronized relative to the input signal. They can thus respond as a whole to the system dynamics.

This configuration of oscillators was suggested by Schaal and Sternad (1998), although they suggested using oscillators at the joints rather than fixed gains, did not comment on the difficulty of producing phases, and also did not show how feedback could be applied to the master oscillator.

This configuration was implemented on the robot to perform a sawing task. A single oscillator was used to drive the two shoulder joints and two elbow joints of the arm using fixed gains. The gains were tuned to produce an approximately linear motion of the hand. A saw was attached and used to cut a 2-by-4 beam of wood. Stills from a video of this motion are shown in figure 5-8. The linear motion of the arm is not exact, however the inherent compliance of the arm allows the arm to deflect without causing the saw to jam. Because the relative phasing of the arm joints is fixed at $\pm\pi$ by the signs of the gains, the feedback to the master oscillator cannot change the relative phasing between the joints as in chapter 4. It does however alter the speed of the motion, giving a behavior similar to the resonance tuning described in chapter 3. The steady state motion thus exploits the natural dynamics by storing and releasing energy in the springs of the arm.




Figure 5-8: REAL Eight stills from the robot sawing a block of wood. The picture sequence runs from left to right. The pictures were taken at 0.1 second intervals and show the saw moving forward and backward, cutting the wood.

Figure 5-9 shows the energy used during sawing, plotted as a function of the oscillator natural frequency ($\propto 1/\tau_1$). The energy was calculated by integrating the instantaneous work done for all the arm joints over one oscillator cycle. The instantaneous power is defined as

$$P(t) = \tau(t)\dot{\theta}_v(t)$$

where $\tau(t)$ is the torque at the joint, and $\dot{\theta}_v(t)$ is the rate of change of the oscillator outputs. The top plot shows the real energy used in sawing, and the lower plot the “reactive” energy, the energy that is stored and released in the arm springs during the sawing. When feedback to the oscillators is used, more energy is stored in the springs of the arms over all frequencies. At low frequencies the arm does more real work per cycle, and at high frequencies the work is about the same.

The energy required for sawing does vary with frequency, due both to the damping in the arm (which is constant), and perhaps different resistances of the saw at different frequencies. The frequency under feedback is higher than without feedback, as shown in figure 5-10. There is also more variation, because the oscillator responds to the sawing resistance, cutting slower when the resistance is high, and faster when it is easier.

The energy plot is replotted in figure 5-11 as a function of the measured sawing frequency. The data is noisy, but suggests that under feedback more work is done, and more energy is stored in the joints of the arm, even for the same frequency. However, this increase in energy does not result




Figure 5-9: REAL Plot of energy expended in sawing versus oscillator natural frequency ($\propto 1/\tau_1$). The top graph shows the real power, i.e. the energy to overcome dissipation in the arm, and actually cut the wood, and the lower graph shows the imaginary power, which is stored and released in the springs in the arm. The $\square$'s refer to oscillator feedback on, and the $*$'s to feedback off. When the feedback is on, the imaginary power is larger, meaning more energy stored in the arm. The real power is larger for lower frequencies, and about the same at higher frequencies. This plot is misleading because the actual frequencies of the points are different, and the energy to move the arm varies with frequency. Figure 5-11 shows the same data plotted against frequency.




Figure 5-10: REAL Sawing frequency versus oscillator frequency for feedback on ($\square$) and feedback off ($*$). With feedback, the sawing is faster, with a greater range of frequencies. This variation is because the oscillators respond to the varying resistance of the sawing task. When the resistance is high, the oscillators drive slower and vice versa.

in a noticeable change in the sawing efficiency, as shown in figure 5-12. This data was collected by measuring the number of strokes taken to saw through a 0.5" by 1.5" block of wood. The data is noisy, but shows that the sawing efficiency is approximately the same with and without feedback.

In summary, using feedback the arm stores more energy in the springs of the arm, requires about the same energy to move, and cuts wood at the same rate as without feedback. Since the oscillator is only scaling frequency (the coordination between the joints is fixed), one might expect the energy stored to be the same in both cases. The reasons for the discrepancy might be the slightly different amplitudes of the sawing motion under feedback, or perhaps a second order effect due to the constant phase between the drives to the joints and the weighted average of the the joint motions enforced by the oscillators.

While the effect of the feedback on the sawing performance is not particularly clear, anecdotally the feedback causes the arm to behave in a sensible manner. The arm slows when the resistance is high, and increases in speed when the resistance is low. The oscillator ensures that the command to the joints matches the joint motion, which acts to prevent jamming. In addition, if the robot is aided by a human, the oscillators can respond and adjust to a new frequency, with cooperation between the human and the robot.




Figure 5-11: **REAL** Plot of energy expended in sawing versus sawing frequency. The top plot shows the real energy required to perform the sawing. The real energy is slightly higher with feedback. The lower plot shows the energy stored and released in the arm springs, which is greater for the oscillator with feedback than without it.

A simple extension of this method which would combine using connections and using mechanical coupling would be to drive different combinations of joints with separate oscillators. The coordination between these joint units or synergies could then be determined by the natural dynamics of the arm. This would allow more control of the arm motion without losing the power of coupling through the natural dynamics.

Although this method of connecting joints is more robust than the network described in section 5.2.1, it is limited in that the only possible phases are $\pm\pi$, and the overall motion is determined by rhythmic commands at the joints.




Figure 5-12: **REAL** Percentage of wood cut per stroke plotted against average sawing frequency. The time and number of strokes to cut a 0.5" by 1.5" block of wood was measured, and the amount of wood cut per stroke calculated as 1/number-of-strokes. The plot shows results using feedback (*), and without feedback ($\square$). There is very little difference between the two methods, although the data is extremely noisy.

5.4 Case (c): Modifying the natural dynamics

Chapter 4 showed that the periodic solutions found by sets of uncoupled oscillators corresponded to the natural modes of the underlying dynamics. This section considers the design approach of artificially altering the natural dynamics of the arm to manipulate the position of the resonant mode. This method exploits the mode-finding properties of the oscillators, and so has greater robustness to parameter changes than the previous two connection methods. On the other hand, it requires the design and implementation of the extra forces.

The dynamics of the arm can be easily altered by applying extra forces to the joints of the arm, perhaps computed from a defined potential field. If a potential field is defined as a function of joint angles, e.g., $V(\Theta)$, the extra torques at the joint can be calculated from the relation:

$$\tau_{potential} = -\frac{\partial V}{\partial \theta_i} \quad (5.5)$$

This can be easily added to the usual joint torque:

$$\tau_i = k_i(\theta_{vi} - \theta_i) - b_i\dot{\theta}_i + \tau_{potential} \quad (5.6)$$

The oscillator can then be used to drive the system, and the final motion is expected to be a resonant mode of the natural dynamics of the system plus the extra potential field. Unlike cases (a) and (b), this method exploits the oscillator properties to find the coordination between the joints and so should have the same kind of robustness and self-organizing properties as in the tasks described in chapter 4.

This method offers the promise of creating phase differences between the joints which are dependent on the task, rather than being restricted to $\pm\pi$. It also creates the possibility of producing motions which are more complicated than that possible using connections. This is because the

range of possible output motions is limited using connections alone by the fact that the oscillator is restricted to producing rhythmic movements at the joints. The possible motions of the robot arm are only those which can be created by sinusoidal-like commands at the joint level. If external constraints exist, such as the crank in chapter 4, or the potential field in this section, the motion of the arm is defined by the interaction of the oscillator outputs and the external constraints. The compliance of the arm is thus exploited to produce motions which are more complicated than given by rhythmic commands.

Since the oscillator behavior is insensitive to parameter values when coupled together using the natural dynamics, one would expect that insensitivity to carry through to the augmented system.

In addition, since the oscillator output is bounded (see chapter 6), the oscillator driving the augmented system is not expected to be unstable. If the potential field is passive, then adding it to the normal dynamics of the arm which are also passive (Takegaki and Arimoto, 1981, Arimoto and Miyazaki, 1985) will make the overall arm system passive. Driving this system with an oscillator will result in oscillations which have a stable, fixed amplitude.

The main disadvantage of this approach is that the potential field can be difficult to design, calculate and implement. Calculating a full potential field could also make the oscillators redundant, since a more traditional controller could be used to move the arm through the field. Potential field methods are well known in robotics and there is a considerable body of research on methods of calculating, controlling and dealing with local minima in potential fields (Latombe, 1991). Once a potential field has been calculated, it may not make sense to use an oscillator to negotiate it, unless the other aspects of the oscillator behavior are important: working with external constraints, and interaction with the properties of manipulated objects.

One advantage of the oscillator solution is that it is very sensitive to the shape of the resonant mode, not requiring a very stiff field. The coupling can be very gentle as in the example of the Slinky toy (chapter 4 and appendix D). This means that as long as the potential field cancels out the undesirable natural dynamics, the actual stiffness of the field ($\partial V / \partial \theta_i$) need not be very high. This makes the potential field easier to implement, since the calculation can be made at a lower rate without causing stability problems (Franklin et al., 1992).

One method which might be appropriate is to build up the potential field gradually. If the natural coupling between the joints is almost correct for a task, then the potential field need only correct the motion in a few locations. A natural way to locate where the field is needed is in the context of an ongoing imitation task. Imitation could be used to correct the errors (i.e. stiffen the field in certain places) and slowly build up a field which would achieve the required motion while still relying on the underlying natural dynamics.

The following section gives a simple example of a potential field used to create a circular motion of the arm.

5.4.1 Example: Circular motion of two joints

This section describes the implementation of a potential field to force the end of the arm to move in a circle, using two joints of the arm. The desired motion is shown in figure 5-13, where the desired circle radius is r^* . The field to produce this motion must provide a force outwards when the radius r is less than the desired radius r^* , and a force inwards when it is greater. The radius of the arm is instantaneously:

$$r = \sqrt{\theta_1^2 + \theta_2^2} \quad (5.7)$$

The appropriate potential field is radially symmetric, and is given by the expression:

$$V(\theta_1, \theta_2) = \frac{1}{2}k(r^* - r)^2 \quad (5.8)$$




Figure 5-13: Application of potential field. The desired coordination between the two joints is that they describe a circle. The potential field required to create this applies forces to push the arm (represented by the black spot) out from the center, and in towards the dashed circle, with no forces when the radius $r = r^*$.

where k is “stiffness” of the field. Differentiating this to calculate the torques gives

$$\begin{aligned}\tau_i &= k(r^* - r) \frac{\partial \sqrt{\theta_1^2 + \theta_2^2}}{\partial \theta_i} \\ &= -\frac{k(r^* - r)}{r} \theta_i\end{aligned}\tag{5.9}$$

This potential field was implemented on the two elbow joints of the arm. It had the effect of creating a circular motion of the joints. Figure 5-14 shows the effect of adding the field to the natural dynamics of the arm. Two oscillators with position feedback were used to actuate the arm joints, and the two graphs show the effect of turning on the potential field. Without the field there is no coordination between the joints, but adding the field provides a strong resonant mode which the oscillators find, resulting in the circular motion. The stiffness of the field k is low, in fact lower than the stiffness of the individual joints, so the radius of the motion is not exactly correct, nor the circle perfectly round. These could be made more precise by increasing the stiffness of the field. The phase difference between the joints is approximately 90° , a phase difference that is difficult to achieve robustly using explicit connections.

Figure 5-15 shows that the feedback to the oscillators is necessary. When the field is on, and the oscillators have no feedback, the arm is repelled from the center of the circle, but does not move smoothly. When the feedback is switched on, the oscillator quickly finds the circular motion.

On Cog, the potential field is implemented digitally, calculating the extra torques and applying them at the joints. The stiffness of the field is limited by the sampling rate of the joint angle information (Franklin et al., 1992), which on Cog is only 50Hz. This limits the application of this method. In theory, and in practice if the hardware on Cog is upgraded, more complex fields can be computed and stably implemented. This should pave the way for more complex movements than that demonstrated here.




Figure 5-14: **REAL** Effect of potential field. The graphs show the arm position of the real robot arm in terms of the two joint angles θ_1 and θ_2 , with and without the potential field. When the field is on, it makes a resonant mode of the system which the oscillators find. When the field is off, the coupling between the oscillators is just through the natural dynamics, which are not enough in this case to give a steady motion.




Figure 5-15: **REAL** Effect of oscillator feedback. The graphs show transients of the robot arm motion from a fixed position with the oscillator feedback on (left graph), and off (right graph). When the feedback is on, the oscillators use the feedback to converge on the circular motion. Without feedback, there is no coordination.

5.5 Conclusion

This chapter has presented a number of methods for designing motions which are not completely determined by the natural dynamics of the arm. The difficulty faced by all these methods is maintaining the robustness and self-organizing properties of the naturally-coupled oscillators while having control over the arm motion.

The first method involved using explicit connections between the oscillators. This was shown to be difficult to analyze and tune due to the complex non-linear interactions between the oscillators and the actuated systems. The fundamental difficulty is the tension between entrainment of the oscillators with the actuated system, and entrainment within the oscillator network. In addition, the overall range of motions possible with this technique is limited due to the sinusoidal joint level commands

An alternative which generalizes the single degree of freedom motions analyzed in chapter 3 to multiple degree of freedom motions was also presented. This is still limited in terms of the range of possible motions, but has greater robustness than the first method, since the feedback is used to coordinate the drive to the joints with their actual motion. This method was used for a sawing task with the oscillators automatically responding to the resistance of the saw, producing an energy efficient motion with energy stored and released in the joints of the arm.

The third method involved augmenting the natural dynamics of the arm with an artificial potential field. The field was designed to produce a resonant mode in a desirable configuration. Oscillators driving the joints without any explicit connections could then find and drive this resonant mode. Adding fields in this way is a potentially powerful way of retaining the oscillator properties, while at the same time increasing the complexity of possible motions.

By using either of the last two methods, the oscillator control can be extended to more complex tasks, and to tasks where the natural dynamics of the arm is not sufficient to fully constrain the motion. The final motion using both of these methods is robust to parameter and system changes as in the previous two chapters.

Chapter 6

Oscillator analysis

6.1 Introduction

This chapter returns to the question of analysis for oscillator driven systems. It presents some theoretically exact results concerning the oscillator behavior, both alone and coupled to various systems. This complements the less rigorous but more practically useful analysis in chapters 3 and 4.

The chapter begins by considering bounds on the oscillator state. This result is important since it shows that the oscillator is well-behaved when coupled to a broad class of stable linear systems. The chapter then presents an analysis of the oscillator (and any linear system coupled to it) as a piecewise-linear (PL) system. A PL system is a non-linear system comprising of a number of different linear systems whose dynamics are switched dependent on the value of the state. Looking at the oscillator in this way provides intuition about the the shape of the limit cycle, and makes it easy to calculate the fixed points of the system. The fixed points can be used together with the boundedness result to shed light on and sometimes predict the final system behavior.

The prediction is somewhat weak, predicting “oscillatory” motion rather than convergence to a limit cycle. The chapter shows how the exact results can be combined with describing function analysis to predict the existence of limit cycle solutions.

Once a limit cycle has been found, casting the system as a PL system makes it easy to test the local stability of the cycle. A rigorous method to do this is presented at the end of the chapter, together with some examples.

The tools and results presented in this chapter are useful for predicting the result of coupling the oscillator to new systems and are thus invaluable for design.

6.2 Boundedness of oscillator

This section will demonstrate that the oscillator output is bounded with and without inputs, and that the oscillator itself is bounded when driven by a bounded input. These two facts are important when considering with what systems the oscillator is well behaved.

Matsuoka (1985) demonstrated that the oscillator states are bounded without any input. This section describes how his analysis can be extended to include the effect of an input g weighted by

an input gain h . The equations of the oscillators in this case are given by

$$\tau_1 \dot{x}_1 = c - x_1 - \beta v_1 - \gamma [x_2]^+ - h[g]^+ \quad (6.1)$$

$$\tau_2 \dot{v}_1 = [x_1]^+ - v_1 \quad (6.2)$$

$$\tau_1 \dot{x}_2 = c - x_2 - \beta v_2 - \gamma [x_1]^+ - h[g]^- \quad (6.3)$$

$$\tau_2 \dot{v}_2 = [x_2]^+ - v_2 \quad (6.4)$$

$$[u]^+ = \max(u, 0) \quad (6.5)$$

$$[u]^- = -\min(u, 0) \quad (6.6)$$

$$y_{out} = [x_1]^+ - [x_2]^+ \quad (6.7)$$

Proposition 6.1 *Consider the oscillator described by equations (6.1)–(6.7), with an input $g(t)$ bounded by $g_{min} \leq g(t) \leq g_{max}$. Then*

(a) *The oscillator state is bounded*

(b) *The oscillator output y_{out} is bounded independent of $g(t)$.*

Proof: This proof is based on the proof in (Matsuoka, 1985), which showed that the oscillator state is bounded without inputs. The proof is extended here to include the effect of input g applied to the oscillator.

The proof proceeds by integrating each of the oscillator equations in turn and determining maximum and minimum values for the states given the oscillator parameters and a starting state $(x_1(0), v_1(0), x_2(0), v_2(0))'$. Integrating (6.2) gives

$$v_1(t) = v_1(0)e^{-t/\tau_2} + \frac{1}{\tau_2}e^{-t/\tau_2} \int_0^t [x_1(u)]^+ e^{u/\tau_2} du \quad (6.8)$$

which since $[x_1(u)]^+ > 0$,

$$v_1(t) \geq -|v_1(0)| \quad (6.9)$$

The lower bound for $v_2(t)$ is similar

$$v_2(t) \geq -|v_2(0)| \quad (6.10)$$

Integrating (6.1) gives

$$\begin{aligned} x_1(t) = & x_1(0)e^{-t/\tau_1} + c(1 - e^{-t/\tau_1}) - \frac{\beta}{\tau_1}e^{-t/\tau_1} \int_0^t v_1(u)e^{u/\tau_1} du \\ & - \frac{\gamma}{\tau_1}e^{-t/\tau_1} \int_0^t [x_2(u)]^+ e^{u/\tau_1} du - \frac{h}{\tau_1}e^{-t/\tau_1} \int_0^t [g(u)]^+ e^{u/\tau_1} du \end{aligned} \quad (6.11)$$

Since $[x_2(u)]^+ > 0$ and $[g(u)]^+ > 0$ the maximum value of (6.11) is independent of those terms.

$$x_1(t) \leq |x_1(0)| + c + \frac{\beta}{\tau_1}e^{-t/\tau_1} \int_0^t v_1(u)e^{u/\tau_1} du \quad (6.12)$$

The maximum value of $x_1(t)$ can then be calculated by applying (6.9)

$$\begin{aligned} x_1(t) & \leq |x_1(0)| + c + \frac{\beta}{\tau_1}e^{-t/\tau_1} \int_0^t |v_1(0)|e^{u/\tau_1} du \\ & = |x_1(0)| + c + \beta e^{-t/\tau_1} |v_1(0)| (e^{t/\tau_1} - 1) \\ & \leq |x_1(0)| + c + \beta |v_1(0)| \end{aligned} \quad (6.13)$$

The maximum value of $x_2(t)$ is thus similar

$$x_2(t) \leq |x_2(0)| + c + \beta|v_2(0)| \quad (6.14)$$

The maximum value of $v_1(t)$ can be found by applying (6.13) to (6.8)

$$\begin{aligned} v_1(t) &\leq |v_1(0)| + \max \left[\frac{1}{\tau_2} e^{-t/\tau_2} \int_0^t [x_1(u)]^+ e^{u/\tau_2} du \right] \\ &= |v_1(0)| + \frac{1}{\tau_2} e^{-t/\tau_2} \int_0^t (|x_1(0)| + c + \beta|v_1(0)|) e^{u/\tau_2} du \\ &\leq (1 + \beta)|v_1(0)| + |x_1(0)| + c \end{aligned} \quad (6.15)$$

with again a similar bound for $v_2(t)$.

$$v_2(t) \leq (1 + \beta)|v_2(0)| + |x_2(0)| + c \quad (6.16)$$

Looking at the minimum value of $x_1(t)$ from equation (6.11)

$$\begin{aligned} x_1(t) &\geq -|x_1(0)| - \max \left[\frac{\beta}{\tau_1} e^{-t/\tau_1} \int_0^t v_1(u) e^{u/\tau_1} du \right] \\ &\quad - \max \left[\frac{\gamma}{\tau_1} e^{-t/\tau_1} \int_0^t [x_2(u)]^+ e^{u/\tau_1} du \right] - \max \left[\frac{h}{\tau_1} e^{-t/\tau_1} \int_0^t [g(u)]^+ e^{u/\tau_1} du \right] \\ &= -|x_1(0)| - \frac{\beta}{\tau_1} e^{-t/\tau_1} \int_0^t ((1 + \beta)|v_1(0)| + |x_1(0)| + c) e^{u/\tau_1} du \\ &\quad - \frac{\gamma}{\tau_1} e^{-t/\tau_1} \int_0^t (|x_2(0)| + c + \beta|v_2(0)|) e^{u/\tau_1} du \\ &\quad - \frac{h}{\tau_1} e^{-t/\tau_1} \int_0^t g_{max} e^{u/\tau_1} du \\ &\geq -|x_1(0)| - \beta((1 + \beta)|v_1(0)| + |x_1(0)| + c) \\ &\quad - \gamma(|x_2(0)| + c + \beta|v_2(0)|) - hg_{max} \end{aligned} \quad (6.17)$$

or

$$x_1(t) \geq -(1 + \beta)|x_1(0)| - \beta(1 + \beta)|v_1(0)| - \gamma|x_2(0)| - \gamma\beta|v_2(0)| - (\beta + \gamma)c - hg_{max} \quad (6.18)$$

with a similar expression for $x_2(t)$, where the input term is g_{min} .

$$x_2(t) \geq -\gamma|x_1(0)| - \gamma\beta|v_1(0)| - (1 + \beta)|x_2(0)| - \beta(1 + \beta)|v_2(0)| - (\beta + \gamma)c - hg_{min} \quad (6.19)$$

(a) Since $g_{min} \leq g(t) \leq g_{max}$, equations (6.9), (6.10), (6.15), (6.16), (6.18), (6.19), (6.13), (6.14) establish the bounds on the states of the oscillator.


(b) The oscillator output $y_{out} = [x_1]^+ - [x_2]^+$, the maximum value of which is $\max(x_1)$, and the minimum $-\max(x_2)$. By equations (6.13) and (6.14) neither of these terms depend on the value of $g(t)$.

■

This result leads to the more useful proposition

Proposition 6.2 *Consider the oscillator connected to a linear system $G(s)$, as shown in figure 6-1. If $G(s)$ is Linear Time Invariant and stable, then all the states of the coupled system are bounded.*

Proof:

Figure 6-1: Oscillator coupled to a linear system $G(s)$.

1. The oscillator output y_{out} is strictly bounded by equations (6.13) and (6.14). This implies that the L_∞ norm for the oscillator is bounded.
2. If $G(s)$ is stable, then $\|G\|_1 \leq \infty$. This means that the signal g in figure 6-1 cannot become unbounded.
3. Since the input to the oscillator is bounded, all the states in the oscillator are bounded (by Proposition 6.1).

■

These two propositions determine bounds on the oscillator and the driven system states. They show that the coupled system will be bounded if the driven system is stable. If the driven system is unstable, then the overall system is not guaranteed to be stable. The behavior of the coupled system is not necessarily to converge to a limit cycle, or even have an oscillatory solution. To determine what type of solution occurs, it is necessary to examine the fixed points of the coupled non-linear system, as shown in the following sections.

6.3 Oscillator as piecewise-linear system

This section shows that the equations for the oscillator can be written as a piecewise-linear (PL) system. This format is easier to analyze than considering the non-linear equations directly. It also gives considerable intuition about the behavior of the system.

The origin of the piecewise-linearity is the max operator used in the oscillator equations. For example, consider the term $[g]^+$. When $g > 0$ this term is just g , and when $g < 0$ this term is zero. In both regimes, the term is linear, with the system behavior switching when $g = 0$. The oscillator has three variables (x_1, x_2, g) that appear in the equations with the max operator. This results in three switching surfaces $(x_1 = 0, x_2 = 0, g = 0)$, which divide the state space into $2^3 = 8$ regions. Each region has linear dynamics, which change whenever the state crosses a switching surface. Indexing the regions by $\alpha = 1, 2, \dots, 8$, the system can be written in the form:

$$\dot{\mathbf{x}} = A_\alpha \mathbf{x} + b + H_\alpha g \quad (6.20)$$

$$y_{out} = c_\alpha \mathbf{x} \quad (6.21)$$

where $\mathbf{x} = (x_1, v_1, x_2, v_2)'$ is the vector of state variables, A_α is the state transition matrix, and c_α is the output matrix.




Figure 6-2: Figure showing the different values of A_α and c_α in different regions of the state space. The state space has been projected onto the x_1, x_2 plane. These matrices were calculated for $\beta = 2, \gamma = 2, \tau_1 = 0.1, \tau_2 = 0.2$

Writing

$$\{p\}_{q>0} = \begin{cases} p & \text{if } q > 0 \\ 0 & \text{otherwise} \end{cases}$$

The values of these matrices are

$$\begin{aligned} A_\alpha &= \begin{pmatrix} -1/\tau_1 & -\beta/\tau_1 & \{-\gamma/\tau_1\}_{x_2>0} & 0 \\ \{1/\tau_2\}_{x_1>0} & -1/\tau_2 & 0 & 0 \\ \{-\gamma/\tau_1\}_{x_1>0} & 0 & -1/\tau_1 & -\beta/\tau_1 \\ 0 & 0 & \{1/\tau_2\}_{x_2>0} & -1/\tau_2 \end{pmatrix}, \\ H_\alpha &= \begin{pmatrix} \{-hg/\tau_1\}_{g>0} \\ 0 \\ \{hg/\tau_1\}_{g<0} \\ 0 \end{pmatrix}, \\ b &= \begin{pmatrix} c/\tau_1 \\ 0 \\ c/\tau_2 \\ 0 \end{pmatrix}, \\ c_\alpha &= (\{1\}_{x_1>0} \ 0 \ \{-1\}_{x_2>0} \ 0) \end{aligned} \quad (6.22)$$

If the input is ignored, there are four linear systems, dependent only on (x_1, x_2) . The values of A_α, c_α for these regions are shown in figure 6-2, for the case where $\beta = 2, \gamma = 2, \tau_1 = 0.1, \tau_2 = 0.2$. Since it is impossible to plot the 4-dimensional state space, it is visualized by projecting the v_1 and v_2 axes down on to the x_1, x_2 plane.

The behavior of the oscillator in the different parts of the state space can be determined by looking at the eigenvalues, eigenvectors and fixed points of each linear subsystem. The fixed points




Figure 6-3: SIM Linear regions, fixed points and eigenvalues. Each sub-plot corresponds to one of the four linear subregions of the oscillator. The titles of each sub-plot give the eigenvalues, the *'s mark the fixed points, and the patches indicate the regions for each subsystem. For regions (2), (3), (4) the system is stable with all the eigenvalues having negative real parts. The fixed points for these systems are all in region (1) and are thus not fixed points of the complete system. Region (1) has two eigenvalues which are unstable and a fixed point at $(0.2, 0.2, 0.2, 0.2)'$. This is the fixed point of the complete system.

are calculated from the point where $\dot{\mathbf{x}} = 0$, by solving (6.20)

$$\mathbf{x} = -A_{\alpha}^{-1}b \quad (6.23)$$

Figure 6-3 shows the eigenvalues and fixed points for the four different systems, projected onto the plane $v_1 = 0, v_2 = 0$. For regions (2), (3) and (4), the systems are stable, with either real negative eigenvalues or complex eigenvalues with negative real parts. The fixed points for these systems are all in region (1), so they are not fixed points of the complete system. Region (1) has two eigenvalues which are real and negative, as well as two which are complex with positive real parts, corresponding to a “saddle” fixed point. This fixed point is in region (1), and is thus a fixed point of the complete non-linear system.

Examining the eigenvalues and vectors of the system can give valuable intuition as to the overall system behavior, as shown in figure 6-4. In region (4), all the eigenvalues are negative and real, so that the state should move directly towards the fixed point, i.e. move towards region (1). Regions (2) and (3) are symmetrical, with fixed points in region (1), two real eigenvalues and two complex eigenvalues. The system response is to converge towards the fixed point, although the complex eigenvalues create a response which is a damped oscillation resulting in the “loops” in figure 6-4. The unstable fixed point is in region (1). There are two stable real eigenvectors with eigenvectors in the $(1, -, 1, -)'$ direction, ignoring the v_i terms. The two unstable eigenvalues are complex with eigenvectors in the $(1, -, -1, -)'$ direction. The combined effect of these modes is to move the state towards the plane $x_1 = -x_2$ and force the state out of region (1) along that plane. The limit cycle motion shown in the right hand graph of figure 6-4 shows this behavior. The trajectory converges into region (1), where it is pushed out of the other side by the unstable fixed point, loops back and across and so on.




Figure 6-4: SIM The left hand graph shows the approximate shape of trajectories projected onto the $x_1 - x_2$ plane. The right hand graph shows two trajectories plotted on the same projection, with the unstable fixed point marked with a $\bullet$. If the trajectory starts in regions (2), (3), or (4) it ends up in region (1). The unstable fixed point in region (1) makes the state converge to the direction $x_1 = -x_2$, and is unstable in that direction, causing the trajectory to leave region (1), entering either (2) or (3). The trajectories are unique and do not intersect one another, although they appear to in this plot because the variation of the other two states (v_1 and v_2) is not shown.

The only starting state that is observed to not go to the limit cycle is one where the unstable mode is not excited, i.e. with a starting state $x_1 = x_2, v_1 = v_2$. These trajectories converge to the fixed point in region (1).

6.4 Oscillatory solution

Given the boundedness shown by Proposition 6.1, and the fixed points determined in the previous section, we are now in a position to prove that the oscillator has an oscillatory solution. Unfortunately the proof is not strong, in that the oscillator state can only be shown to not become unbounded and not to stop (the only fixed points are unstable). The motion of the state is likely to be oscillatory, i.e. continuously varying, while not necessarily converging to a limit cycle or periodic solution. The term used to describe this solution is “oscillatory but not necessarily periodic” and is taken from Matsuoka (1985).

Proposition 6.3 *The oscillator described by equations (6.1)–(6.7), with no input ($g(t) = 0$) has oscillatory but not necessarily periodic solution.*

Proof: This proof is copied from Matsuoka (1985), the only difference being the technique used to determine the fixed points, and more detail on uniqueness. To prove the proposition, we need to show:

- (a) That the oscillator is bounded, which is true by Theorem 6.1.
- (b) That the oscillator has only unstable fixed points, which was demonstrated in the previous section.




Figure 6-5: Oscillator coupled to a linear system. The individual state equations shown in this figure can be combined to create a single piecewise-linear system as shown in equation (6.28).

(c) That the oscillator solutions are unique.

Uniqueness is demonstrated by checking the Lipschitz condition (Slotine and Li, 1991, p. 152), which states that the dynamical system $\dot{x} = f(x, t)$ has unique solutions if $f(x, t)$ is continuous and a strictly positive constant L exists such that

$$\|f(x_2, t) - f(x_1, t)\| \leq L\|x_2 - x_1\|$$

for all x_1, x_2 within a finite neighborhood of the origin and all t in the interval $[t_0, t_0 + T]$ (with T also strictly positive).

Since the oscillator is piecewise-linear, it meets this criteria within each linear region, and is so locally Lipschitz:

$$\begin{aligned} \|\dot{\mathbf{x}}_2 - \dot{\mathbf{x}}_1\| &= \|A_\alpha(\mathbf{x}_2 - \mathbf{x}_1)\| \\ &\leq \|A_\alpha\| \|\mathbf{x}_2 - \mathbf{x}_1\| \end{aligned}$$

where $L = \|A_\alpha\|$ which is strictly positive. ■

The result that the system is bounded, together with the ease of finding the type and position of the fixed points of the system sheds light on the overall behavior of the system. This is illustrated in the following section.

6.5 Connecting oscillators to systems

When the oscillator is connected to the linear system shown in figure 6-5, the coupled system remains piecewise-linear. Because there are three variables whose values are important (x_1, x_2, g) there are 8 linear subsystems. The coupled equations of the system are

$$\dot{\mathbf{x}} = A_\alpha \mathbf{x} + b + H_\alpha g \quad (6.24)$$

$$y_{out} = h_o c_\alpha \mathbf{x} \quad (6.25)$$

$$\dot{\mathbf{x}}_s = A_s \mathbf{x}_s + B_s y_{out} \quad (6.26)$$

$$g = c_s \mathbf{x}_s + d_s \quad (6.27)$$

where h_o is an optional output gain on the oscillator. These be rearranged to form the augmented system

$$\begin{pmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{x}}_s \end{pmatrix} = \begin{pmatrix} A_\alpha & H_\alpha c_s \\ h_o B_s c_\alpha & A_s \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{x}_s \end{pmatrix} + \begin{pmatrix} b + H_\alpha d_s \\ 0 \end{pmatrix} \quad (6.28)$$

where the switching surfaces are $x_1 = 0, x_2 = 0$ and $g = c_s \mathbf{x}_s + d_s = 0$.

Examining the eigenvalues and fixed points of the system then allows the prediction of the final system behavior. The fixed points are only fixed points of the whole system if they are inside the region of state space where the linear model is defined.

6.5.1 Example: No oscillation

For example, taking the system

$$G(s) = \frac{s + 0.1}{2s + 0.1} \quad (6.29)$$

This system can be easily written in the state space formulation (Ogata, 1970) and the augmented system in equation (6.28) constructed. The coupled system has two fixed points $\mathbf{x}_1, \mathbf{x}_2$ with eigenvectors λ_1, λ_2 :

$$\mathbf{x}_1 = \begin{pmatrix} 0 \\ 0 \\ 0.33 \\ 0.33 \\ -6.67 \end{pmatrix}, \lambda_1 = \begin{pmatrix} -10 \\ -0.05 \\ -7.5 + 9.68j \\ -7.5 - 9.68j \\ -5 \end{pmatrix}; \quad \mathbf{x}_2 = \begin{pmatrix} +0 \\ -0 \\ 0.33 \\ 0.33 \\ -6.67 \end{pmatrix}, \lambda_2 = \begin{pmatrix} -25 \\ 2.5 + 6.63j \\ 2.5 - 6.63j \\ -10 \\ -0.07 \end{pmatrix} \quad (6.30)$$

The first fixed point is stable (all eigenvalues have negative real parts), and the second is unstable. Since the overall system is bounded (equation (6.29) is stable), the states of the system cannot blow up. Since the overall system is high dimensional, more work on the details of the piecewise-linear systems would be needed to determine the final motion. The system might converge to the stable fixed point, and might also have an oscillatory, not necessarily periodic solution from the unstable fixed point. Some transients from the system are illustrated in figure 6-6 showing convergence to the first fixed point $\mathbf{x}_1$.

As this example illustrates, calculating the fixed points does not always give enough information to determine the system motion. They do however give some information about the system properties.

6.5.2 Example: Constant input

If the oscillator is not connected to a system, but has a constant input g , then the fixed points are given by the solution of the equation

$$\mathbf{x} = -A_\alpha^{-1}(b + H_\alpha g) \quad (6.31)$$

If $g = 1$, and the other parameters of the oscillator are set as usual, there is only one fixed point of the system:

$$\mathbf{x}_1 = \begin{pmatrix} -0.67 \\ 0 \\ 0.33 \\ 0.33 \end{pmatrix}, \quad \lambda_1 = \begin{pmatrix} -10 \\ -7.5 + 9.68j \\ -7.5 - 9.68j \\ -5 \end{pmatrix} \quad (6.32)$$




Figure 6-6: SIM Two transients of the oscillator connected to the system in equation (6.29). The plot shows the trajectories plotted on the $x_1 - g$ plane, and illustrates convergence to the stable fixed point at $(0, 0, 0.33, 0.33, -6.67)'$, marked with a $\square$. The other fixed point of the system is unstable resulting in the oscillatory motion at the top of the graph. The eigenvalue for convergence to the fixed point is very low (-0.05), so the system converges slowly.

This fixed point is stable. Unfortunately, even though this is the only fixed point of the system, more work is needed to determine the system behavior. Since adding the input moves all the fixed points in the system, but does not change any of the linear subsystems, one might expect the system to converge to the fixed point. This is because in order to get an oscillatory solution, the trajectory has to pass through region (3), and in the usual case, the fixed point for this region is in region (1). Thus with all the other linear systems having a similar behavior, one should expect the trajectory to enter region (3). Region (3) has some complex eigenvalues so the trajectory can leave (3), but given the convergence behavior of the oscillator to its limit cycle, one would expect the system not to get stuck somewhere on the boundary of region (3), but converge to the fixed point.

The transient of the system with and without constant input is shown in figure 6-7, showing the convergence to the point $(-0.67, 0, 0.33, 0.33)'$.

Incidentally, this result gives theoretical backing to the use of a high pass filter to remove the offset from any input signal (see chapter 3). The filter effectively removes constant inputs, allowing the oscillator to oscillate. This greatly increases the robustness of the oscillator in practice.

6.5.3 Example: Mass-spring system

Consider the mass spring system

$$G(s) = \frac{k}{ms^2 + bs + k} \quad (6.33)$$

where for example $m = 1$, $k = 50$, $b = 4.24$, giving a system with natural frequency 7.07 rad/s, and damping factor $\zeta = 0.3$.

If this is connected to an oscillator, there are two unstable fixed points at $(0.2, 0.2, 0.2, 0.2, 0, 0)'$. Since the overall system is bounded, this implies that the coupled system will have a oscillatory but




Figure 6-7: SIM Transients of the system with a constant input. The input changes the system so there is no longer a unstable pole in region (1), but a stable fixed point at $(-0.67, 0, 0.33, 0.33)$, to which the trajectory (dash-dot) converges. Without the input (solid line) the trajectory converges to the usual limit cycle.

not necessarily periodic solution. In practice, the system appears to always converge to a limit cycle, as shown in figure 6-8. The local stability of this cycle is analyzed in section 6.8.1.

6.6 Prediction of entrainment

If by examining the boundedness and fixed points an oscillatory but not necessarily periodic solution is predicted, then describing function analysis can be used to further predict the existence of a limit cycle. This is an approximate method, but is quick to apply and useful in practice.

Describing function analysis applied to the oscillators was presented in chapter 3. To recap, the oscillator transfer function $N(j\omega, A)$ is measured at a range of frequencies (ω) and input amplitudes (A) and compared to the transfer function $G(j\omega)$ of the driven system. Since the two systems are connected in a loop (as shown in figure 6-9), the condition for oscillations is that the loop gain is unity.

$$N(j\omega, A)G(j\omega) = 1 \quad (6.34)$$

Solutions to this equation are calculated graphically, by plotting $G(j\omega)$ and $1/N(j\omega, A)$ on the complex plane, and finding intersections at the same frequency. At the solution frequency ω_f and amplitude A_f , the magnitudes of the two transfer functions multiply to one, and the phases are equal and opposite:

$$|N(j\omega_f, A_f)||G(j\omega_f)| = 1 \quad (6.35)$$

$$\angle N(j\omega_f, A_f) + \angle G(j\omega_f) = 0 \quad (6.36)$$

A sensible heuristic to determine whether the oscillator and a system can have an entrained solution is whether they produce compatible phases, such that equation (6.36) can be satisfied. This




Figure 6-8: SIM Two transients of the oscillator driving the mass spring system in equation (6.33). The plot shows the 6-dimensional state space projected onto the $x_1 - x_2$ plane. Both transients converge to the limit cycle.




Figure 6-9: Figure showing the oscillator tightly coupled to a driven system. Describing function analysis is used to compute the transfer function $N(j\omega, A)$ for the oscillator as a function of frequency ω and input amplitude A . The transfer function for the system $G(j\omega)$ is easily calculated if the system is linear. Since the two systems are connected in a loop, the condition for oscillations is that the gain around the loop is unity, or $N(j\omega, A)G(j\omega) = 1$.

is because the range of phases across the oscillator cannot be changed by varying parameters. The range is approximately $180 \leftrightarrow 50$ degrees, and so also the range $-180 \leftrightarrow -50$ if the output is simply inverted (multiplied by a negative output gain h_o , see equation (6.25)). Thus if the system is stable, and has a frequency response with phases in this range it is possible for there to be a limit cycle oscillation.

If the phases do match, there will probably not be a limit cycle for all values of the oscillator parameters. It is possible for the system to not be entrained, with the oscillator oscillating at its natural frequency, or partly entrained because the input amplitude to the oscillator is too low to sustain entrainment. Both of these can be diagnosed using the describing function plot, by determining whether the plots overlap.

If the plots do not overlap, but the phases are in the correct range, then the oscillator the parameters can be easily tuned (as described in chapter 3) to move the $N(j\omega, A)$ plot so that an intersection and limit cycle is possible.

Using the eigenvalue/fixed point analysis described above together with the describing function analysis gives a powerful set of tools to determine the final motion. The eigenvalue analysis shows the overall quality of the solution (oscillatory or not), and the describing function analysis predicts whether there is an entrained limit cycle solution. This is illustrated in the following example.

6.6.1 Example: Non-minimum phase system

If the non-minimum phase system (a system with zeros in the right half plane)

$$G(s) = \frac{8(s-1)}{s^2 + 6s + 18} \quad (6.37)$$

is connected to the oscillator, the overall system is bounded because $G(s)$ is stable, with poles at $s = -3 \pm 3j$. Examining the fixed points shows that there are two unstable fixed points. This suggests that the final solution will be oscillatory, but not necessarily periodic.

Plotting the describing function for the oscillator will predict whether a limit cycle oscillation will occur. The phase of $G(j\omega)$ is in the range $30 \leftrightarrow -60$ degrees, which overlaps with the oscillator phase, so a limit cycle is expected to be possible.

Figure 6-10 shows the describing function plotted together with the Nyquist plot of $G(s)$. The plot shows $1/N(j\omega, A)$ and $G(j\omega)$ plotted on the complex plane, and also $-1/N(j\omega, A)$, which is the frequency response of the oscillator when the output is reversed (output gain $h_o = -1$). The graph shows two possible solutions, one for each value of output gain. The prediction in this case was correct, with limit cycles for each solution as shown in Figure 6-11.

6.7 Local stability of the limit cycle

This section examines the local stability of the oscillator limit cycle. The previous sections have covered the general characteristic of the motion, determining whether the motion is oscillatory, and giving some heuristics to determine if the two systems will entrain with one another. Proving that a limit cycle exists in a given system is hard if the system has more than two dimensions. However, testing the local stability of a candidate limit cycle is relatively easy. The local stability can be predicted using the describing function analysis (chapters 3 and 4), and although it is possible in some cases to extend this work (Bergen et al., 1982), the result remains a prediction. This section shows how the local stability can be determined in a more rigorous fashion. The main reference used for this work was Gonçalves (1999).

To analyze the local stability of a limit cycle for a multi-dimensional system, the most suitable technique is to analyze the properties of a Poincaré map. The general idea is illustrated in Figure 6-12, see also Strogatz (1994), Khalil (1996). The n dimensional state space of the dynamical system




Figure 6-10: SIM Plot of $1/N(j\omega, A)$, $-1/N(j\omega, A)$ and $G(j\omega)$ on the complex plane. $G(j\omega)$ is the frequency response of the non-minimum phase system in equation (6.37), and is indicated with the thick line. The plot predicts limit cycle motions at two different frequencies. When $h_o = 1$, the solution is at the lower dot, at $\omega_f = 11.85$, $A_f = 0.31$. When the gain is negative ($g_o = -1$) there is a different solution at the upper dot with $\omega_f = 4.67$, $A_f = 0.67$.




Figure 6-11: SIM Two transients of the non-minimum phase system, with different values of oscillator output gain h_o . When the gain is positive (top graph), the oscillation is fast, with $\omega_f = 11.4$, $A_f = 0.32$, compared to the prediction (figure 6-10) of $\omega_f = 11.8$, $A_f = 0.31$. With a negative gain (lower graph), there is a slower oscillation, $\omega_f = 4.5$, $A_f = 0.7$, compared to $\omega_f = 4.67$, $A_f = 0.67$.




Figure 6-12: Poincaré map. The map turns the n dimensional dynamics $\dot{x} = f(x)$ into the $n - 1$ discrete system $y_{k+1} = g(y_k)$ by cutting the dynamics with a $n - 1$ dimensional hyper plane P . The limit cycle γ in the full dynamics appears as a fixed point of g , i.e. x^* . The local stability of the limit cycle can be determined by linearizing the Poincaré map around x^* , as shown in this section.

$\dot{x} = f(x)$ is cut by a $n - 1$ dimensional hyper plane P . This converts the system to a $n - 1$ dimensional discrete system $y_{k+1} = g(y_k)$. A limit cycle in the original system will then appear as a fixed equilibrium point of the reduced system. By examining the local stability of this equilibrium point, the local stability of the limit cycle is determined. The proof of this can be found in Khalil (1996).

For a PL system, the Poincaré section is naturally taken to be one of the switching surfaces, and the map can be calculated by following a trajectory through the limit cycle.

Considering the trajectory γ in figure 6-13 which starts on the switching surface S_0 , at state x_0^* and at time $t = 0$. The trajectory intersects the next switching surface S_1 at time $t = t_1^*$, and state x_1^* . The trajectory between these two surfaces is given by integrating (6.20):

$$x(t) = e^{A_1 t} (x_0^* + A_1^{-1} B_1) - A_1^{-1} B_1 \quad (6.38)$$

and so

$$x(t_1^*) = x_1^* = e^{A_1 t_1^*} (x_0^* + A_1^{-1} B_1) - A_1^{-1} B_1 \quad (6.39)$$

The trajectory continues through a number of switches, eventually returning to the point x_0^* after say k switches. The switching surfaces can be written as

$$S_i \equiv c_i x + d_i = 0 \quad (6.40)$$

Given this notation, a necessary and sufficient condition for the local stability of the limit cycle can be determined as shown below. This result was taken from Gonçalves (1999), where it is presented in more detail.

Proposition 6.4 *Consider the piecewise-linear system described above. Assume there exists a periodic solution γ with period t^* . Let $x_0^* \in S_0$ be the initial state that generates the periodic motion. The Jacobian of the Poincaré map is given by $W = W_k W_{k-1} \dots W_2 W_1$ where*

$$W_i = \left(I - \frac{v_i c_i}{c_i v_i} \right) e^{A_i t_i^*}$$

with $v_i = A_i x_i^* + B_i$. The limit cycle of γ is locally stable if and only if W has all its eigenvalues inside the unit disk.




Figure 6-13: The limit cycle trajectory γ starts in state x_0^* , and travels through the first linear system ($\dot{\mathbf{x}} = A_1\mathbf{x} + B_1$) until it crosses the first switching surface S_1 at time t_1^* . It continues into region 2, 3, etc. The limit cycle is complete after k switches, so that the state on the k th switch x_k^* is the same as the starting state i.e. $x_k^* = x_0^*$.

Proof: This proof is taken verbatim from Gonçalves (1999).

The proposition is proved by perturbing the initial state by a small amount, and measuring the size of that perturbation on the trajectory when it reaches the next switching surface. The perturbation is expressed as a linear function of the initial perturbation and higher order terms are ignored. The perturbation is followed around the limit cycle back to the starting surface in order to compute the linear part of the map.

If the starting point is x_0^* , then the point on the next switching surface S_1 is given by equation (6.39). The starting point is perturbed along the switching surface, i.e. $x(0) = x_0^* + \delta_1 x_0^*$, so that $c_1(x_0^* + \delta_1 x_0^*) + d_1 = 0$. The trajectory starting from this new position will intersect S_1 at time $t_1^* + \delta_1 t_1^*$:

$$x(t_1^* + \delta_1 t_1^*) = e^{A_1(t_1^* + \delta_1 t_1^*)}(x_0^* + \delta_1 x_0^* + A_1^{-1}B_1) - A_1^{-1}B_1$$

Taking the series expansion in $\delta_1 x_0^*$ and $\delta_1 t_1^*$ gives

$$\begin{aligned} x(t_1^* + \delta_1 t_1^*) &= x_1^* + e^{A_1 t_1^*} \delta_1 x_0^* + e^{A_1 t_1^*} (A_1 x_0^* + B_1) \delta_1 t_1^* + O(\delta_1^2) \\ &= x_1^* + e^{A_1 t_1^*} \delta_1 x_0^* + v_1 \delta_1 t_1^* \end{aligned} \quad (6.41)$$

where $e^{A_1 t_1^*} (A_1 x_0^* + B_1) = A_1 x_1^* + B_1 = v_1$. Since $x(t_1^* + \delta_1 t_1^*)$ is on the switching surface S_1 , $c_1 x(t_1^* + \delta_1 t_1^*) + d_1 = 0$. Ignoring higher order terms this gives

$$c_1 x_1^* + c_1 e^{A_1 t_1^*} \delta_1 x_0^* + c_1 v_1 \delta_1 t_1^* + d_1 = 0$$

which since $c_1 x_1^* + d_1 = 0$,

$$c_1 v_1 \delta_1 t_1^* = -c_1 e^{A_1 t_1^*} \delta_1 x_0^*$$

Assuming that the Poincaré map is continuous in the region of the solution implies that $c_1 v_1 \neq 0$, so that the previous equation can be written

$$\delta_1 t_1^* = -\frac{c_1 e^{A_1 t_1^*}}{c_1 v_1} \delta_1 x_0^*$$

This can now be replaced in equation (6.41), giving

$$\begin{aligned} x(t_1^* + \delta_1 t_1^*) &= x_1^* + \left(I - \frac{v_1 c_1}{c_1 v_1} \right) e^{A_1 t_1^*} \delta_1 x_0^* + O(\delta_1^2) \\ &= x_1^* + W_1 \delta_1 x_0^* + O(\delta_1^2) \end{aligned}$$

Similarly, for the next switching surface

$$\begin{aligned} x(t_2^* + \delta_2 t_2^*) &= x_2^* + W_2 \delta_2 x_1^* + O(\delta_2^2) \\ &= x_2^* + W_2 W_1 \delta_1 x_0^* + O(\delta_2^2) \end{aligned}$$

Repeating this until the trajectory arrives back at the original switching surface $S_k = S_0$, where $x_k^* = x_0^*$, gives

$$\begin{aligned} x(t_k^* + \delta_k t_k^*) &= x_k^* + W_k \delta_k x_{k-1}^* + O(\delta_k^2) \\ &= x_0^* + W_k W_{k-1} \dots W_2 W_1 \delta_1 x_0^* + O(\delta_2^2) \end{aligned}$$

This is an expression for the linear part of the Poincaré map, so looking at the eigenvalues of $W = W_k W_{k-1} \dots W_2 W_1$ will determine the local stability. Since the system is discrete, this corresponds to the eigenvalues lying within the unit disk.

■

6.8 Local stability for the oscillator

To show local stability for the oscillator limit cycle, the values of W_i in Proposition 6.4 can be calculated. The limit cycle was simulated to recover the values of the states at each switching surface (x_i^*) and the times between switches (t_i^*), as shown in figure 6-14.

The states are

$$x_0^* = \begin{pmatrix} 0.3991 \\ 0.1415 \\ -0.0000 \\ 0.2587 \end{pmatrix}, x_1^* = \begin{pmatrix} 0.3414 \\ 0.3535 \\ 0.0000 \\ 0.0612 \end{pmatrix}, x_2^* = \begin{pmatrix} -0.0000 \\ 0.2587 \\ 0.3991 \\ 0.1415 \end{pmatrix}, x_3^* = \begin{pmatrix} 0.0000 \\ 0.0612 \\ 0.3413 \\ 0.3535 \end{pmatrix} \quad (6.42)$$

with times

$$t_0^* = 0.2884, \quad t_1^* = 0.1604, \quad t_2^* = 0.2884, \quad t_3^* = 0.1604 \quad (6.43)$$

These values result in the eigenvalues of W :

$$\begin{pmatrix} 0.1265 \\ -0.1196 + 0.0412j \\ -0.1196 - 0.0412j \\ 0 \end{pmatrix} \cdot 10^{-3} \quad (6.44)$$

which are all well within the unit disk, proving that the limit cycle is locally stable. For different oscillator parameters this process would have to be repeated, but one would expect similar results.




Figure 6-14: SIM Oscillator limit cycle plotted in the $x_1 - x_2$ plane. The points where the cycle crosses the switching surfaces $x_1 = 0$ (horizontal line), and $x_2 = 0$ (vertical line) are marked with $\bullet$'s. Using these points, and the times for the trajectory to go between them, the local stability of the limit cycle can be determined.

6.8.1 Example: Mass-spring system

A similar procedure can be followed when the oscillator is driving a system, using the 8 different linear systems described in section 6.5. For example, taking the mass-spring system analyzed in section 6.5.3, and plotting the limit cycle in x_1, x_2, g space gives figure 6-15. Calculating W_i as before, and calculating the eigenvalues gives

$$\begin{pmatrix} 0.1049 \\ 0.0028 - 0.0015i \\ 0.0028 + 0.0015i \\ -0.0036 - 0.0024i \\ -0.0036 + 0.0024i \\ 0 \end{pmatrix} \quad (6.45)$$

which are all inside the unit disk, so that the limit cycle is proved to be locally stable.

6.9 Global stability

The previous sections in this chapter have described anecdotally the convergence of the oscillator to a limit cycle, have shown some qualitative results for the oscillator behavior (oscillatory or not), and have shown the local stability of the limit cycle. The part that remains is to show that the system does in fact converge to a limit cycle, or alternatively to quantify the range of states over which this contraction occurs.

If the oscillator system had only two states, then it would be possible to take the quantitative results and go directly to a limit cycle using the Poincaré-Bendixson theorem (Strogatz, 1994, p. 203). However, since the oscillator alone has four state variables, and when coupled to systems




Figure 6-15: SIM Plot of limit cycle for mass spring system and oscillator. The cycle is plotted in the three switching parameters x_1, x_2, g , with the switching states marked.

the number of states increases further, this approach will not work. Unfortunately there is no generalization of this theorem to higher dimensions, with proof of convergence being difficult.

There are a number of avenues which are worth pursuing. The most promising is to establish a contraction region for the limit cycle by quantifying the size of the non-linear part of the Poincaré map (Gonçalves et al., 1998). By writing the Poincaré map as a linear part and a non-linear part (instead of ignoring the higher order terms as in section 6.7), and finding the size of the non-linear part in terms of the size of the perturbation, the contraction region can be determined. There are some technical difficulties associated with systems with many switching surfaces, although these are not thought to be insurmountable (Gonçalves, 1999).

Other approaches include contraction analysis (Lohmiller and Slotine, 1998) although that method has difficulties with systems with complex eigenvalues (such as the oscillator), or the graphical technique presented by Pettit (1995). This method converts the piecewise-linear system into a graph where the nodes on the graph are the regions on the switching surfaces where the flow is going in the same direction, and the links of the graph show where that flow goes. It thus gives a graphical representation of the structure of the system. This does not show convergence but might give a way to determine global entrainment to a bound calculated perhaps as above.

Proving convergence to a particular limit cycle is difficult, as is finding out more precisely what systems will and will not have limit cycle solutions, Methods based on the small gain theorem seem appropriate here, although are also difficult to apply in this case (Megretski, 1997).

6.10 Conclusion

This chapter has addressed some theoretical questions about the oscillator, and its behavior coupled to systems. The first part of the chapter considered the general behavior of the system, showing that the oscillator state is bounded when connected to a bounded system, implying that the system state will not become unbounded if the driven system is stable.

By looking at the system as a piecewise-linear system it is easy to calculate the fixed points for the non-linear system, which can be used together with the bounded result to predict the final

behavior. This prediction is somewhat weak: the final motion can be predicted to be oscillatory, but the existence of a limit cycle cannot be shown. Combining this information with describing function analysis allows prediction of the limit cycles. This is useful and easy to apply in practice.

The last part of the chapter addressed the specific behavior of the limit cycle, giving a method to determine local stability. The local stability of a number of examples was determined.

The difficult part that remains is to bridge the gap between the specific local results and the imprecise global results with a proof that the system converges to the limit cycle. This was shown to be difficult due to the high dimensionality of the system, and possible methods to determine this were presented. Anecdotally the system is well behaved converging very quickly to the limit cycle, although this is not proven.

In all, the results presented in this chapter are useful for design, determining what systems will produce oscillatory behavior. They extend the analysis in previous chapters with a rigorous determination of whether oscillations are possible, and if they occur, whether they are locally stable.

Chapter 7

Examples

This chapter describes a number of example applications using oscillators which further demonstrate their versatility. The tasks differ from the ones described in previous chapters by consisting of a single motion, rather than continuous motions. The oscillator dynamics are used to produce these motions, with feedback being used to modify the oscillator outputs in the usual way.

The first half of the chapter describes using the oscillators for a task which is partly rhythmic and partly discrete: the task of drumming. The “beat” is a continuous signal, while the individual drum beats are discrete. The oscillator system uses auditory feedback of the sound of the drumming to create and maintain a steady beating rhythm. The chapter shows how the oscillator properties make it easy to beat in time with another drummer, and construct a variety of different rhythms, using both one and two arms.

The second half of the chapter describes the tasks of throwing, hitting and hammering. The oscillators are used to drive the arm back and forward in a single motion to achieve a good throw or a powerful hit. These motions exploit the natural dynamics by storing and releasing energy in the joints of the arm, and also exploit the oscillator properties, the oscillators responding to the system dynamics. The effect of feedback in these discrete tasks is less clear than in the continuous case, although it appears to have approximately the correct effect: making the arm move faster.

7.1 Drumming

The first example in this chapter is drumming. This task combines elements of the rhythmic motions described in previous chapters with discrete motions. The overall rhythm or beat is a continuous signal, while the individual hits on the drum are discrete. The example also shows how other sensory modalities can be used as feedback to the oscillators. In this case auditory feedback (the sound of the drum) is used to modify the oscillator behavior. This has interesting and useful properties as described in this section.

The overall setup is illustrated in figure 7-1, and a picture of the robot drumming is included in figure 7-2. An oscillator is used to drive either the wrist or the elbow of the robot to hit the drum. The master-slave oscillator system described in section 5.3 was also used to drive the elbow and wrist joints as a single unit. The sound of the drumming was recorded by a microphone, processed and applied as an input to the oscillators. The auditory processing is simple, with the raw auditory signal sampled at 8kHz, and thresholded to find the sound of the drum. This is then low-pass filtered and subsampled to reduce the sampling rate to around 50Hz, an appropriate rate for the oscillators.

The drumming action of the robot is accomplished by moving the arm up and down, with the drum being hit on the down stroke. When the arm is still, the drum stick is held slightly above the drum, so it actually hits the drum on the follow through. This is similar to the hitting and throwing examples in the previous sections. Human drummers grip the drumstick lightly and exploit the dynamics of the stick on the drum to achieve high frequency rolls. The light grip has the added advantage of reducing the inertia of the stick so it can bounce on the drum and make a good note. To achieve this on the robot, a simple passive hand was built to hold the drumstick, as shown in




Figure 7-1: Schematic of drumming. The system is set up so that the drumming task is an auditory-motor loop, where the action of the oscillators is to drive the arm to hit the drum, and the sound of the drum is recorded and used as an input to the oscillator. Details of the auditory processing are included in the text.




Figure 7-2: Picture of Cog hitting a drum. The robot uses its elbow and wrist joints to move the arm up and down, so hitting the drum.




Figure 7-3: Drumstick holder. This special purpose hand was made for the robot to hold the drumstick. The stick swings freely on a pivot, rotating between two pads made of either rubber or felt. The whole hand is rotated by the wrist joint of the robot. When the hand is moved up and down, the stick swings on the pivot, and bounces on the drum. The interaction of the drum, stick and pads can be used to get single or multiple hits on the drum, exploiting the bounce of the drum head.

figure 7-3. The stick is pivoted so it can swing freely, its motion damped by two felt or rubber pads. By using a piece of tape to modulate the free motion of the stick, the number of bounces of the stick on the drum could be controlled. This passive system worked well. Other researchers have considered active low impedance holders for drumsticks with similar results (Hajian et al., 1997).

The overall system (shown in figure 7-1) forms a closed auditory-motor loop. There are considerable delays in the system, mostly in transferring the auditory signals to the oscillators, and from the command to move the arm to the actual hit on the drum.¹ These delays combine with the oscillator dynamics to create a limit cycle drumming solution with a steady beat, as shown in figure 7-4. The main effect of the feedback is to make the drumming faster, and because the stick is moving faster the sound is slightly louder. The auditory signal is a series of spikes, but is strong enough to cause entrainment of the oscillators.

By entraining the oscillators to the sound of drumming, the oscillator system produces a steady beat without requiring a timer. This is important for the particular implementation of the robot arm, because there is no global clock. The system does not require any model of the delays in the system, and in fact requires the delays in order to work. The tuning of the system is also simple, because of the robustness to the value of the input gain (see chapter 3). As long as the gain is large enough to cause entrainment, its exact value is not important.

Since the oscillators use the raw auditory signal as an input it can respond to other noises, such as an additional drum beat. The oscillator adjusts its phase and frequency to beat in phase with the extra beat. This effect is shown in figure 7-5. One arm of the robot was used to produce a constant beat, and the other arm was controlled with an oscillator. The effect of the feedback is make the oscillator driven arm hit in phase with the extra beat.

Hitting in phase is the only possible solution because the robot “hears” both beats. The stable limit cycle without the extra noise has the arm hitting just before the noise, so the stable entrained cycle with the extra beat must have the same characteristic. Thus the only stable cycle is one where the oscillator hits in phase with the extra beat. Because the speed of the limit cycle is determined by the delays in the system, the range of frequencies with exact synchronization is quite small (in the range 9.3 to 10.3 rad/s about a “natural” frequency of 9.8 rad/s). If the oscillator only responds to the extra beat and ignores the sound that it makes then it can entrain over a wider range of frequencies, although the hits are not exactly in phase. Changing time constants and delays in the

¹The architecture for the system is described in Appendix A.




Figure 7-4: **REAL** Plot showing the effect of feedback on the drumming system. The plots show the auditory signal (top trace in both plots), the command to the joints (solid line) and the actual joint motion (dashed). The drum is hit when the arm is moving down. The main effect of the feedback is to speed up the drumming, which has the consequence of making more noise. The final speed of the drumming is determined by the delays in the system, which are dominated here by the delay from the command to the arm to the arm actually moving. The shape of the auditory signal is quite different from a sine wave, yet the oscillators can still entrain to it.




Figure 7-5: **REAL** Two time plots for the oscillator entraining with an external beat. One arm of the robot was used to create a constant beat (dash-dot line), while the other arm used the auditory signal (top trace) to control the arm motion (solid line). The feedback gain is indicated by the dashed line. When the feedback is on, the oscillator drives the arm in phase with the external beat, and when it is off there is no coordination.




Figure 7-6: REAL Plot showing examples of different rhythms possible using counting. The top plot shows the auditory signal (top trace), the motion of the dominant arm (middle trace), and the motion of the second arm which is hitting every second beat. The lower plot shows the same, but the second arm is hitting every fourth beat. Since the underlying oscillator for the second arm is entrained to the auditory signal, the two arms hit in phase. The up and down motion of the arm is visible in the lower traces, the drum being hit on the follow through, after the arm has been lifted up and down.

system, either by introducing new delays, or changing the stiffness of the arm can change the limit cycle frequency.

This example illustrates the power of coupling systems using feedback from the environment. It is only because the sound of the two drums is combined by being detected with a single microphone that the oscillator beats in phase.

7.1.1 Counting

This section describes how the drumming can be extended to using two arms, and how the auditory-motor loop can be used to synchronize the arms with one another, even when the arms are beating at different rates.

Synchronizing two arms using the oscillators is easy, by using an oscillator for each arm and feeding both oscillators the same auditory signal. Since the coordination of the system is determined by how the oscillators entrain with the auditory signal, both arms will make the same motion, and so hit the drum in phase. The architecture of the system (see Appendix A) has the controllers for each arm in separate processors, with low bandwidth communication between the processors. Applying the feedback separately allows both arms to be coordinated with one another using a distributed control system.

Getting the arms to hit at different rates can be achieved by using something analogous to counting. Say for example one arm is hitting every cycle, and we want the other arm to hit on every other cycle. A simple way to achieve this is to keep the oscillator for the second arm synchronized to the auditory signal (so oscillating in phase with the first arm), but only use its output to drive the arm every other cycle. The counting is easily implemented by counting the cycles of the oscillator




Figure 7-7: REAL Plot showing two stable limit cycles when the arms both hit on every other beat. The top graph shows the slow cycle, the arms hitting together, and the lower graph shows the arm hitting alternately. The speed of the faster cycle is the same as if it were made by one arm drumming on every beat.

since the last hit. The robotic performance of this rhythm is shown in figure 7-6.

An interesting limit cycle is created when both arms hit on alternate cycles. This results in two auditory-motor limit cycles, one slow cycle where the arms beat in phase, and one faster one where the arms beat alternately. These are shown in figure 7-7. Both cycles appear to be stable, and can be switched from one to another by disturbing the system with an extra auditory signal as shown in figure 7-8. The oscillators for the two arms are completely independent, but are coordinated with one another using the common auditory signal.

7.1.2 Hitting on the offbeat

Another extension of the system allows the arms to hit between the beats. This is easily achieved by inverting the sign of the feedback signal to one of the arms. This will cause that arm's oscillator to entrain out of phase with the other arm, and so hit the drum halfway between the other arm's beats. Unfortunately, hitting in this way creates a more complex auditory signal, to which the oscillators cannot easily entrain. The system needs to ignore the extra beats, and only pay attention to the dominant beat.

One solution to this problem is to improve the auditory processing to ignore the extra beat. There is a considerable literature on beat tracking (e.g. Scheirer (1998)), and the task here is simplified by the clean sound of the drum.

An alternative and simpler solution is to inhibit the auditory feedback except when a sound is expected, i.e. only listen for a short period after the motion of the dominant arm. This was implemented on the arm, and was successful at producing a range of rhythms, with the entrainment of the system undisturbed by the offbeats. Rhythms hitting every fourth, second and alternate beats are shown in figure 7-9. It is also easy to change the rhythm produced by the arms, by altering the number of beats "counted", as shown in figure 7-10. Since the system only listens for a short period after the motion of the dominant arm, changing the second arm does not disturb the steady




Figure 7-8: REAL Plot showing how an auditory disturbance can alter the limit cycle found by the oscillator when each arm hits on alternate beats. The disturbance was from a third drumstick banging the drum, which caused the system to switch from fast to slow (top graph) or slow to fast (lower graph).

entrained beating pattern.

7.1.3 Hitting more than one drum

It is also possible to use an extra oscillator to enable the system to hit more than one drum. As described in chapter 3, the setpoint for the whole arm is generated by adding the normal oscillator output y_{out} to a fixed posture θ_p :

$$\theta_v = \theta_p + y_{out} \quad (7.1)$$

The posture does not need to be fixed, and can be generated by the output of another oscillator. If that oscillator provides a series of steps, then the arm will move backwards and forwards between two locations. If this movement is superimposed on the drumming, the arm will hit two different drums. The synchronization between the drumming and the posture change is obtained by using the auditory signal to entrain the posture oscillator.

This method was used to hit a drum and a cymbal as shown in figure 7-11. Unfortunately, the auditory feedback was not clean enough in this case to get the oscillators to entrain. A fake sine wave signal was used instead, which was easier to entrain with.

The oscillator system was used for a variety of percussion instruments, including tambourines, maracas, cow bells etc. For some of these (e.g. the maracas), the auditory signal was sufficiently complex for the simple auditory processing to not be able to extract a signal which the oscillators could entrain with. In those cases a sine wave was used to produce the rhythms.

7.2 Discrete from rhythmic

The discrete motions which are considered in the rest of this chapter are single, dynamic motions which contrast with the continuous dynamic tasks in the previous chapters. For example, a golf




Figure 7-9: **REAL** Three different rhythms which use the motion of the dominant arm to inhibit the auditory processing. The plots show the raw auditory signal (dashed), together with the signal applied to the oscillators (top solid trace), and the two arm motions. Since the oscillator only “listens” for a short period (approx 0.25s) following the motion of the arm, the sound of the second arm drumming on the offbeat is ignored. This allows the system to remain entrained.




Figure 7-10: REAL Switching between different rhythms. By changing the number of beats which are “counted” by the oscillator, the system can change the rhythm produced. The top plot shows the second arm changing from hitting every other beat, to hitting every second, and the lower plot shows the change from every fourth, to alternate beats. The auditory signals are shown as described in figure 7-9.

swing is a discrete dynamical motion, while the actions during eating are also discrete but less obviously dynamic! The commands required for this kind of motion are pulses and steps, which can be obtained by starting and stopping the oscillator at peaks in its output, as shown in figure 7-12. Pulses are created by stopping at peaks of the same sign, while steps by stopping at peaks of opposite sign.²

Feedback can be applied to this discrete oscillator system exactly as in the rhythmic case. The effect of feedback is more complex because the final motion is not a converged limit cycle, but the transient trajectory towards that limit cycle. The oscillator entrains quickly, which implies that the feedback will have an effect. Because the oscillator tends to find resonant-like motions when run continuously, one would expect the transient to have the quality of “moving towards resonant-like motions”, or injecting energy into the system. Like the continuous case, the feedback should make the oscillators sensitive to the arm dynamics, and create movement where an appreciable amount of energy is stored and released in the arm springs.

Schöner (1990) also suggested using a rhythmic system to produce discrete motions. His system was perhaps more mathematically elegant than the solution presented here, although it was sensitive to parameter values and not used for any practical purpose.

7.3 Throwing

The throwing task was implemented by using the oscillators in discrete mode to drive the joints of the arm. There was no connection between the joints except through mechanical coupling, and the feedback to the oscillators was from the joint angles of the arm.

Figure 7-13 shows a sequence through an underarm throwing action implemented on the robot.

²The peaks can be easily detected by checking the zero crossings of the rate of change of oscillator output $\dot{y}_{out}$.




Figure 7-11: REAL Four stills from a video of the robot drumming. The pictures are in no particular order. The robot uses its right arm to hit the drum, and its left arm alternates between the drum and the cymbal.




Figure 7-12: Schematic of discrete motion. A single cycle of the oscillator can be used to create single motions by detecting peaks in the oscillator output and stopping and starting the integration. The thin line shows the usual oscillator output, and the thick line shows how this can generate (a) a pulse and (b) a step.




Figure 7-13: REAL Series of stills from sequence of the arm throwing a ball. The picture sequence is left to right and then top to bottom. The overall time of the throw is about 1.2 s, although these pictures are not equally spaced over that time. The arm holds the ball up, swings the elbow and shoulder back, and then swings the straight arm forward and releases the ball.




Figure 7-14: **REAL** Plot of joint angles (θ) and setpoints (θ_v) during throwing. Two cases are plotted, using thick lines indicate when feedback is on, and thin when it is off. The plots show the shoulder moving back and forward and the elbow moving straight back. Some photographs of this motion are shown in figure 7-13. The ball is released at the end of the shoulder motion, as indicated by the vertical lines, although the ball does not actually leave the hand until about 0.5 seconds later. The most obvious effect of the feedback is to increase the speed of the setpoints to the joints (solid lines), and thus change the speed of the arm. The feedback also makes the hand release earlier in the motion, which throws the ball forward rather than up.

The postures of the arm were chosen to be similar to a human action. The arm starts holding the object up, then the elbow is driven back using a step command, and the shoulder back and forward using a pulse. This has the effect of swinging the arm back firmly, so hyperextending the shoulder joint, and storing energy in the shoulder spring. The arm comes forward straight and releases the object. Since the robot has no hand, the object is gripped between the wrist and a sheet of metal attached to the forearm. The wrist releases grip when the shoulder oscillator has finished its pulse.

The effect of the feedback on the throwing action is not very great, at least for the example throw shown in figure 7-14. The feedback makes the arm move faster, and also releases the object earlier in the throw. When using feedback, the wrist starts to release the ball when the arm is extended all the way back, but the ball is carried by the hand and not actually thrown until some time later (see photographs in figure 7-13). Without feedback, the ball is released much later. Releasing earlier may give a better follow through on the throw, and contributes to throwing forward rather than up.

In any case, the effect of the feedback is not large. One might expect it to be greater, and this is an area for further work. The oscillator's ability to quickly adapt to the dynamics of the arm might make it particularly appropriate for throwing tasks, by adjusting automatically to the mass of the object. Work by Bingham et al. (1989) suggests that humans heft objects to determine their properties before tuning the stiffnesses of their arms to achieve a powerful throw.

Throwing objects is sensitive to the release of the object, which was somewhat arbitrary in this implementation. Even a good swing can produce a poor throw if the release is fumbled. An overarm throw was also implemented, using four joints of the arm, with a sequence shown in figure 7-15. For this case, the shape of the hand was redesigned to give a better throw.




Figure 7-15: REAL Overarm throw sequence, with pictures taken every 0.2 seconds. The sequence goes from left to right and top to bottom. Two shoulder joints and two elbow joints are used to produce this motion, all activated with pulses except for one elbow joint which was driven using a step. Throwing is sensitive to the release of the object, and in this case a redesigned hand was used to improve the throw.

7.4 Hitting

The task of hitting or hammering requires a similar motion to throwing but avoids the difficulty of releasing the object. Two applications are described here, with the robot hitting a “punchbag” and hammering a nail. These motions exploit the natural dynamics of the arm, using energy stored in the arm springs and gravity to produce effective hitting. They also exploit the ability of the arm design to withstand considerable shock loads.

The punchbag was constructed to measure the energy transferred from the arm to the target for tuning purposes. It consists of a rubber ball on a sprung boom. When the arm hits the ball, it causes the boom to rotate against the spring. By measuring the maximum deflection of the boom, the energy of the impact can be determined.

Figure 7-16 shows a sequence from the oscillator hitting the punch bag. The elbow and shoulder are activated by oscillators with the wrist having a low stiffness. The energy stored in the wrist is visible in the lower left frame, it being swung back by the action of lifting the arm, before rebounding onto the punchbag. The hitting action is much stronger when the arm is lifted up and then down, rather than just going down, due to this extra stored energy. The hit is also more powerful when the wrist stiffness is low, the mass of the hand acting like a hammer.

Figure 7-17 shows the shape of the hitting motion in the vertical plane, computed by using the joint angles and an approximation to the forward kinematics of the arm. When the feedback is on, the arm is moved up and then directly down to hit the punchbag (represented by the black rectangle). Without feedback the motion is more curved and much slower (the arm is drawn at 0.02 second intervals). This is shown more clearly in figure 7-18 which plots the height of the hand against time. The position of the punchbag is shown with the horizontal line, and the increased speed at impact is apparent.

The same configuration was used to hammer nails into wood, as shown in figure 7-19. Unfortunately no visual guidance was used so careful alignment of the nail with the hammer was required!




Figure 7-16: REAL Four stills from a video of the robot hitting the punchbag. The sequence goes from left to right and top to bottom. The arm starts resting on the punchbag, and is raised up as shown at the top. The arm whips down, overextending the wrist joint (bottom left) and hitting the target powerfully. The shock tolerance of the arm is exploited to produce this movement without damage.




Figure 7-17: REAL Motion of the arm in the vertical plane during hitting. The target is the punchbag which is indicated by the rectangle. The trajectory of the hand is marked, together with the position of the arm at 0.02 second intervals on the downward stroke. In the left hand graph, feedback is used and the arm moves faster and more directly to the target, as opposed to the non-feedback case in the right hand graph. The hyperextension of the wrist and its contribution to the hitting is also apparent when feedback is used.




Figure 7-18: REAL Graph of the height of the hand plotted against time for hitting the punchbag. The full trajectory is shown in figure 7-17. The hand moves up at about the same speed regardless of the feedback, but with feedback it accelerates down rapidly, with a much higher speed at impact (where the lines cross the starting height, marked by the horizontal line). The impact does not affect the arm motion a great deal because the inertia of the punchbag is small, and because it is sprung and so presents a small resisting force.




Figure 7-19: REAL Four stills from a video of the robot hammering a nail. The sequence goes from left to right and top to bottom. The elbow and shoulder joints drive up and down to give a powerful hit with the hammer. The motion causes the hammer to be swung fairly high, and extend the wrist joint (lower left picture). This stored energy is then released into the nail (lower right).

The arm swings the hammer up, storing gravitational potential energy in the hammer, and storing energy in the springs of the arm. These are both released into the nail.

As with throwing, the effect of the feedback is predominately that of changing the speed of the motion, which could easily be achieved by using a pulse with no feedback. The oscillator does make the motion faster, which is the correct action for both tasks. Anecdotally the feedback was useful from a safety point of view. If the arm was restrained, the oscillator was affected by the feedback from the slowly moving arm, causing it to drive the arm in a less violent manner. As with the throwing task, more research is needed to elucidate the effect of the feedback, and determine whether it is useful.

7.5 Conclusion

This chapter has shown that the oscillators are versatile, and can be used for a wide range of tasks, both rhythmic and discrete, and use other sensory modalities to provide feedback and entrainment.

The oscillator solutions in these examples are all adaptive, altering the drive to the arm depending on the motion of the arm, or in the case of the drumming, the sound created by the arm. For the discrete tasks the nature of adaptation is not very clear, although it appears to have roughly the right effect. The feedback causes the arm to move out and back faster, so creating a more powerful throw or a more powerful hit. The effect of the feedback is more clear for the drumming, allowing the arm to synchronize with an external beat, and create a stable drumming rhythm.

The drumming example shows the simplicity of exploiting the natural dynamics of the system for a complex task. The arm synchronizes with an external beat because the sound of that beat and the arms own drumming are combined by being heard by the same microphone. The synchronization occurs even though the overall system is not calibrated, has unknown delays and uses a low gain controller to move the arm. It is synchronized because hitting in phase is the limit cycle solution of the auditory-motor loop set up by the oscillators. The system requires the delays and low stiffness in order to work, which reduces the pressure on making efficient auditory processing or having high bandwidth connections between the system processors.

The examples also illustrate the ease of implementation of the oscillators for these tasks. The oscillators can entrain with a range of input signals, and are robust to changes in the parameters. Once the input gain is large enough that the system is entrained, changing it has little effect on the overall system. The method of inhibiting the auditory signal when it is not expected is another simple and robust method to obtain complex behavior.

The particular implementation of the arms uses a number of processors with low bandwidth communication between them. The entrainment properties of the oscillators make them appropriate for this distributed system, and more robust than attempting to create a synchronized global clock.

Finally, it is interesting in the drumming example that the two phases that were easy to create were in phase and out of phase. These were made by simply flipping the sign of a feedback signal, and modifying the auditory input. Humans also appear to find these phases easy to produce, and others more difficult. There is considerable literature (e.g. Haken et al. (1985), Kay et al. (1987)) which shows that humans have a preference for these phases. The preference varies with frequency, with out of phase motions being only possible at low frequencies, and in phase motions possible at all frequencies. It would be interesting to explore the frequency dependence of the oscillator system from this respect.

Chapter 8

Conclusion

8.1 Further work

There are many topics arising from the thesis which deserve further study. These include:

- **Complex timing constraints.** There exist tasks which fit into the general framework exploited by the oscillators but which have complex timing constraints. For example, consider the cases of a Yoyo and a Slinky toy. Both exert forces on the arm, and in both cases the steady state motion requires coordination between these forces and the arm motion. This suggests that the oscillator solution should work well. While the oscillators work well with the Slinky, they do not appear to work with the Yoyo. This is because the coordination required of the arm is more complex and dependent on timing for the Yoyo case. This is true both in steady state, and particularly in the startup transient. Extending or manipulating the timing of the oscillator system would allow it to perform tasks with complex timing constraints.
- **Potential fields.** The idea of applying potential fields to the arm was mentioned in chapter 5 and deserves further study because the work was restricted to a simple example due to hardware and time limitations. One interesting aspect of the potential field is that it need only be strong enough to cancel the undesirable natural dynamics of the arm. This means that it need not be stiff (which facilitates implementation) and need only be specified where necessary. This makes it particularly appropriate for incremental learning through imitation. One might envisage a system where a human guides the arm through a desired motion, and then corrects the motion as the robot repeats it. The correction would be implemented by increasing the strength of the field locally, so constraining the arm motion at that point. Learning by imitation has recently become an intensive field of study because it promises to provide simple and flexible ways to control robots. Recent examples include Miyamoto et al. (1996), Miyamoto and Kawato (1998), Brooks et al. (1999) and Schaal (1999). The robustness and ease of implementation of the oscillator systems should make them appropriate for imitating rhythmic and dynamic discrete tasks.
- **Sequences and compound motions.** One problem not addressed in the thesis is how to switch among a number of oscillators to produce compound motions or sequences. This is another possible method to increase the complexity of tasks that are possible using oscillator control. Switching between controllers raises a number of interesting questions:
 - How should the oscillators be synchronized across the switch? and is that necessary?
 - When should the switch be made? (after a certain amount of time, or dependent on the state)
 - How should the subcontrollers be designed?
 - Will the final system be stable?

There is theoretical work on hybrid systems (e.g. Branicky (1994)) which is relevant here. One robotic example which combines switching and exploiting the natural dynamics is the work of Kuniyoshi and Nagakubo (1997) where a simulated humanoid was controlled to stand up from a supine position using a sequence of simple controllers.

- **Discrete motions.** The oscillator behavior for single motions was described in chapter 7. An analysis method analogous to the describing function analysis developed for the rhythmic case is required to gain improved insight into the discrete behavior. This would hopefully shed light on the exact role of the feedback in these motions. One idea for an evaluation criterion is to determine the energy imparted to a robot link by a single cycle of the oscillator. This could then be compared to the optimal pulse, given the system stiffness and inertial properties.
- **Theoretical questions.** The theoretical results in chapter 6 indicate that many questions concerning the oscillators remain unanswered. These include results on the classes of systems that have limit cycle solutions when driven by oscillators and the global or semi-global stability of these limit cycles. The oscillators are a compelling system for further study in this respect. They exhibit a range of interesting behaviors, their piecewise-linear dynamics open them to a range of techniques, and any new theoretical results will have an impact on the practical design of systems.

8.2 Vision

This thesis has demonstrated an exciting range of tasks with a simple control system. While it has demonstrated the power both of exploiting natural dynamics and using oscillators to control arms, the method has a long way to go to reach the kind of versatility and flexibility of actions that humans routinely demonstrate. This section considers how to create that kind of performance, considering other ways to exploit the arm dynamics, exploiting the arm statics, tidying up details, and looking at the problem of learning.

- **Generalization.** This thesis has argued strongly for exploiting natural dynamics using the Matsuoka oscillator. This approach, while perhaps surprisingly versatile, is limited due to the dynamics of the oscillator. To achieve anything like human versatility other controllers and control methods need to be explored. Investigating other controllers which exploit natural dynamics should open up a variety of different tasks, for example discrete motions, slow motions, compound rhythmic motions, or asymmetric motions.

This thesis has concentrated on exploiting the resonant properties of the arm, when there are other aspects of natural dynamics which can be exploited to give wider ranges of tasks. These include the use of the skeleton for load carrying, and also the use of internal inertia for explosive tasks (as in the work of Brown (1994)).

Generalizing the ideas in this thesis to other controller types which can exploit the natural dynamics, as well as generalizing to other aspects of the robot dynamics should give greater versatility.

- **Exploiting task statics.** All the applications in this thesis are dynamic tasks, where precise motion is not required. Humans are generally skilled both at dynamic motions and precise manipulation tasks, such as threading a needle, sewing, writing, or peeling potatoes.

For these tasks the static behavior of the arm is more important than its dynamical behavior. The static properties of the arm such as its stiffness, whether it is braced to remove excess degrees of freedom, and whether the active degrees of freedom are aligned with the task can all be exploited to simplify the control required. The sensory modalities for these tasks are also different from the dynamic case, relying more on touch, force and vision sensing. These

different sensory modalities might have to be applied in different ways to modify the arm behavior.

The questions of how arm properties can be exploited for precise tasks, and how the sensory feedback should be used are challenging. Any progress towards performing these tasks with the same robustness and computational simplicity as the work in this thesis would be an exciting step forward. It would also greatly increase the versatility of robots.

- **Details.** Most of the work in this thesis ignores the details of arm control which would be needed for a flexible working system. A good example of this is starting and stopping the oscillators cleanly. This was not considered during the thesis but could be important for some tasks. A complete system would need details of this kind to be worked out more carefully.
- **Programming/learning.** The combination of arm and task dynamics is complicated, particularly if a wide range of tasks are considered. For the final system to be versatile and flexible it is inevitable that some learning will be required. This raises the questions of what should be learned and how should it be learned. While these specifics remained to be worked out, the type of tuning used in this thesis is a good model. Static parameters were tuned which determined a complete dynamic behavior, with the oscillator behavior reducing the complexity of the learning problem (by automatically being coordinated with the task). In addition the parameters themselves were easy to set, with good performance for a wide variety of parameter settings. While not all learning problems will be so straightforward these characteristics make the learning easy and so are worthwhile to bear in mind.

One of the major issues will be learning and scaling the complexity while retaining the useful properties of exploiting natural dynamics. This issue was touched on in chapter 5 where two different methods of connecting oscillators had rather different robustness properties.

Attacking the learning issue should shed light on how to increase the complexity of the robot behavior. Questions which are relevant here are the role of imitation in teaching the robot new tasks, and perhaps the role of memory and internal models in remembering and applying previously learnt behaviors to new situations. The transfer of skills to different parts of the workspace, to different arms and to different tasks is also important. Thinking of these questions in the context of exploiting natural dynamics may facilitate finding answers. For example, transfer of skills may only be useful if the underlying dynamics are similar.

8.3 Summary

This thesis has presented an approach to robot arm control which exploits natural dynamics. The approach uses a set of simple non-linear oscillators to drive the joints of a compliant arm. The oscillators use feedback to adjust their commands relative to the actual joint motion, and so tune into the arm's natural dynamics.

Chapter 3 introduced the oscillator and described its behavior when coupled to single degree of freedom systems. This behavior is generally coordinated with the motion of the system and robust to system and oscillator parameter changes. An analysis method using describing functions was described which is accurate and powerful; it enables detailed design of oscillator systems by predicting their behavior. It can also be used to explain the inherent robustness of the oscillators. The design tool is simple and easy to apply. It enables other researchers to understand, design and compare oscillator solutions to other approaches.

The oscillator robustness properties carry over to the multiple degree of freedom tasks which were considered in chapter 4. Using mechanical coupling between the oscillators as opposed to any explicit connections, the oscillators find and drive the system in its resonant mode, automatically adjusting themselves relative to the task. This is both remarkable and useful. It is remarkable that

independent control at the joints can give such complex coordinated behavior of the whole arm. It is also remarkable that this works when the arm is redundant, with more than one arm, and when the coupling is fairly weak, as in the Slinky toy example. This property is useful because a variety of tasks are defined as resonant modes of the arm system.

The robustness of this system is also compelling. Most of the parameters can be changed by factors of three without changing the overall performance. This makes tuning parameters trivial. The oscillator natural behavior removes the need to specify coordination between the joints, which makes tuning easier still, because there are less parameters to tune.

Not all tasks are conveniently specified by the resonant mode of the mechanical system and chapter 5 considered methods to extend the oscillator system to some of those cases. The main difficulty is retaining the robustness and self-organizing properties of the oscillators while influencing their final behavior. The chapter showed that coupling the oscillators through the natural dynamics was more robust and more powerful than connecting the oscillators into networks. This result is a compelling argument for exploiting natural dynamics.

Chapter 6 returned to the issue of analysis first raised in chapters 3 and 4. It presented a variety of theoretically exact results and tools for predicting the final behavior of oscillator driven systems. These methods together with the earlier results provide a powerful set of tools for understanding the behavior of oscillator driven systems.

Chapter 7 discussed various applications of the oscillator system. The chapter showed that the oscillators can produce both discrete and rhythmic motions. The oscillator system is generally easy to use. The parameters are easy to set and the feedback gives powerful and sensible behavior. It is indicative of the merits of the oscillators that all the tasks in the thesis were accomplished with the same oscillator architecture. The only changes were adjustments to the posture of the arm, the sizes of the joint commands, and the object attached to the robot's hand.

To conclude; the approach taken in this thesis of exploiting the natural dynamics of the arm using non-linear oscillators has been shown to work well in a number of respects. The approach can be applied over a wide range of discrete and rhythmic tasks, is robust and is easy to implement. These results suggests that the general approach of exploiting natural dynamics is a powerful method for obtaining coordinated dynamic behavior for robots.

Bibliography

- Adolfsson, J., Dankowicz, H., and Nordmark, A. (1998). 3-D stable gait in passive bipedal mechanisms. *Proceedings of 357 Euromech*.
- Alexander, R. M. (1990). Three uses for springs in legged locomotion. *International Journal of Robotics Research*, **9** (2), 53–61.
- An, C. H., Atkeson, C. G., and Hollerbach, J. M. (1988). *Model-based control of a robot manipulator*. Cambridge, MA, MIT Press.
- Arimoto, S. and Miyazaki, F. (1985). Stability and robustness of PID feedback control for robot manipulators of sensory capability. In *Proceedings of Third International Symposium of Robotics Research*, Gouvieux, France.
- Arshavsky, Y. I., Grillner, S., Orlovsky, G. N., and Panchin, Y. V. (1991). Central generators and the spatio-temporal pattern of movements. In J. Fagard and P. H. Wolff (Eds.), *The Development of Timing Control and Temporal Organization in Coordinated Action*, pp. 93–113. Elsevier Science.
- Asteroth, A., Fischer, M. S., Möller, K., and Schnepf, U. (1992). Tracking and grasping of moving objects - a behaviour-based approach. In *Proceedings of the 5th International Conference IEA/AIE-92. Industrial and Engineering Applications of AI and Expert systems*, Vol. 1, pp. 195–204.
- Atkeson, C. G. and Hollerbach, J. M. (1985). Kinematic features of unrestrained vertical arm movements. *Journal of Neuroscience*, **5**, 2318–2330.
- Bäck, T. and Schwefel, H.-P. (1993). An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation*, **1** (1), 1–23.
- Barbeau, H. and Rossignol, S. (1994). Enhancement of locomotor recovery following spinal cord injury. *Current Opinion in Neurology*, **7**, 517–524.
- Beccari, G. and Stramigioli, S. (1998). Impedance control as merging mechanism for a behaviour-based architecture. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, pp. 1429–1434, Leuven, Belgium.
- Bennett, D. J., Hollerbach, J. M., Xu, Y., and Hunter, I. W. (1992). Time-varying stiffness of human elbow joint during cyclic voluntary movement. *Experimental Brain Research*, **88**, 433–442.
- Bergen, A. R., Chua, L. O., Mees, A. I., and Szeto, E. W. (1982). Error bounds for general describing function problems. *IEEE Transactions on Circuits and Systems*, **CAS-29** (6), 345–354.
- Bernstein, N. S. (1967). *The coordination and regulation of movements*. Oxford, Pergamon.
- Bingham, G. P., Schmidt, R. C., and Rosenblum, L. D. (1989). Hefting for a maximum distance throw: A smart perceptual mechanism. *Journal of Experimental Psychology: Human Perception and Performance*, **15** (3), 507–528.

- Bingham, G. P., Schmidt, R. C., Turvey, M. T., and Rosenblum, L. D. (1991). Task dynamics and resource dynamics in the assembly of a coordinated rhythmic activity. *Journal of Experimental Psychology: Human Perception and Performance*, **17** (2), 359–381.
- Bird, J. O. (1997). *Electrical Circuit Theory and Technology*. Butterworth-Heinemann.
- Branicky, M. S. (1994). Stability of switched and hybrid systems. In *Proceedings of the 33rd Conference on Decision and Control*, pp. 3498–3503, Lake Buena Vista FL.
- Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, **RA-2**, 14–23.
- Brooks, R. A. (1995). *The L Manual*. IS Robotics, Cambridge MA.
- Brooks, R. A., Breazeal, C., Irie, R., Kemp, C. C., Marjanović, M., Scassellati, B., and Williamson, M. (1998). Alternate essences of intelligence. In *Proceedings of American Association of Artificial Intelligence (AAAI-98)*, Madison, WI.
- Brooks, R. A., Breazeal, C., Kemp, C. C., Marjanović, M., Scassellati, B., and Williamson, M. (1999). Imitation through social interaction. *Applied Artificial Intelligence*. In preparation.
- Brooks, R. A. and Stein, L. A. (1994). Building brains for bodies. *Autonomous Robots*, **1** (1), 7–25.
- Brown, K. L. (1994). *Design and Analysis of Robots that Perform Dynamic Tasks Using Internal Body Motion*. PhD thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA.
- Canavier, C. C., Butera, R. J., Dror, R. O., Baxter, D. A., j. W. Clark, and Byrne, J. H. (1997). Phase response characteristics of model neurons determine which patterns are expressed in a ring circuit model of gait generation. *Biological Cybernetics*, **77** (6), 367–380.
- Cohen, A. H., Holmes, P. J., and Rand, R. H. (1982). The nature of coupling between segmental oscillators of the lamprey spinal generator for locomotion: A mathematical model. *Journal of Mathematical Biology*, **13**, 345–369.
- Colgate, J. E. and Hogan, N. (1988). Robust control of dynamically interacting systems. *International Journal of Control*, **48** (1), 65–88.
- Colgate, J. E. and Hogan, N. (1989). An analysis of contact instability in terms of passive equivalents. In *Proceedings of the 1989 IEEE International Conference on Robotics and Automation*, pp. 404–409.
- Collins, J. J. and Richmond, S. A. (1994). Hard-wired central pattern generators for quadrupedal locomotion. *Biological Cybernetics*, **71**, 375–385.
- Connell, J. H. (1988). A behavior-based arm controller. Technical Report 1025, Massachusetts Institute of Technology Artificial Intelligence Lab, Cambridge, Massachusetts.
- Connell, J. H. (1994). Get me that screwdriver! developing an sensory-action vocabulary for fetch-and-carry tasks. Technical Report RC 19473, IBM T. J. Watson.
- der Linde, R. Q. V. (1998). Active leg compliance for passive walking. *IEEE Conference on Robotics and Automation*, pp. 2339–2345.
- Doya, K. and Yoshizawa, S. (1989). Adaptive neural oscillator using continuous-time back-propagation learning. *Neural Networks*, **2**, 375–385.

- Doya, K. and Yoshizawa, S. (1992). Adaptive synchronization of neural and physical oscillators. In J. E. Moody, S. J. Hanson, and R. P. Lippmann (Eds.), *Advances in Neural Information Processing Systems*, Vol. 4, pp. 109–116, San Mateo, CA. Morgan Kaufmann.
- Drake, S. H. (1977). *Using Compliance in Lieu of Sensory Feedback for Automatic Assembly*. PhD thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology.
- Dror, R. O., Canavier, C. C., Butera, R. J., Clark, J. W., and Byrne, J. H. (1999). A mathematical criterion based on phase response curves for stability in a ring of coupled oscillators. *Biological Cybernetics*, **80** (1), 11–23.
- Ekeberg, Ö. (1993). A combined neuronal and mechanical model of fish swimming. *Biological Cybernetics*, **69**, 363–374.
- Epstein, S. and Kopell, N. (1999). Resonance tuning revisited. In preparation.
- Ermentrout, B. and Kopell, N. (1994). Learning of phase lags in coupled oscillators. *Neural Computation*, **6** (2), 225–241.
- Fowble, J. and Kuo, A. (1996). Stability and control of passive locomotion in 3D. *Proceedings of the Conference on Biomechanics and Neural Control of Movement*, pp. 28–29.
- Franklin, G. F., Powell, J. D., and Workman, M. L. (1992). *Digital Control of Dynamic Systems*. Addison-Wesley, 2nd edition.
- Gandolfo, F., Mussa-Ivaldi, F. A., and Bizzi, E. (1996). Motor learning by field approximation. *Proceedings of the National Academy of Sciences*, **93** (9), 3842–3846.
- Garcia, M., Chatterjee, A., and Ruina, A. (1998). Speed, efficiency, and stability of small-slope 2D passive dynamic bipedal walking. *IEEE International Conference on Robotics and Automation*, pp. 2351–2356.
- Gelb, A. and Vander Velde, W. E. (1968). *Multiple-input Describing Function and Nonlinear System Design*. McGraw-Hill.
- Gershon, D. (1990). Parallel process decomposition of a dynamic manipulation task: Robotic sewing. *IEEE Transactions of Robotics and Automation*, **6** (3), 357–367.
- Getting, P. A. (1988). Comparative analysis of invertebrate central pattern generators. In A. H. Cohen, S. Rossignol, and S. Grillner (Eds.), *Neural Control of Rhythmic Movements in Vertebrates*, chapter 4, pp. 101–127. New York, Wiley.
- Goldfield, E. C., Kay, B. A., and Warren, W. H. (1993). Infant bouncing—the assembly and tuning of action systems. *Child Development*, **64** (4), 1128–1142.
- Gonçalves, J. M. (1999). Analysis of switching systems. PhD thesis proposal, MIT Dept of Electrical Engineering and Computer Science.
- Gonçalves, J. M., Megretski, A., and Dahleh, M. (1998). Semi-global analysis of relay feedback systems. In *Proceedings IEEE Conference on Decision and Control*.
- Gribble, P. L., Ostry, D. J., Sanguineti, V., and Laboissiere, R. (1998). Are complex control signals required for human arm movement? *Journal of Neuroscience*, **79** (3), 1409–1424.
- Hajian, A. Z., Sanchez, D. S., and Howe, R. D. (1997). Drum roll: Increasing bandwidth through passive impedance modulation. In *Proceedings IEEE International Conference on Robotics and Automation*, pp. 2294–2299, Albuquerque, NM.

- Haken, H., Kelso, J. A. S., and Bunz, H. (1985). A theoretical model of phase transitions in human hand movements. *Biological Cybernetics*, **51**, 347–356.
- Hatsopoulos, N. G. (1996). Coupling the neural and physical dynamics in rhythmic movements. *Neural Computation*, **8**, 567–581.
- Hatsopoulos, N. G. and Warren, W. H. (1996). Resonance tuning in rhythmic arm movements. *Journal of Motor Behavior*, **28** (1), 3–14.
- Hatsopoulos, N. G., Warren, W. H., and Sanes, J. N. (1992). A neural pattern generator that tunes into the physical dynamics of the limb system. In *Proceedings of the International Joint Conference on Neural Networks*, Vol. 1, pp. 104–109, Baltimore, MD.
- Herr, H. (1993). *Human Powered Elastic Mechanisms*. Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Hogan, N. (1985a). Impedance control: An approach to manipulation. *Journal of Dynamic Systems, Measurement, and Control*, **107**, 1–24.
- Hogan, N. (1985b). The mechanics of multi-joint posture and movement control. *Biological Cybernetics*, **52**, 315–331.
- Hollerbach, J. M. (1981). An oscillator theory of handwriting. *Biological Cybernetics*, **39** (2), 139–156.
- Hunter, I. W., Hollerbach, J. M., and Ballantyne, J. (1991). A comparative analysis of actuator technologies for robotics. *Robotics Review*, **2**.
- Ijspeert, A. J., Hallam, J., and Willshaw, D. (1998). From lampreys to salamanders: Evolving neural controllers for swimming and walking. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S. W. Wilson (Eds.), *Fifth International Conference on Simulation of Adaptive Behavior*, pp. 390–399, Cambridge, MA. The MIT Press.
- Kay, B. A., Kelso, J. A. S., Saltzman, E. L., and Schöner, G. S. (1987). The space-time behavior of single and bimanual movements: Data and model. *Journal of Experimental Psychology: Human Perception and Performance*, **13**, 178–192.
- Khalil, H. K. (1996). *Nonlinear Systems*. New Jersey, Prentice-Hall, 2nd edition.
- Kimura, H., Sakurama, K., and Akiyama, S. (1998). Dynamic walking and running of the quadruped using neural oscillators. In *Proceedings of the IEEE/RSJ International Workshop on Intelligent Robots and Systems (IROS-98)*, Vol. 1, pp. 50–57, Victoria, B.C.
- Kimura, S., Yano, M., and Shimizu, H. (1993). A self-organizing model of walking patterns of insects. *Biological Cybernetics*, **69**, 183–193.
- Kuniyoshi, Y. and Nagakubo, A. (1997). Humanoid as a research vehicle into flexible complex interaction. In *Proceedings of 1997 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '97)*, Vol. 2, pp. 811–820.
- Latombe, J.-C. (1991). *Robot Motion Planning*. Norwell, Massachusetts, Kluwer Academic Press.
- Lohmiller, W. and Slotine, J.-J. (1998). On contraction analysis for non-linear systems. *Automatica*, **34** (6), 683–696.
- Lynch, K. M., Shiroma, N., Arai, H., and Tanie, K. (1998). The roles of shape and motion in dynamic manipulation: The butterfly example. In *Proceedings of the 1998 IEEE International Conference on Robotics and Automation*, pp. 1958–1963, Leuven, Belgium.

- Marjanović, M. J., Scassellati, B., and Williamson, M. M. (1996). Self-taught visually-guided pointing for a humanoid robot. In P. Maes, M. Matarić, J.-A. Meyer, J. Pollock, and S. Wilson (Eds.), *From Animals to Animats: Proceedings 1996 Society of Adaptive Behaviour*, pp. 35–44, Cambridge, MA. The MIT Press.
- Mason, M. T. (1982). Compliant motion. In M. Brady, J. M. Hollerbach, T. L. Johnson, T. Lozano-Pérez, and M. T. Mason (Eds.), *Robot Motion: planning and control*, chapter 5, pp. 305–322. MIT Press.
- Mason, M. T. and Lynch, K. M. (1993). Dynamic manipulation. In *Proceedings of the IEEE/RSJ International Workshop on Intelligent Robots and Systems (IROS-93)*, pp. 152–159, Yokohama, Japan.
- Massone, L. L. E. and Myres, J. D. (1996). The role of plant properties in arm trajectory formation: A neural network study. *IEEE Transactions on Systems, Man and Cybernetics*, **26**, 719–732.
- Matsuoka, K. (1985). Sustained oscillations generated by mutually inhibiting neurons with adaptation. *Biological Cybernetics*, **52**, 367–376.
- Matsuoka, K. (1987). Mechanisms of frequency and pattern control in neural rhythm generators. *Biological Cybernetics*, **56**, 345–353.
- McGeer, T. (1990). Passive dynamic walking. *International Journal of Robotics Research*, **9** (2), 62–82.
- Megretski, A. (1997). Global stability of oscillations induced by a relay feedback. In *Proceedings of 13th World Congress, International Federation of Automatic Control*, Vol. E, pp. 49–54, Oxford, UK. Pergamon.
- Miyakoshi, S., Taga, G., Kuniyoshi, Y., and Nagakubo, A. (1998). Three dimensional bipedal stepping motion using neural oscillators - towards humanoid motion in the real world. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-98)*, Vol. 1, pp. 84–89, Victoria, B.C.
- Miyakoshi, S., Yamakita, M., and Furata, K. (1994). Juggling control using neural oscillators. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-94)*, Vol. 2, pp. 1186–1193, Munich, Germany.
- Miyamoto, H. and Kawato, M. (1998). A tennis serve and upswing learning robot based on bi-directional theory. *Neural Networks*, **11** (7-8), 1331–1344.
- Miyamoto, H., Schaal, S., Gandolfo, F., Gomi, H., Koike, Y., Osu, R., Nakano, E., Wada, Y., and Kawato, M. (1996). A kendama learning robot based on bi-directional theory. *Neural Networks*, **9** (8), 1281–1302.
- Morasso, P. (1981). Spatial control of arm movements. *Experimental Brain Research*, **42**, 223–227.
- Mussa-Ivaldi, F. A., Hogan, N., and Bizzi, E. (1985). Neural, mechanical, and geometric factors subserving arm posture in humans. *Journal of Neuroscience*, **5** (10), 2732–2743.
- Niemeyer, G. and Slotine, J.-J. E. (1997). A simple strategy for opening an unknown door. In *Proceedings of 1997 IEEE International Conference on Robotics and Automation*, pp. 1448–1453.
- Nishii, J. (1995). An adaptive control model of a locomotion by the central pattern generator. In *Proceeding of IWANN'95. International Workshop on Artificial Neural Networks*, pp. 151–157, Malaga-Torremolinos, Spain.

- Nishii, J. (1998). A learning model for oscillatory networks. *Neural Networks*, **11**, 249–257.
- Ogata, K. (1970). *Modern Control Engineering*. Prentice-Hall.
- Pagano, C. and Turvey, M. T. (1995). The inertia tensor as a basis for the perception of limb orientation. *Journal of Experimental Psychology: Human Perception and Performance*, **21** (5), 1070–1087.
- Pearson, K. G. (1993). Common principles of motor control in vertebrates and invertebrates. *Annual Review of Neuroscience*, **16**, 265–297.
- Pettit, N. B. O. L. (1995). *Analysis of Piecewise Linear Dynamical Systems*. New York, NY, Wiley.
- Playter, R. R. (1994). *Passive Dynamics in the Control of Gymnastic Maneuvers*. PhD thesis, Department of Aeronautical and Astronautical Engineering, Massachusetts Institute of Technology, Cambridge, MA.
- Playter, R. R. and Raibert, M. H. (1992). Control of a biped somersault in 3D. In *Proceedings IROS '92*, Raleigh, NC.
- Playter, R. R. and Raibert, M. H. (1994). Passively stable layout somersaults. In *Proceedings Eighth Yale Workshop on Adaptive and Learning Systems*, pp. 66–71, New Haven, CT.
- Pratt, G. A. and Williamson, M. M. (1995). Series elastic actuators. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-95)*, Vol. 1, pp. 399–406, Pittsburg, PA.
- Pratt, J., Dilworth, P., and Pratt, G. (1997). Virtual model control of a bipedal walking robot. In *Proceedings of the IEEE International Conference on Robotics and Automations (ICRA '97)*, Albuquerque, NM.
- Pratt, J., Torres, A., Dilworth, P., and Pratt, G. (1996). Virtual actuator control. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, Osaka, Japan.
- Pratt, J. E. and Pratt, G. A. (1998). Exploiting natural dynamics in the control of a planar bipedal walking robot. In *Proceedings of the Thirty-Sixth Annual Allerton Conference on Communication, Control and Computing*, Monticello, IL.
- Pribe, C., Grossberg, S., and Cohen, M. A. (1997). Neural control of interlimb oscillations. II. Biped and quadruped gaits and bifurcations. *Biological Cybernetics*, **77**, 141–152.
- Raibert, M. H. (1986). *Legged Robots That Balance*. Cambridge, Massachusetts, MIT Press.
- Raibert, M. H. and Craig, J. J. (1981). Hybrid position/force control of manipulators. *Journal of Dynamic Systems, Measurement, and Control*, **103**, 126–133.
- Rich, E. and Knight, K. (1991). *Artificial Intelligence*. McGraw-Hill, second edition.
- Rizzi, A. A. and Koditschek, D. E. (1994). Further progress in robot juggling: Solvable mirror laws. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Vol. 4, pp. 2935–2940, San Diego, CA.
- Rossignol, S. (1996). Neural control of stereotypic limb movements. In L. B. Rowell and J. T. Sheperd (Eds.), *Handbook of Physiology, Section 12. Exercise: Regulation and Integration of Multiple Systems*, chapter 5, pp. 173–216. American Physiological Society.
- Rossignol, S. and Dubuc, E. (1994). Spinal pattern generation. *Current Opinion in Neurobiology*, **4**, 894–902.

- Sabes, P. N. and Jordan, M. I. (1997). Obstacle avoidance and a perturbation sensitivity model for motor planning. *Journal of Neuroscience*, **17** (18), 7119–7128.
- Salisbury, J., Townsend, W. T., Eberman, B. S., and DiPietro, D. M. (1988). Preliminary design of a whole arm manipulation system (WAMS). In *Proceedings 1988 IEEE International Conference on Robotics and Automation*.
- Saltzman, E. L. and Kelso, J. A. S. (1987). Skilled actions: A task dynamic approach. *Psychological Review*, **94** (1), 84–106.
- Scassellati, B. (1998). A binocular, foveated active vision system. Technical Report 1628, MIT Artificial Intelligence Lab Memo.
- Schaal, S. (1999). Is imitation learning the way to humanoid robots? *Trends in Cognitive Sciences*. In submission.
- Schaal, S. and Atkeson, C. G. (1993). Open loop stable control strategies for robot juggling. In *Proceedings 1993 IEEE International Conference on Robotics and Automation*, Vol. 3, pp. 913–918.
- Schaal, S. and Sternad, D. (1998). Programmable pattern generators. In *Proceedings of International Conference on Computational Intelligence in Neuroscience (ICCIN '98)*, Triangle Park, NC.
- Scheirer, E. D. (1998). Tempo and beat analysis of acoustic musical signals. *J. Acoust. Soc. Am.*, **103** (1), 588–601.
- Schneider, K., Zernicke, R. F., Schmidt, R. A., and Hart, T. J. (1989). Changes in limb dynamics during the practice of rapid arm movements. *Journal of Biomechanics*, **22** (8–9), 805–817.
- Schöner, G. (1990). A dynamic theory of coordination of discrete movement. *Biological Cybernetics*, **63**, 257–270.
- Slotine, J.-J. E. (1988). Putting physics in control—the example of robotics. *IEEE Control Systems Magazine*, **8** (6), 12–18.
- Slotine, J.-J. E. and Li, W. (1991). *Applied nonlinear control*. Englewood Cliffs, N.J., Prentice Hall.
- Smithers, T. and Malcolm, C. (1988). Programming robotic assembly in terms of task achieving behavioural modules. DAI Research Paper 417, University of Edinburgh, Edinburgh, Scotland.
- Strang, G. (1993). *Introduction to Linear Algebra*. Wellesley, MA, Wellesley-Cambridge Press.
- Strogatz, S. H. (1994). *Nonlinear dynamics and chaos: with application to physics, biology, chemistry and engineering*. Reading, MA, Addison-Wesley.
- Taga, G. (1995a). A model of the neuro-musculo-skeletal system for human locomotion. I. Emergence of basic gait. *Biological Cybernetics*, **73** (2), 97–111.
- Taga, G. (1995b). A model of the neuro-musculo-skeletal system for human locomotion. II. Real-time adaptability under various constraints. *Biological Cybernetics*, **73** (2), 113–121.
- Taga, G. (1998). A model of the neuro-musculo-skeletal system for anticipatory adjustment of human locomotion during obstacle avoidance. *Biological Cybernetics*, **78** (1), 9–17.
- Taga, G., Yamaguchi, Y., and Shimizu, H. (1991). Self-organized control of bipedal locomotion by neural oscillators in unpredictable environment. *Biological Cybernetics*, **65** (3), 147–159.

- Takanobu, H., Kuchiki, N., and Takanishi, A. (1995). Control of rapid closing motion of a robot jaw using nonlinear spring mechanism. In *Proceedings 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems.*, Vol. 1, pp. 372–377, Los Alamitos, CA.
- Takegaki, M. and Arimoto, S. (1981). A new feedback method for dynamic control of manipulators. *Journal of Dynamic Systems, Measurement, and Control*, **102**.
- Terman, D., Kopell, N., and Bose, A. (1998). Dynamics of two mutually inhibitory neurons. *Physica D*, **117**, 241–275.
- Thelen, E., Zernicke, R., Schneider, K., Jensen, J., Kamm, K., and Corbetta, D. (1992). The role of intersegmental dynamics in infant neuromotor development. In G. E. Stelmach and J. Requin (Eds.), *Tutorials in Motor Behavior II*, pp. 533–548. Elsevier Science Publishers.
- Trevelyan, J. (1992). Force controlled surface following without stability problems. In *TENCON '92. 'Technology Enabling Tomorrow' 1992 IEEE Region 10 International Conference. Computers, Communications and Automation towards the 21st Century*, Vol. 1, pp. 247–251, New York, NY.
- Turvey, M. T. and Carello, C. (1995). Dynamic touch. In W. Epstein and S. Rogers (Eds.), *Perception of Space and Motion*, pp. 401–490. Academic Press.
- Ulrich, N. T. (1990). Mechanical design optimization of robot manipulator performance. Technical Report MS-CIS-90-72, GRASP LAB 238, Department of Computer and Information Science, School of Engineering and Applied Science, University of Pennsylvania, Philadelphia, PA.
- Wadden, T. and Ekeberg, Ö. (1998). A neuro-mechanical model of legged locomotion: single leg control. *Biological Cybernetics*, **79**, 161–173.
- Whitney, D. E. (1977). Force feedback control of manipulator fine motions. *Transactions of the ASME Journal of Dynamical Systems, Measurement and Control*, **99**, 91–97.
- Whitney, D. E. (1982). Quasi-static assembly of compliantly supported rigid parts. *Journal of Dynamic Systems, Measurement and Control*, **104**, 65–77.
- Williamson, M. M. (1995). Series elastic actuators. Technical Report 1524, Massachusetts Institute of Technology Artificial Intelligence Lab, Cambridge, Massachusetts.
- Williamson, M. M. (1996). Postural primitives: Interactive behavior for a humanoid robot arm. In *Fourth International Conference on Simulation of Adaptive Behavior*, pp. 124–131, Cape Cod, Massachusetts.
- Williamson, M. M. (1998a). Neural control of rhythmic arm movements. *Neural Networks*, **11**, 1379–1394.
- Williamson, M. M. (1998b). Rhythmic robot control using oscillators. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '98)*.
- Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). Are arm trajectories planned in kinematic or dynamic coordinates—An adaptation study. *Experimental Brain Research*, **103** (3), 460–470.
- Zajac, F. E. (1989). Muscle and tendon: Properties, models, scaling, and application to biomechanics and motor control. *CRC Critical Reviews of Biomedical Engineering*, **17** (4), 359–411.
- Zielińska, T. (1996). Coupled oscillators utilised as gait rhythm generators of a two-legged walking machine. *Biological Cybernetics*, **74**, 263–273.

Appendix A

Experimental Apparatus

This chapter described the design of the arms used throughout this thesis. It details the design of the actuators, the mechanical design of the various joints, the sensors used at each joint, the wiring of the robot, the control architecture and control algorithms.

A.1 Overall arm design

The two arms used both have six degrees of freedom, and are mounted on the humanoid robot Cog (Brooks and Stein, 1994, Brooks et al., 1998). The arms are mirror images of one another and are otherwise identical. The two arms are illustrated in figure A-1. The kinematics of the arms are designed to be similar to a human arm, and are shown in figure A-2. There are two joints each at the shoulder, elbow and wrist, although the axis of the first elbow joint (elbow-a in figure A-2) is coincident with the axes of the shoulder joints. The arm is approximately the same length as a human arm.




Figure A-1: Picture of the two robot arms. Each arm has 6 degrees of freedom, arranged in a similar manner to a human arm. The arms are approximately the same length as a human arm




Figure A-2: Arm kinematics. The thick lines indicate the structure of the arm, and the thin lines the axes of joints which do not line up with the arm structure. There are two joints for each of the shoulder, elbow and wrist, making six degrees of freedom in total. The axes of the shoulder-a, shoulder-b, and elbow-a are coincident at the shoulder, and the axes at the wrist also intersect. The elbow-b joint is offset to allow the arm to curl up so that the hand can touch the shoulder.

Each joint of the arm is actuated by a series elastic actuator (Williamson, 1995, Pratt and Williamson, 1995), which consists of an ordinary electric motor and gearbox with a spring in series with the motor output. These actuators give good quality force control, shock tolerance, and stable behavior of the whole arm. They thus allow the whole arm to be robust and compliant.

The following sections describe in more detail the actuator and arm design, and the design of the controllers.

A.2 Actuators

The motivation for the use of series elastic actuators comes from the consideration that electric motors, which are a commonly available actuator for robotics, are most efficient when spinning at high speed with low torque, while most robotic applications require motors which operate at low speed and high torque. The planetary gearheads which are commonly used as a compact method for converting the speeds and torque unfortunately suffer from a number of disadvantages. The gear teeth introduce friction and backlash into the motor drive, and the teeth are weak and often break under shock loads (Hunter et al., 1991). Alternatives such as stictionless cable drives are bulky (as used in the WAM robot (Salisbury et al., 1988)), as are direct drive motors (An et al., 1988). The characteristics of gearboxes make them most suitable for position controlled tasks, rather than for force control. However, for the kind of robust interaction desired for the arms used in this thesis, force control at the joints is necessary.

Another consideration comes from robot force control. Early researchers found that controlling stiff robot arms to contact hard surfaces was difficult, but that by covering the robot is some

compliant material (foam rubber), the task was made somewhat easier (Whitney, 1977).

These two considerations come together in the series elastic actuator, shown schematically in figure A-3. The actuator consists of an ordinary electric motor and gearbox with a spring connected between the motor output and the actuator output. The spring is thus in series with the load. The actuator controls the force output by controlling the deflection of the spring. The spring thus turns the force control problem into a position control problem, which is better suited to the motor type. It can also be thought of as moving the compliance from the end of the arm into each joint.




Figure A-3: Schematic of actuator design. The series elastic actuator consists of an ordinary electrical motor and gearbox, with a spring connected between the motor output and the actuator output. By controlling the deflection of the spring, the output force of the actuator can be controlled. The spring naturally filters out the friction and backlash introduced by the gearbox, contributing to the quality of the force control. It also absorbs shock loads, protecting the gearbox teeth from damage.

Adding the spring has some advantages. The spring acts as a low pass filter, naturally filtering out the friction and backlash introduced by the gearbox. It also absorbs any shock loads, storing the energy of impact in the spring so protecting the gearbox teeth from damage. Finally the spring makes it easy for the overall actuator to be passive, which guarantees its stability when contacting any passive environment (Colgate and Hogan, 1988). Intuitively this is because at low frequencies the control law will work well so the actuator will be stable, while at high frequencies above the bandwidth of the controller, the actuator behavior will be dominated by the physical spring. The spring is a passive element which ensures that the actuator will be passive at all frequencies.

The main disadvantage of adding the spring is that the bandwidth of the actuator is limited for some motions. For an electric motor, the force is proportional to the current through the motor windings which can be changed quickly, without the motor actually moving. For the series elastic actuator, one end of the spring has to be moved in order to generate the force. The current through the motor has to move the mass of the motor to generate this deflection, which results in a loss in bandwidth. The other disadvantage is the extra design complexity introduced by incorporating the spring. Even a compact spring involves extra coupling components and bearings, and thus complicates the mechanical design.

Series elastic actuators are increasing in popularity, having also been used in the Leg Lab at MIT e.g. Pratt et al. (1996, 1997).

A.3 Spring design

The major difficulty in implementing a series elastic actuator is the design of the spring, particularly if a torsional spring is required. This is because the spring carries the whole load of the actuator but must also be compliant enough to get the benefits of the series elasticity. This combination of properties makes the spring design difficult. See Williamson (1995) for exhaustive details.

The design chosen consists of a steel bar with a cross shaped cross-section, as shown in figure A-4. This section can be designed to be strong yet not particularly stiff because the thin sections of the cross shape allow large deflections without yielding. The deflection of the spring is measured using strain gauges mounted on the flats of the spring, as shown in figure A-5.




Figure A-4: Spring design. The figure shows a torsional spring with a cross-shaped cross-section. This shape gives a spring which is strong while still being flexible.

The model of the spring developed in (Williamson, 1995) compared the cross shaped spring to an equivalent flat plate spring as shown in figure A-6. The parameters which determine the spring behavior are the length l , web thickness t and the combined breadth and height of the cross shape $b + h$. The behavior of a flat plate under torsion is governed by the following equations

$$\theta_{max} = \frac{l\tau_{max}}{tG} \quad (\text{A.1})$$

$$k_{spring} = \frac{(b+h)t^3G}{3l} \quad (\text{A.2})$$

$$T_{max} = k_{spring}\theta_{max} = \frac{(b+h)t^2\tau_{max}}{3} \quad (\text{A.3})$$

where θ_{max} is the maximum deflection of the spring, τ_{max} is the yield shear stress (half the yield stress




Figure A-5: The deflection of the spring is measured using strain gauges mounted on the flats of the spring, as shown here. Four gauges are used in a Wheatstone Bridge configuration to measure the strain in the spring, and so the deflection. The deflection is proportional to the torque transmitted through the spring.




Figure A-6: Figure comparing the cross-shaped spring with a flat plate in torsion. Since the cross-shaped spring is composed of flat plates in torsion, its behavior is related to the behavior of a flat plate. The equivalent flat plate spring has the same length (l), web thickness (t) with a width the same as the sum of the width and height of the cross spring ($b + h$).

σ_{yield}), G the shear modulus, k_{spring} the spring stiffness and T_{max} the maximum load sustainable by the spring.

The cross shaped spring behavior is similar to that of a flat plate, with some linear factor differences as shown in figure A-7 and figure A-8. These plots show the comparison between the actual stiffness and maximum load for various springs and the stiffness and loads calculated for an equivalent flat plate spring. The relationship in both cases is linear, with the cross spring having approximately the same stiffness as the flat plate, but being 1.6 times as strong. The cross section is also more compact, which is important for design purposes.

The linear relationships in figures A-7 and A-8 allow the design of new springs, since new values of t , l and $b + h$ can be chosen using a design specification and the following equations:

$$\theta_{max} = 1.65 \frac{l \tau_{max}}{tG} \quad (A.4)$$

$$k_{spring} = 0.98 \frac{(b+h)t^3 G}{3l} \quad (A.5)$$

$$T_{max} = 1.63 \frac{(b+h)t^2 \tau_{max}}{3} \quad (A.6)$$




Figure A-7: REAL Comparison of cross-spring stiffness and flat plate stiffness. The figure shows the measured stiffness of a number of springs plotted against the stiffness of an equivalent flat plate of similar dimensions and material. The points lie on a straight line, with gradient 0.98. This suggests that the stiffness of a cross-shaped spring is almost the same as for a flat plate of the same dimensions.

A.4 Actuator control law

The control law for the actuator is shown in figure A-9. The input to the control loop is the desired deflection of the spring, which is proportional to the desired torque. A conventional PID controller is used to control the spring deflection, with the integral term being rolled off at low frequencies to preserve passivity (see Williamson (1995) for more details). A feedforward term of the desired




Figure A-8: **REAL** Comparison of maximum torque before yielding for a cross-shaped spring and the calculated value for an equivalent flat plate. The points lie on a straight line, indicating that the flat plate is a good model of the cross-shaped spring's behavior. The gradient of the line is 1.62, showing the higher load-bearing capacity of the cross-shaped cross-section.

deflection/force is used to increase the speed of the system. The output of the control law is the desired force or current in the motor. A simple motor model is used to convert this to the voltage to be applied to the motor windings. The motor model compensates for the back e.m.f. generated by the motor motion, using the encoder sensor to determine the motor velocity. The physical implementation of this control law is described later in this Appendix.




Figure A-9: Figure showing the control law for the actuator. The desired force of spring deflection is the input, and a PID controller is used to control the actual deflection to track the desired deflection. The integral term is rolled off at low frequencies to preserve passivity (see Williamson (1995) for more details). A feedforward term is used to speed up the system. The output from the PID is the desired motor current, which is transformed using a motor model to a voltage which is applied to the actuator. The motor model removes the effect of back e.m.f. caused by the motor motion.

The series elastic actuators provide force control, giving a force output at the joint which is close to the desired force. To control the position of the joint, a low gain proportional-derivative (PD) controller is used, whose output is the desired torque or force at the joints, as shown in figure A-10. The stiffness and damping at the joints can be varied by changing stiffness K or damping B , and the arm moved around by changing the setpoint θ_v of the controller. The PD control makes the arm behave as if its links are connected by springs and dampers, but because the forces in the joints are

accurately controlled by the series elastic actuators, the overall motion is smooth and compliant. The force control bandwidth of the actuators is fairly low, which forces the stiffness of the arm to be low. This gives robustness when interacting with objects.




Figure A-10: The joint level control consists of a position control loop, with stiffness K , damping B , and setpoint θ_v organized around an inner torque control loop provided by the series elastic actuators. The inner loop controls the force output of the joint u to accurately track the desired force u_d . Organizing the control in this fashion results in smooth, compliant motion of the arm, where the overall stiffness and damping can be altered by changing the values of K and B .

A.5 Arm design

A.5.1 General

As described previously, each arm has six degrees of freedom, as shown in figure A-2. As far as possible the arm was designed to be lightweight distributing the weight as far up the arm as possible to minimize loads at the joints due to the arm itself. It was also designed to be modular, each arm being easily split into three parts, with a minimum of mechanical or electrical disassembly.

A.5.2 Joint configurations

A major part of the design concerned mounting and positioning the springs in the series elastic actuators. The spring, as described in section A.3, is approximately as long as a motor, and generally has an outside diameter of 1". For the series elastic actuators to work well, the spring needs to be between any transmission and the actuator output, which implies that the spring has to lie on the axis of each joint. To give added flexibility in the positioning of motors, cable drives were used as an extra reduction after the motor gearbox. This allowed the choice of smaller gearboxes for the motors. The cable reduction stage adds complexity, requiring extra bearings, pulleys, cables and cable tensioners, but mechanically the reduction is preferable to using a larger gearbox. This is because the cables introduce no backlash, have no stiction (static friction), and are stronger and more resilient to shock loads than gearbox teeth. Unfortunately they are bulky, which is why a combination of gearboxes and cables were used.

The arm joints were designed to have one of two similar configurations, the first of which is shown in figure A-11. The motor (gearbox - 1, motor - 2, encoder - 3) is placed parallel to the spring (8), which lies inside a tube (7) along the axis of the joint. The motor is connected to a motor pulley (11), which is connected through a cable drive to the large pulley (9). The pulley floats on the tube supported by a bearing, and the tube itself is supported on two bearings (6) by supports (5), which also clamp the motor. One end of the spring is rigidly connected to the pulley through a connecting piece (10), and the other end is rigidly connected to the tube (12). If the right hand end of the tube is grounded, and the motor turned on, the whole motor-support-cover piece will walk round

the outside of the tube. The output of the actuator is the whole arm structure, with the spring placed between the output and the tube which is grounded. The elbow-a and wrist-a joints use this configuration.

Reversing inputs and outputs gives the configuration in figure A-12. This has the motor and supports grounded, with the actuator output either being the end of the tube (7), or from a clamp attached to the tube (27). The shoulder-a, shoulder-b, elbow-b and wrist-b joints all use this second configuration.

These configurations are compact, since the length of the spring is folded beside the length of the motor, strong because the load is transferred from the tube to the arm structure using two bearings which are widely separated, and easy to align because all the bearings are on the same shaft. Because the tube is hollow, it is also possible to run electrical cables through the joint to minimize bending and twisting.

A schematic of the bearing design used for all the joints is shown in figure A-13. The tube (7) is supported on two supports (5), through angular contact bearings (15). These bearings can withstand radial loads, but only axial loads in one direction. They are normally used in pairs which “push” against one another. The further apart the bearings are, the stiffer the connection between the supports and tube, and the stronger the arm. The bearings are preloaded using a shoulder on the tube (7a), shoulders on the supports (5a), a spacer (13), and a tightening ring (14). Tightening up the ring (normally with screws into the tube itself) tightens up the bearing surfaces against one another. The pulley (9) is attached to the spring and rotates relative to the tube when the spring deflects. It is mounted on a single radial bearing (16).

A schematic of the cable mechanism used for all the joints is shown in figure A-14. As mentioned above, the cables are used to give flexibility in the positioning of the motor, as well as provide extra reduction. Each joint has a cable, with an associated cable mechanism. One end of the cable (21) is terminated (22) in a slot in the large pulley (9a), from which it wraps around the large pulley (9), around the motor pulley (11), then terminates in the termination block (23). The tension in the cable is created by using the tensioner screw (24) to pull the block (23) backwards and forwards. The tensioning mechanism is part of a metal component (25), which also doubles as a spring holder, rigidly attaching the spring (8) to the large pulley (9). Four of the joints in the arm have the motor pulley close to the large pulley, however for the elbow-b and wrist-b joints extra pulleys were used to position the motor away from the joints. The cable routes in those cases are illustrated in figure A-15.

The shoulder and upper arm of the robot is shown in figure A-16, which shows clearly the shoulder-a and shoulder-b joints. Each of these have the motor grounded, with the tube providing the actuator output. The upper arm is clamped to the tube of the shoulder-b joint. The clamp is used to allow simple assembly and disassembly of the arm. The pulley for the shoulder-b joint is not completely circular to save space.

Both elbow joints are shown in figure A-17. The tube and pulley for the elbow-a joint are visible, connecting to the shoulder-b joint. The elbow-a motor is connected to a support and the two plates which form the structure of the upper arm. When the motor turns the upper arm rotates around the pulley. The motor for the elbow-b joint is mounted high up the arm, and uses extra pulleys to drive the elbow-b joint. The lightweight construction of the arm is also visible, with the side plates hollowed out to save weight, and most of the arm volume being air. The elbow joint is offset in order to allow the arm to curl up, as shown in figure A-18.

The similar configuration for the wrist joints is shown in figure A-19. The upper motor drives wrist-a, and the lower motor drives wrist-b through some extra pulleys. The lower arm attaches to the upper arm using a clamp, and the hand attaches to the arm also using a clamp. These clamps allow the arm to be easily dismantled and assembled.





Figure A-11: Arm joint configuration I. The figure shows schematically the layout of the series elastic actuator at the elbow-a and wrist-a joints. The motor (gearbox - 1, motor -2, encoder -3), is arranged to be parallel to the spring (8). The motor drives a small pulley (11), which is connected with a cable drive to the large pulley (9). One end of the spring is connected to the pulley with the piece (10), and the other end is rigidly connected to the tube (7), (12). The pulley supported on the tube with a bearing (6), and so is free to move relative to the tube. The tube is itself held with two bearings (6), between two supports (7), with side plates (4) holding the whole thing together. If the right hand end of the tube is grounded, the action of the motor is to walk round the pulley, moving the whole motor-support-side-plates block around the tube. The spring is between the motion of this block, and the fixed tube.






Figure A-13: The same bearing configuration was used for all the joints of the arm, the only difference being the sizes of the various tubes. The tube is supported on two angular contact bearings (15). These are preloaded by shoulders on the tube (7a), shoulders on the supports (5a), a spacer (13) and a lock ring (14). The bearings are far apart, and can withstand axial and radial loads between the tube and the supporting structure. The pulley (9) needs to move relative to the tube to accommodate deflections in the spring. It does not have large axial or radial loads, so it is supported on a radial bearing.




Figure A-14: Cable tensioning mechanism. The same tensioning mechanism was used for all the cables in the arm. One end of the cable (21) is terminated in a slot in the large pulley (9a). It wraps around the large pulley (9), around the motor pulley (11), and terminates in the tensioning block (23). The block can be moved by the tensioning screw (24), so adjusting the cable tension. The screw is part of a metal piece (25), which doubles as a connection between the end of the spring (8), and the large pulley.




Figure A-15: The left hand figure shows the cable configuration used for the elbow-b joint. An extra pulley (26) is used to transfer the cable drive from the motor pulley (11) to the large pulley (9). This allows the motor shaft to be perpendicular to the axis of the joint, freeing up the motor position. A similar organization is used for the wrist-b joint, shown in the right hand figure. Here two pulleys are used to achieve the same effect.

A.6 Arm technical data

The performance of the various joints is illustrated in table A.1.

Joint	Power (W)	Volts	Max torque (Nm)	Max speed (rpm)	Range of motion (degrees)
shoulder-a	70	48	20	30	230
shoulder-b	70	48	20	30	180
elbow-a	35	48	13	40	180
elbow-b	35	48	13	40	180
wrist-a	6	12	4	30	180
wrist-b	6	12	4	30	160

Table A.1: REAL Table showing technical data for the individual arm joints

A.7 Sensors

Each joint of the arm has a strain gauge mounted on the spring which measures the deflection of the spring, and so the torque transmitted by the motor, as well as a potentiometer to measure the joint position. These sensors are shown in figure A-20. Hollow shaft potentiometers are used with a plastic conductive resistive element (18) attached to the supports (5). The tube of each joint (7) has a plastic ring attached (20), which supports the potentiometer wiper (19). As the joint turns, the wiper moves on the potentiometer and records the joint position. Mounting the potentiometer




Figure A-16: Robot shoulder joint. The photograph shows the shoulder joint and the upper arm. The shoulder-b joint is shown most clearly with the motor, pulley, tensioning mechanism tube and clamp all in the middle of the picture. The upper arm is clamped onto the shoulder-b tube in order to allow quick assembly.




Figure A-17: Photograph of the upper arm of the robot. Both motors are placed high up on the arm, so reducing the inertia of the link. The elbow-a motor is arranged as described in figure A-11 walking around the tube, and driving the whole upper arm around. The elbow-b motor drives the joint through a cable mechanism with an extra pulley. The lightweight construction of the arm is visible, with the side plates hollowed out, and most of the volume of the link being air.




Figure A-18: Figure showing the arm curled up. The elbow-b joint is offset from the axis of the arm to allow the arm to curl up into this position.




Figure A-19: Lower arm design. The upper motor drives the wrist-a joint, moving the whole lower arm section around its longitudinal axis. The lower motor drives the wrist-b joint using two extra pulleys. The lower arm is attached to the elbow-b joint using a clamp to facilitate assembly.

in this way is neat; the sensor is tucked behind the motor pulley. Each motor also has an encoder measuring the position of the motor shaft. This is only used by the control law as part of the motor model (see figure A-9).

Each of these sensors requires wiring for both power and the signal. A complete actuator needs four separate cables for the strain gauge, encoder, potentiometer and motor power. With six degrees of freedom, and the need to pass these cables through or around the moving joints of the arm, electrical cabling is a fundamental design challenge.

To reduce the number of duplicated power cables for the different sensors, and in order to perform some sensor processing locally on the arm, a sensor board was constructed. This board is shown in figure A-21, and serviced the sensors for two joints on the arm. The board was mounted between the two side plates of the lower and upper arms, and between the plates on the shoulder. The connections to the board are shown in figure A-22 and the schematic for the board is included in figure A-23. The board fulfilled the following functions:

- **Strain gauge processing.** The integrated circuit which provided a stable voltage for the strain gauges, and processed the strain gauge output was included. Tuning potentiometers were included to alter the gain of the amplifier, as well as its offset. The Analog-to-Digital (A/D) converter used on the motor controller boards has a voltage range of 0 – 5 Volts. This means that input signals must have a zero reading of 2.5v and swing by ± 2.5 v around that zero. The sensor board also included a buffer on the strain gauge chip output in order to drive the signal along the wire to the motor controller.
- **Strain velocity.** The force control loop requires signals recording the strain as well as the rate of change of strain in order to control the force accurately. The rate of change of strain (Strain velocity) could be calculated digitally, but with only an 8 bit A/D on the motor controller, the errors would be large. The sensor board thus included an analog differentiation circuit and buffer to produce this signal directly.
- **Joint angle.** The board included an amplifier and level shifter so that the maximum, minimum and zero voltage for the joint angle motion could be set. It also included a differentiating circuit to calculate joint velocity. Setting the joint angle on the board sets the calibration for all the joints. This means that the robot can be powered on and off without having to recalibrate the joint limits. This means that the calibration is relative to the voltage in the system, so a well regulated power supply is required.
- **Encoders.** The board provided power distribution for the motor encoders.
- **Spare channels.** The board provided two extra connectors, connected back to A/D channels on the motor controller boards. The connectors for these have power and ground also, which makes adding extra potentiometers, touch sensors etc. easy.

A.8 Electrical wiring

Even with the sensor board, there are still a large number of cables required for each set of joints. The sensor board reduces the number of cables to a bundle of 25 wires for the sensors, and two 2 wire cables carrying the motor power. The wiring diagram is shown in figure A-24. The 25 pin sensor cables are connected to a expansion board which fans out the various signal wires to the motor controller boards. The modularity of the system is preserved by using connectors between the joints. This means that disassembling and arm consists of removing the clamps connecting the link to the other link, and disconnecting a few cables.

In addition to the sensor board, a debugging board was fabricated which could be inserted anywhere along the chain of sensor cables. This board had LED's and test points which could be




Figure A-20: In addition to the encoders on the backs of all the motors, each joint has a strain gauge (17) mounted on the spring (8) which measures the torque at the joint, as well as a potentiometer. The potentiometer is of a hollow shaft construction, with the resistive element (18) forming a ring mounted on the support (5). A plastic ring (20) clamped to the tube (7) holds a wiper (19), which completes the potentiometer. The measurement is of the angle between the tube and supports, so measuring the angle of the joints.




Figure A-21: Photograph of the sensor board. The large chips process the strain gauges, and the tuning potentiometers are used to adjust gains and offsets for all the signals. The power to the board and the outgoing signals are carried by the large connector.

used to test all the signals coming from the sensor board. This board together with the modularity of the whole system proved invaluable for debugging.

Physically routing the cables through the robot is somewhat of an art form. The arm was designed so that the majority of cables could pass through the axes of the joints as shown in figure A-25. This reduces the strain on the cables, and makes them less likely to break. The wires were passed through the joints for the wrist and elbow units, but at the shoulder there were too many wires, so they were passed around the joints.

A.9 Control hardware

The control for the arm joints was implemented using a home made motor board, illustrated in figure A-26. This board contains all that is necessary to control two joints. The heart of the board is a Motorola 6811 microcontroller which runs the control program. This is written in assembly language and runs with a loop sampling rate of 2kHz. The 6811 has A/D converters which are used to sample the analog inputs (joint angle, velocity, strain, rate of change of strain). Extra circuitry processes the signals from the motor encoders. The board uses a PIC to generate the Pulse Width Modulated (PWM) signal which is used to command power to the motors. An H-bridge motor amplifier is also included on board which provides power to the motors.

The board communicates with a network of Motorola 68332 processors which perform all the higher level processing for the arm, as shown in figure A-27. The communication is low bandwidth, with the 68332's and 6811's exchanging 13 byte packets at around 50Hz. This motivated coding the control law for force as well as the spring-like joint behavior in the 6811's. The 68332's thus send commands such as the setpoint of the joint as well as the stiffness and damping at 50Hz. Data from the robot, joint position, torque, etc. is thus also read at 50Hz.

The 68332 processors run a version of Common Lisp called L (Brooks, 1995). L is a multi-




Figure A-22: Each joint has a strain gauge, which requires 6 wires, an encoder (4 four wires) and a potentiometer (3 wires). These all require shields indicated by the + 1. Rather than pass these cables all the way up the arm, a sensor board was used to process the signals and package the signals into a fewer number of wires and cables. Each board manages two joints worth of electrical signals. Although only 20 wires are required to carry all the signals, the connector had 25 pins, so the extra wires were used to carry power and ground to the board.




Figure A-23: Schematic of sensor board. The connectors for potentiometer, strain gauges, encoders and spare channels are shown at the bottom, and the output connector at the top. This schematic shows the circuit for one joint. The circuit for the second is identical.




Figure A-24: Wiring diagram. The top half of the figure shows how the sensor cables from each set of joints are passed back along the arm to connect with expansion boards, and then to the motor controller boards (marked "6811"). Each sensor board handles the signals from two joints, and each expansion board handles two sensor boards. Each motor controller can drive two motors. The lower half of the figure shows the cables for the motor power. The dashed lines indicate where the arm can be easily split apart, with the boxes aligned with the lines indicating connectors. Thus the wrist joint can be removed by disconnecting only two connectors.




Figure A-25: Figure showing how the sensor board (30) is mounted between the supports. The cables from the sensor board (29), and those carrying power to the motor (28) are passed through the tube and so up the arm. Arranging the cables in this way reduces strain on the wires, and reduces the likelihood of cable breakages.

processing LISP environment, which was convenient as a language to control the arms. The controller for the arms was partitioned into different processes which communicated using shared variables. The L Listener provided a simple way to connect, examine, spawn and kill processes. The 68332's are organized in a network as shown in figure A-27. The connection is also low bandwidth, using dual-ported RAMs (DPRAMs) to share information.

For example, in the configuration shown in figure A-27, the 68332 labeled 'POSTURE' communicates the posture of the arm to the two processors which run the oscillators ('LEFT OSC' and 'RIGHT OSC'). The 'LEFT OSC' and 'RIGHT OSC' processors can communicate to implement connections between the oscillators (as described in chapter 5). The command to the motors is passed to the 'LEFT ARM' and 'RIGHT ARM' processors. Sensory signals from the arms are passed through the 'ARM' processors to the 'OSC' processors.

The network of processors is distributed with low bandwidth communication both within the network, and also with the arm joint controllers. This arrangement motivated some of the solutions found in this thesis. The example has no global clock, or global synchronization. There is also no direct communication between the joints.




Figure A-26: Photograph of motor controller board. Each board contains all the circuitry to drive two series elastic actuators. Two Motorola 6811 microcontrollers implement the control law for the actuators, with extra circuitry to process the signals from the motor encoders, and generate PWM signals for the motor drives. A power amplifier is also included on board (at right hand end in this picture). The 6811's communicate with the higher level processors using a differentially driven serial line, which connects to the large connector at the left hand end.

The code for the 68332's is written and compiled on a Macintosh PowerPC. This computer acts as a Front End Processor (FEP) for the network. The serial line is used to communicate with the network, the serial communication being mediated by another 68332, dubbed the 'InterFEP'. This configuration is illustrated in figure A-28.

The DPRAM interface is used to connect to other parts of the robot. This is also shown in figure A-28. For example, for the drumming in chapter 7, the auditory signal was detected using a sound card in a PC running QNX, a real time UNIX. The auditory signal was communicated through the DPRAM to the 68332 network. Referring to figure A-27, the auditory signal entered the network through the 'POSTURE' processor, then was passed to the 'OSC' processors.




Figure A-27: The figure shows the network of Motorola 68332 processors used for higher level control, connected to the 6811 joint controllers. The connection between the ‘—ARM’ processors and the 6811s is low bandwidth, with 13 byte packages exchanged at around 50Hz. The communication in the network is also slow, with dual ported RAMs (marked DP) used to share information between the processors. The ‘POSTURE’ processor was used to generate the desired posture of the arm, the ‘—OSC’ processors integrated the oscillator equations, and the ‘—ARM’ processors communicated with the joint level controllers.




Figure A-28: The code for the 68332 network is written and compiled on the Front End Processor (FEP), which is a Macintosh PowerPC computer. The FEP communicates with the network using its serial port, the communication mediated by another 68332, the ‘InterFEP’. The 68332 network can also receive input from other computers associated the robot. The robot vision system uses a network of C40 Digital Signal Processors (Scassellati, 1998), which are attached to a network of PC’s running QNX, a real time UNIX operating system. These are patched into the 68332 network using the dual ported RAM interface.

A.10 Hand design

The arm design does not include a design for the hand. In practice it was found to be relatively easy to modify a simple paddle into a special purpose hand to grip cranks, hammers etc. without requiring the extra design, control and weight of a dextrous hand. Velcro was used extensively as a quick and surprisingly strong method of attaching objects to the hand of the robot.

A.11 Conclusion

This appendix has described the design of the arms used in this thesis. They are robust due to their mechanical design and to the choice of actuators. The control system is somewhat slow and low bandwidth, but given the stiffness of the arm this is not a difficulty. The low bandwidth and slow communication forces the control system implemented to be robust, which is what has been demonstrated in this thesis.

Appendix B

Arm and oscillator parameter tuning

This appendix covers some aspects of tuning of the oscillator parameters. It describes the parameters which need to be set, the automatic tuning methods which were used, and a sample procedure for implementing a new rhythmic task using oscillators.

B.1 Parameters

For two arms with six degrees of freedom there are inevitably a large number of parameters that need to be set. However, most of them can be fixed and not altered in normal operation. The parameters are described below.

- **Force control gains.** These parameters control the performance of the force control in the series elastic actuators, described in detail in appendix A. The parameters can be calibrated once and then fixed. The calibration procedure involves tuning the PD control of force, usually by commanding a square wave of force, and observing the closed loop behavior. The calibration of the joint angles and velocities (described in appendix A) does not need to be particularly accurate, mostly because the control implemented on top (the oscillator system) does not require the arm to be accurately calibrated.
- **Arm stiffness/damping.** These parameters control the arm dynamic properties. There is an upper limit on the arm stiffness due to noise in the sensors, as well as bandwidth limitations of the force control. There is also an upper limit on damping, due to noise in the joint angle sensors. The stiffness and damping for each joint of the arm can be set to reasonable values and effectively fixed. Occasionally it is useful to alter the stiffness and damping of individual joints, in order to reduce or accentuate their motion.
- **Arm posture.** The arm posture is important for the overall system behavior, and can be easily set with one parameter per joint.
- **High pass filters.** There is a high pass filter on each oscillator input which removes the DC component from the signal. In fact a low pass filter is used to calculate the offset and this is subtracted from the input signal. A simple single pole digital filter was used with a cut-off frequency of 0.2 Hz. This was fixed for all the oscillator experiments.
- **Choice of oscillator input.** In the work for this thesis there are two possible choices of input parameter to the oscillators. The two signals used are joint angle θ_i and joint torque u_i . There is no reason to restrict the input signal to these parameters (for the juggling task in chapter 3 the ball trajectory was used, and auditory feedback was used for drumming in chapter 7. For

any input signal, the describing function analysis described in chapter 3 is useful to determine the approximate effect of the feedback.

The describing function analysis can also be used to help choose between joint angle and joint torque. Williamson (1998b) and Williamson (1998a) describe the difference in some detail. Appendix D describes the effect of torque feedback on oscillators coupled through the natural dynamics (compared to the model in chapter 4).

The only major difference between the feedback types is that if torque feedback is used, it is possible to get a solution where the arm is not moving, but the oscillators are entrained. The oscillator outputs are entrained with internal forces. When the joint angle is used as input this obviously cannot occur. This means that although the torque feedback could be used for crank turning with one arm, it did not work with two arms. The torque feedback also appears to be more sensitive to mechanical coupling, being appropriate for the Slinky toy described in chapter 4 and appendix D.

- **Oscillator parameters.** The main parameters that need to be set are the tonic c in equation (3.1) which determines the amplitude of the joint level oscillation, and τ_1 which determines the oscillator frequency. The third parameter is the input gain which as discussed in chapter 3 is rather robust. As long as this gain is large enough to cause entrainment, its particular value does not appear to be important.

B.2 Automatic tuning

Since for most of the tasks described in this thesis, the system was set up to drive the resonant mode of the arm-task system, the posture of the arm and the amplitudes of the oscillator motions were the most important parameters to set. These have to be roughly correct to get the required motion, but their exact values are not important. The other oscillator parameters are easy to set, with their values being relatively unimportant.

The initial values of posture and amplitude can be set by hand to match a particular task, say for example a new crank turning configuration. These values can then be modified using an automatic tuning scheme, measuring the performance in the task, and adjusting the parameters accordingly.

The tightly coupled nature of the system makes most supervised learning techniques difficult to apply. It would be difficult to calculate an “error” and then relate that error to the parameters involved, especially given the different time-scales involved. The timing issue, as well as the rather large state space makes reinforcement learning techniques also not appropriate. The best learning techniques for this problem appear to be genetic algorithms, evolutionary strategies (Bäck and Schwefel, 1993) and hill climbing (e.g. Rich and Knight (1991)). Initial experiments on a two degree of freedom simulated arm suggested that the most effective strategy was hill climbing, and this was implemented to tune the parameters of the real arms.

The learning algorithm was then

1. Make a small random change to the parameter vector.
 - A number of strategies were tried here. The most successful was a technique called “run-and-twiddle”. In this method the parameter vector is changed, and if the change causes an increase in performance, then further changes to the parameter vector are made in the same direction. This is continued until changing in that direction no longer gives an increase in performance, at which point a new direction is tested.
2. Apply these parameters and test the performance.
 - The performance evaluation generally included terms directly related to the task, as well as terms related to the energy required to produce the motion. This then optimized for

efficient motion. When using two arms the evaluation balanced the work between the two arms by including terms for the total energy, and the difference in energy between the two arms. The performance evaluation was sometimes noisy making the learning less robust. The noise can be reduced by measuring performance a number of times and averaging the results.

3. If the performance is better than ever before, keep the new parameter set and go to step 1. If not, revert to old parameter vector and go to step 1.
 - To prevent the learning being stuck in local minima, the new parameter setting was accepted if the performance was within some small percentage of the maximum performance. Simulated annealing could be used to apply this idea more rigorously, although the simpler technique appeared to work well.

It is indicative of the robustness of the oscillator parameters that a simple learning algorithm like this was successful for tuning. This was even the case when 10 or 12 parameters were tuned simultaneously. In the crank turning case (chapter 4), and also the hitting application (chapter 7) there appeared to be a variety of parameter setting that gave efficient performance, which the hill climbing algorithm was fairly quick to find.

B.3 Procedure for new task

The procedure of tuning the arm for a new task is described below, taking as an example the task of pumping a bicycle pump.

1. Connect pump to arm, either with hand or special attachment.
2. Manoeuvre arm to approximately correct posture for the pumping motion.
3. Adjust oscillator time constants for approximately the correct frequency.
4. Adjust oscillator amplitudes to create motion.
5. Adjust feedback gains to cause entrainment of the system. At this point it is easy to determine which feedback type (angle or torque) works better.
6. If the system does not operate as intended: adjust posture, oscillator amplitudes, maybe stiffen unused joints, adjust arm stiffness and damping.

If the automatic tuning is then required to further optimize the system performance, then

1. Find way to measure performance of the task. In the bicycle pump example this might be a measure of the linear motion of the pump, or the airflow out the pump, as well as an estimate of the energy expended by the arm to perform the task.
2. Select the parameters which should be tuned. The most important parameters are probably the posture and the tonics (controlling the oscillator amplitudes), not only because these dictate whether motion will occur, but also the quality of the final motion (matched to task, too violent etc.).
3. Run hill climbing on the most important parameters.

Appendix C

Oscillator Inputs

This appendix considers the effect of using linear inputs to the oscillator instead of the usual non-linear inputs. Using linear inputs is not a good idea because they reduce the robustness of the system.

The oscillator used throughout this thesis is governed by the following equations:

$$\tau_1 \dot{x}_1 = c - x_1 - \beta v_1 - \gamma [x_2]^+ - \Sigma_j h_j [g_j]^+ \quad (\text{C.1})$$

$$\tau_2 \dot{v}_1 = [x_1]^+ - v_1 \quad (\text{C.2})$$

$$\tau_1 \dot{x}_2 = c - x_2 - \beta v_2 - \gamma [x_1]^+ - \Sigma_j h_j [g_j]^- \quad (\text{C.3})$$

$$\tau_2 \dot{v}_2 = [x_2]^+ - v_2 \quad (\text{C.4})$$

$$y_i = [x_i]^+ = \max(x_i, 0) \quad (\text{C.5})$$

$$y_{out} = [x_1]^+ - [x_2]^+ = y_1 - y_2 \quad (\text{C.6})$$

The input is applied through a max operator, so that the input inhibits the equation. For example, when the input g_j is positive, its only effect is in the x_1 equation, making $\dot{x}_1$ more negative. The input can be applied without this non-linearity, replacing $\Sigma_j h_j [g_j]^+$ by $\Sigma_j h_j g_j$ and $\Sigma_j h_j [g_j]^-$ by $-\Sigma_j h_j g_j$.

The major effect of not using the non-linearity on the oscillator input is that the oscillator output amplitude becomes dependent on the size of the input. This is illustrated in figure C-1, showing the output amplitude of the oscillator with and without non-linear inputs. This affects the Bode plot for the oscillator (shown in figure C-2), making the gains higher, and with a smaller range in gains than for the non-linear input case. The phases for the two cases are approximately the same.

By reducing the variety of gains possible from the oscillator, using linear inputs reduces the robustness of the system. In addition, it couples together the parameters which control input gain (h_j) with those controlling output amplitude (c), so making the tuning more complicated. Since the output is not dependent on the input, the stability results discussed in chapter 6 do not hold in their present form. A more rigorous determination of the H_∞ norm of the oscillator is required before the behaviors of systems coupled to an oscillator using linear inputs can be predicted.




Figure C-1: SIM Plot showing effect of using non-linearity on oscillator input on the oscillator output amplitude. When the non-linearity is used, the oscillator output is approximately constant (solid lines, where the multiple lines correspond to different frequencies). When the inputs are applied in a linear fashion, the output amplitude is approximately proportional to the input amplitude (dash-dot lines).




Figure C-2: SIM Describing function of oscillator with and without linear inputs. The top plots show the Bode plot of the oscillator with linear inputs, and the lower plots with non-linear inputs. When linear inputs are used, the gain of the oscillator is higher, and covers a lower range of gains than with the non-linearity. This reduces the robustness of the oscillator to changing gains and system parameters. The phases in both cases are similar.

Appendix D

Oscillators and torque feedback

This appendix describes how the simple model of coupling introduced in chapter 4 is modified when the joint torque, rather than the joint angle is used as the feedback signal to the oscillators. The results are the same with the different feedback, the oscillators finding the resonant mode of the underlying mechanical system. The use of this feedback is illustrated with the example of coordinating two arms of the robot to pass a Slinky toy from hand to hand.

D.1 Torque feedback




Figure D-1: A simple model of coupling through the natural dynamics. The model consists of two masses driven by oscillators, connected by a coupling spring k_T .

The model of coupling when force or torque feedback is as shown in figure D-1. The only difference is that instead of using the joint angle θ , the oscillator uses the joint torque $u = k(\theta_v - \theta) - b\dot{\theta}$ as the feedback signal. Since each joint is effectively a mass-spring system, the joint torque is related to the joint angle:

$$u = k(\theta_v - \theta) - b\dot{\theta} = m\ddot{\theta}$$

Using the describing function analysis, the effect of the oscillator can be written:

$$\theta_v = ge^{j\phi}u = ge^{j\phi}(-m\omega^2)\theta \quad (D.1)$$

where m is the mass of the system. This alters the equation of the example coupled system to:

$$\begin{pmatrix} k_1 + k_T + j\omega c_1 - m_1\omega^2 & -k_T \\ -k_T & k_2 + k_T + j\omega c_2 - m_1\omega^2 \end{pmatrix} \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} = \begin{pmatrix} -m_1\omega^2 k_1 g_1 e^{j\phi_1} \Theta_1 \\ -m_2\omega^2 k_2 g_2 e^{j\phi_2} \Theta_2 \end{pmatrix} \quad (D.2)$$

For steady state solutions, the damping terms must be zero:

$$\omega c_1 + m_1\omega^2 k_1 \text{Im}[g_1 e^{j\phi_1}] = \omega c_2 + m_2\omega^2 k_2 \text{Im}[g_2 e^{j\phi_2}] = 0 \quad (D.3)$$

which implies that the imaginary parts of $g_i e^{j\phi_i}$ must be negative, and vary with $1/\omega$. Removing the damping terms results in the eigenvalue problem:

$$\begin{pmatrix} \frac{k_1 + k_T}{m_1(1 - \text{Re}[g_1 e^{j\phi_1}])} & \frac{-k_T}{m_1(1 - \text{Re}[g_1 e^{j\phi_1}])} \\ \frac{-k_T}{m_2(1 - \text{Re}[g_2 e^{j\phi_2}])} & \frac{k_2 + k_T}{m_2(1 - \text{Re}[g_2 e^{j\phi_2}])} \end{pmatrix} \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} = \omega^2 \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} \quad (\text{D.4})$$

The torque feedback makes the effect of the oscillator on the coupled system dynamics more apparent than in the angle feedback case. The effect of the oscillator is to alter the apparent inertia in the system, changing the effect of mass m_1 to $m_1(1 - \text{Re}[g_1 e^{j\phi_1}])$, and similarly for mass m_2 . Because of the oscillator properties, the values of phase for both oscillators will be similar, so $(1 - \text{Re}[g_1 e^{j\phi_1}]) \approx (1 - \text{Re}[g_2 e^{j\phi_2}])$. This results in the effect of the oscillator being a common factor in (D.4), so only effecting the eigenvalues (the final frequency), not the eigenvectors (the mode shape). The solutions to the damping constraint (D.3) have imaginary parts which are negative, which due to the oscillator properties implies that the real parts will also be negative. This means that the oscillator increases the apparent inertia of the system i.e. $(1 - \text{Re}[g e^{j\phi}]) > 1$.

The oscillator solutions are also locally stable using this feedback method, which is seen by looking at the effect of amplitude changes on the damping in the system (see figure D-2).




Figure D-2: SIM Local stability under torque feedback. The plot shows $N(j\omega, A)$ and the line $j2\zeta/(\omega\omega_n) = j c_1/(m_1 k_1 \omega)$ plotted on the complex plane, with a sample limit cycle marked with a $\square$. The limit cycle can be shown to be locally stable by examining changes in the system damping as the amplitude of oscillation A is varied. An increase in amplitude (in the direction of the arrow) results in $N(j\omega, A) > 2\zeta/(\omega\omega_n)$, which is equivalent to an increase in damping by (D.3), and so a reduction in the amplitude. A reduction in amplitude has the opposite effect, showing that the limit cycle is locally stable.

D.2 Slinky example

The task chosen is the coordination of two arms of the robot using a Slinky toy. This is a good example because it illustrates that the oscillators are sensitive to weak coupling as opposed to the strong coupling in the crank turning case. As the Slinky is passed from hand to hand the weight

on the hand varies, with the difference in weight being greatest when the hands are furthest apart. This behavior could be approximated by a coupling spring which is weak and has a negative spring constant, giving a larger force away from the central position. This can then be approximated by the coupling model.

When the Slinky is placed between the two arms of the robot as shown in figure D-3, and the joint torque is used as the feedback signal to the oscillators, the oscillators quickly find the anti-phase motion. An in-phase motion is sometimes observed, but it does not appear to be a stable solution, as small perturbations result in the out of phase motion. A sample time trace for the motion of the two arms is illustrated in figure D-4, showing the effect of the torque feedback in coordinating the two arms.




Figure D-3: Picture of Cog passing the Slinky toy from hand to hand. The two elbow joints are used to move the hands up and down, where the coordination between the hands is given by the interaction between the oscillator dynamics and the coupled arm-slinky system.

Data from the Slinky was taken, recording the frequencies, amplitudes and phases of the joint motions as the natural frequency of the oscillator was varied. This data could then be used to directly calculate the gain of the oscillators, and the mode of the Slinky. The process followed is similar to that for the crank turning example described in chapter 4.

The damping constraint in the coupling model predicts that the imaginary part of the oscillator action $\text{Im}[g \exp j\phi] \propto 1/\omega$, which is plotted in figure D-5. The constant of proportionality should be $-c/(km)$, making the slope negative. The data lies approximately on a straight line, with data at very slow speeds not fitting on the line. The slope is negative as predicted by the theory.

A plot of the real parts of the oscillator actions is included in figure D-6. The model predicts that the real parts should be proportional to $m - k_T/\omega^2$. The graph shows the real part of the oscillator action plotted against $1/\omega^2$. The data falls on a positive slope line at high speeds, indicating that the coupling stiffness k_T is negative. At low frequencies the points do not lie on the line but appear to fan out. This may be due to a change in the coupling type of the Slinky. At high speeds the spring-like properties of the Slinky complement the transfer of mass from arm to arm, while at low frequencies the coupling will only come from the mass. This may reduce the strength of the coupling, giving rise to the points not lying on the same line.

These results show that the Slinky operation is a resonant mode. The oscillator quickly responds to the resonant properties of the system and tunes automatically into the mode shape, which is




Figure D-4: REAL Two examples of Slinky operation. Both plots show the outputs from the oscillators as the torque feedback (dash-dot) is turned on and off. When the traces are in phase, the Slinky is moving in anti-phase. When the feedback is on, the two arms are coordinated and the outputs are synchronized, but when off, the oscillators are no longer synchronized. The only connection between the oscillators is through the physical structure of the Slinky.

passing the Slinky from hand to hand. The behavior is also robust as shown in figure D-7, the oscillators finding the correct coordination even when the oscillator time constant is doubled.

D.3 Conclusion

This appendix has shown that the effect of torque feedback on the oscillators is somewhat similar to joint angle feedback in multi-joint motions. In both cases, the oscillator finds the resonant mode of the system. The effect of the torque feedback is to increase the inertia of the system, as opposed to alter the stiffness in the angle case. This means that the overall motion is often slower, and the mode found by the oscillators is often the slower one in these cases. In addition, because the input to the oscillator is the joint torque, it is possible to have a solution where the arm is completely still, and the oscillators are pushing against one another, entrained with the internal forces.

The example with the Slinky toy showed that small forces can be used to entrain the oscillators, and that the analysis from the model is useful in understanding the nature of the coupling. The robustness of the system to parameter changes was also shown. The sensitivity of the oscillators to the resonant mode makes them useful for tasks where the mechanical coupling is weak as well as strong.




Figure D-5: REAL Plot of $\text{Im}[u_i^T K' G e^{j\Phi} u_i]$ against $1/\omega$ for the Slinky. The model predicts this to be a linear relation with slope $-c/(km)$. The data confirms this prediction, lying approximately on a line of negative slope. At low frequencies (high values of $1/\omega$), the data does not lie on the same line.




Figure D-6: REAL Plot of $\text{Re}[u_i^T K' G e^{j\Phi} u_i]$ against $1/\omega^2$ for the Slinky. The model predicts that this relation should have a slope $-k_T$. If the coupling from the Slinky is equivalent to a negative spring, then the slopes of the data should be positive. This is shown in the plot. At low frequencies the data does not lie on the same line, perhaps indicating that the type of coupling is changing to one dominated by the transfer of the mass of the Slinky from arm to arm.




Figure D-7: REAL Mode of Slinky operation. The top plot shows the amplitude θ_1/θ_2 of the Slinky motion versus frequency as the natural frequency of the oscillator is changed. The lower plot shows the phase between the two hands. Even though the frequency of the system approximately doubles over this plot, the mode remains approximately constant. This indicates not only the robustness of the oscillator solution, but also its ability to find the mode of the system as its final periodic solution.