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 $\dot{\theta}_v(t)$.

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

or

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

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

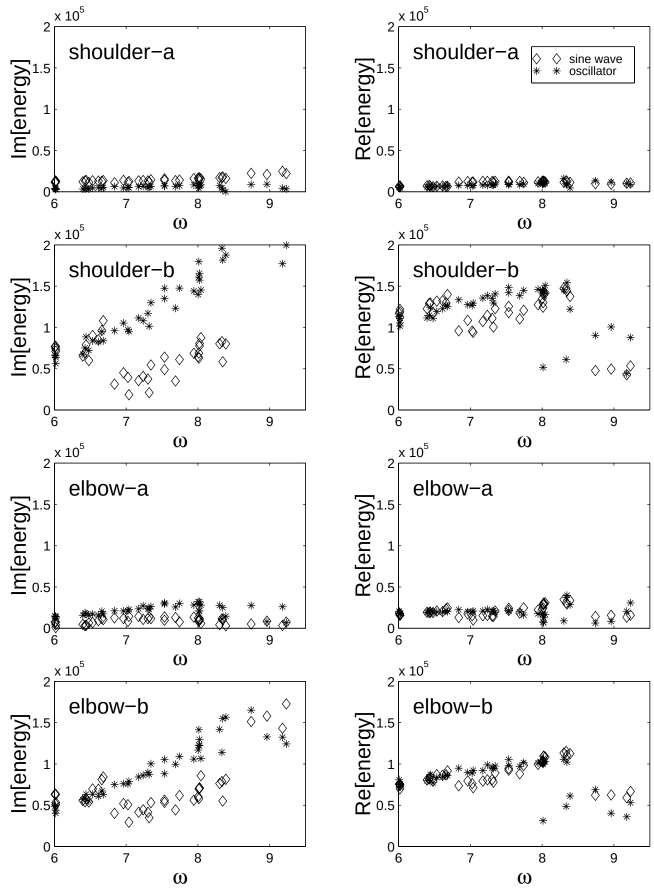

Figure 4-27: $\boxed{\text{REAL}}$ Plot showing the real and imaginary parts of the power plotted against frequency for the different joints of the arm during crank turning. The imaginary part of the power (left hand graphs) refers to energy stored in the system, which is generally higher for the oscillator ($*$) 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.

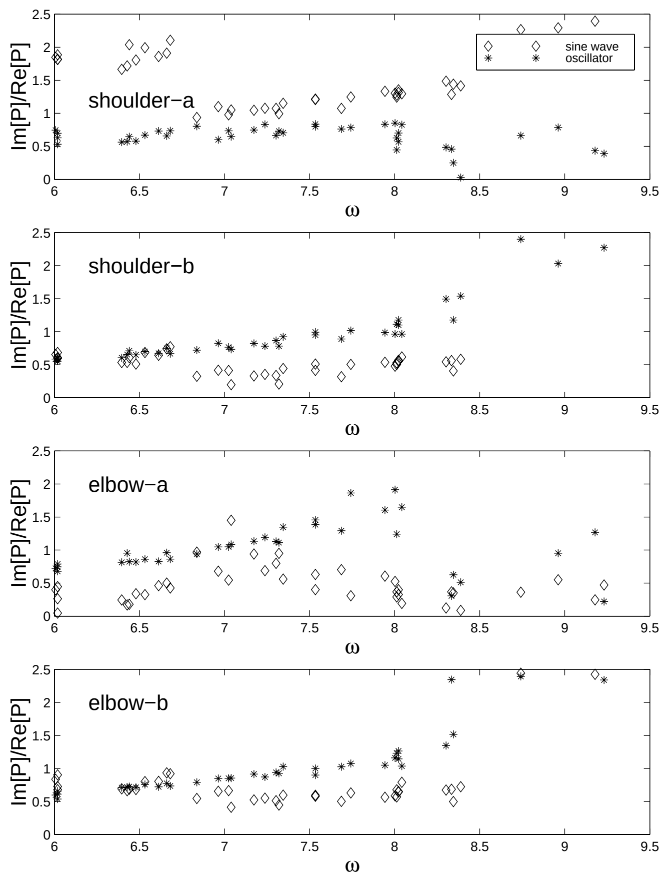

Figure 4-28: $\boxed{\text{REAL}}$ Plot showing the ratio between energy stored and energy used to turn the crank, for the oscillator ($*$), and the sine wave ($\diamond$). The ratio is higher for the oscillator in all cases except for the first shoulder joint. This shows that the oscillator solution 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 were 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

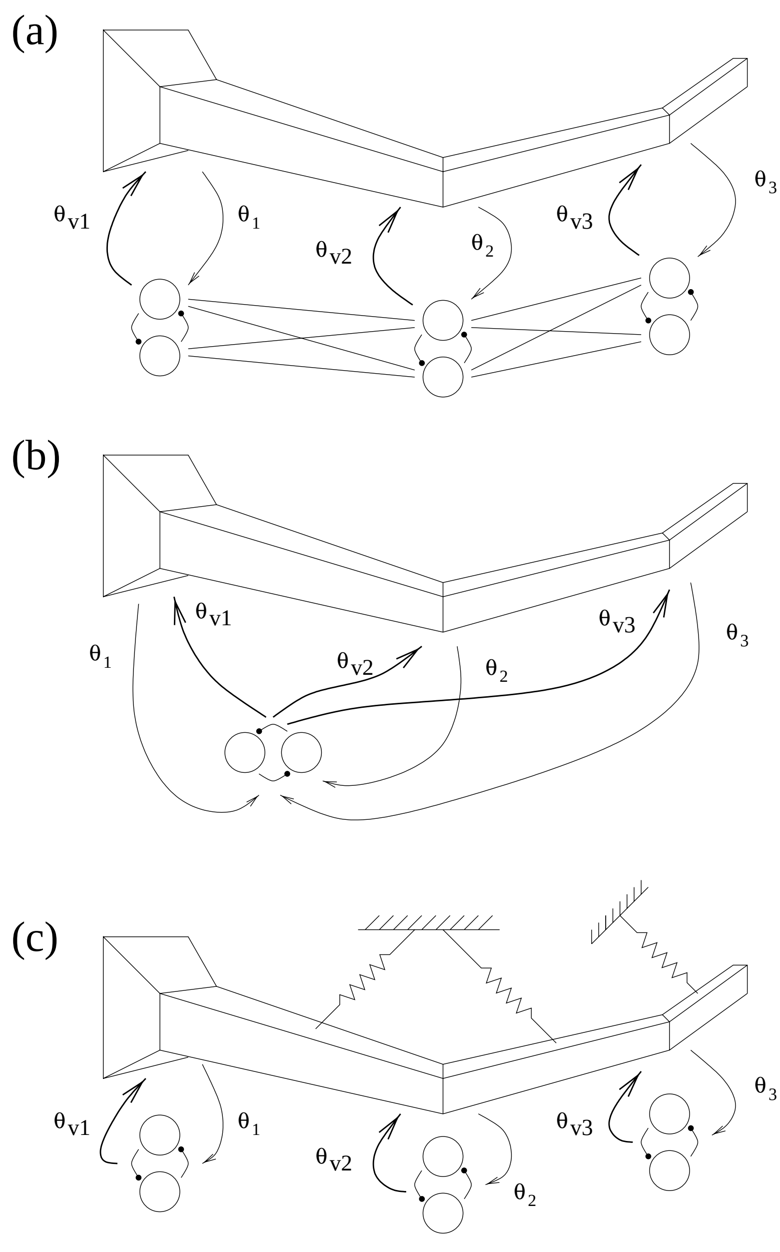

Figure 5-1: Three methods for coupling oscillators. **(a)** shows connections between the oscillators which enforce phase differences between the joints. **(b)** shows a method using a single oscillator to drive a number of joints as a single unit. **(c)** shows the case where the natural dynamics of 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.

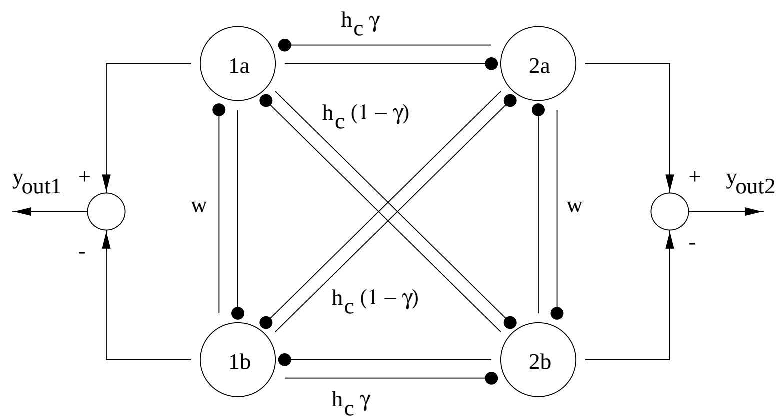

Figure 5-2: Simple oscillator network. The figure shows two oscillators with the usual mutual inhibition weights w and extra connections. The black circles on the end of the links correspond to inhibitory connections. The weights between the a neurons have strength $h_c\gamma$, and the cross coupled 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.

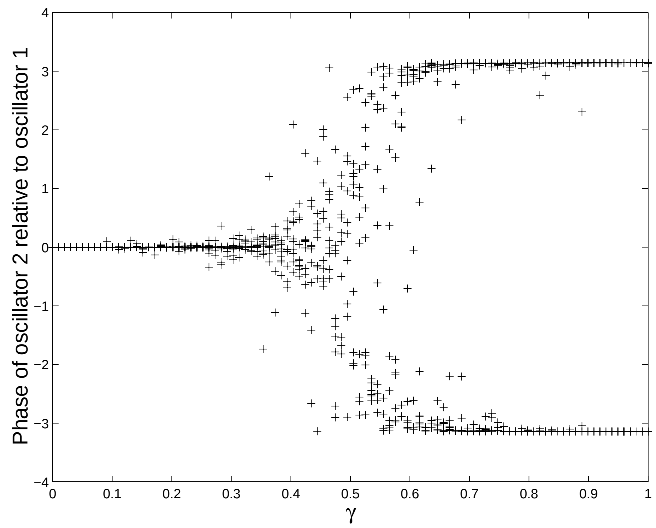

Figure 5-3: $\boxed{\text{SIM}}$ Plot showing how the relative phase between the oscillator outputs varies as a function of γ. The plot was produced by simulating the system from different initial conditions for each value of γ and calculating the phase using a Fourier transform. The oscillator outputs 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.

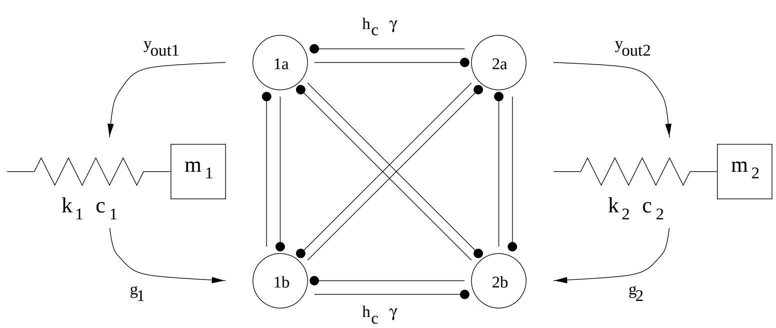

Figure 5-4: Figure showing two oscillators with external connections with parameters h_c, γ, driving two different mass-spring systems. Each oscillator is tightly coupled to a mass-spring system, with output y_{outi}, and input g_i. If the oscillators were not connected together, they would entrain with 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.

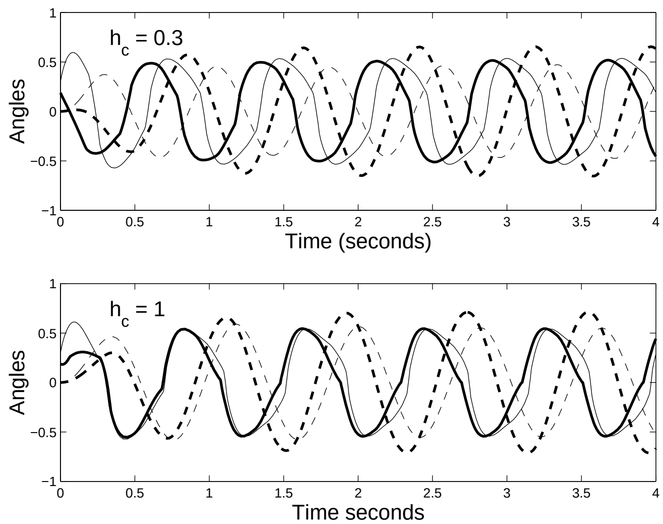

Figure 5-5: $\boxed{\text{SIM}}$ Coupled system transients. The plots show the motion of the the set-points for the two masses (solid lines) and the actual mass motion (dashed), for two values of input gain h_c and $\gamma = 0$. The thick lines refer to one mass-spring system, and the thin lines to the 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

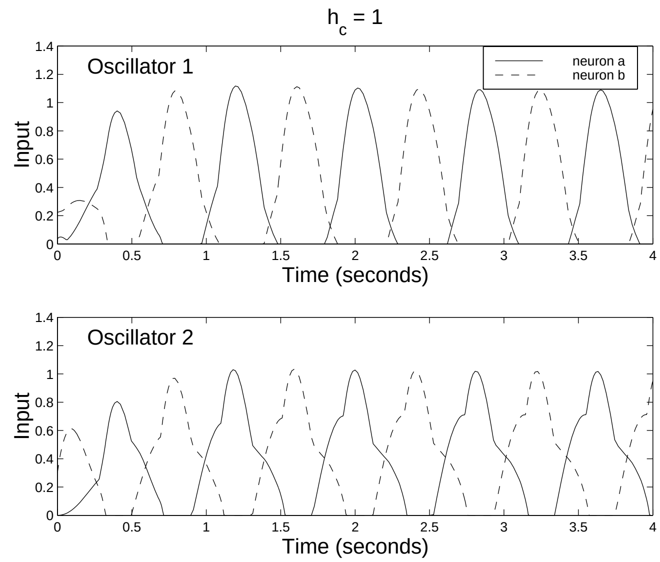

Figure 5-6: $\boxed{\text{SIM}}$ Oscillator inputs. The graphs show the inputs to the individual neurons during the motion in the lower graph of figure 5-5. The top graph refers to oscillator 1, lower to oscillator 2. The inputs are not sine waves, particularly for oscillator 2 where the combination 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.

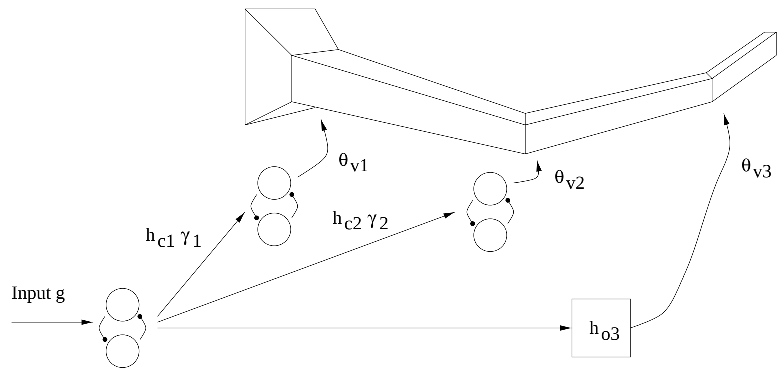

Figure 5-7: Master-slave network. A single oscillator is used to drive multiple joints of the arm using slave oscillators at the joints with explicit but unidirectional connections. Input is applied to the master, which specifies the sizes and phases for all the joints. For unidirectional connections, 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 ith joint thus becomes:

$$\theta_{vi} = h_{oi} y_o + \theta_{pi} \tag{5.3}$$

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:

$$g = h_{in} \sum_i \frac{\theta_i}{h_{oi}} \tag{5.4}$$

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.

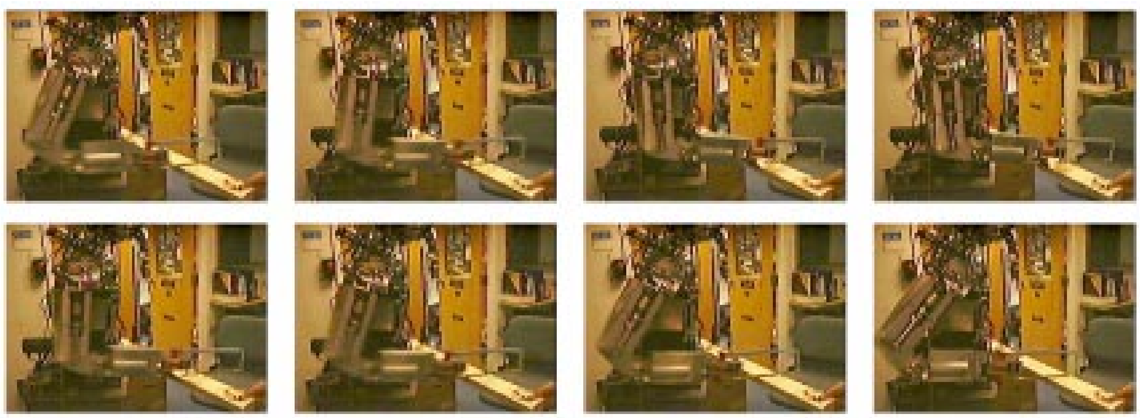

Figure 5-8: REAL Eight stills from the robot sawing a block of wood. The picture sequence runs from left to right. The pictures were taken at 0.1 second intervals and show the saw moving forward and backward, cutting the wood.

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

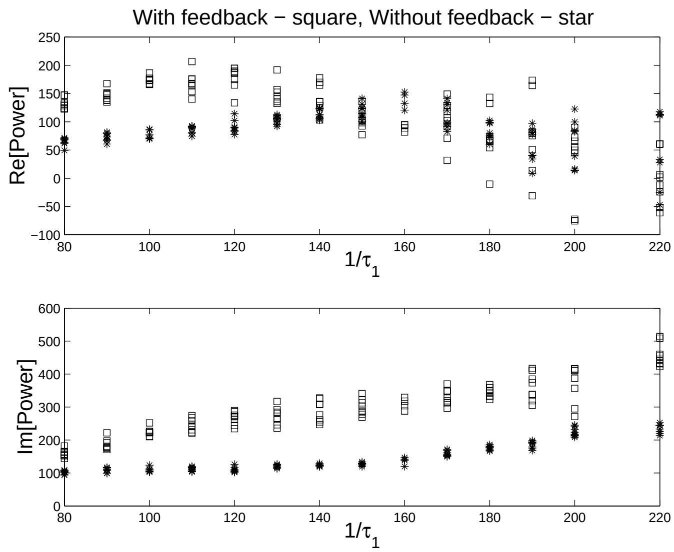

Figure 5-9: REAL Plot of energy expended in sawing versus oscillator natural frequency ($\propto 1/\tau_1$). The top graph shows the real power, i.e. the energy to overcome dissipation in the arm, and actually cut the wood, and the lower graph shows the imaginary power, which is stored and released in the springs in the arm. The □'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.

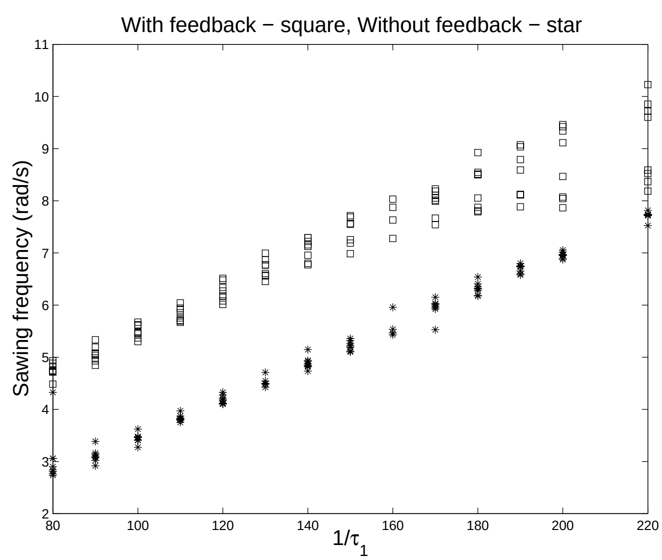

Figure 5-10: REAL Sawing frequency versus oscillator frequency for feedback on (□) and feedback off ($*$). With feedback, the sawing is faster, with a greater range of frequencies. This variation is because the oscillators respond to the varying resistance of the sawing task. When the resistance is high,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

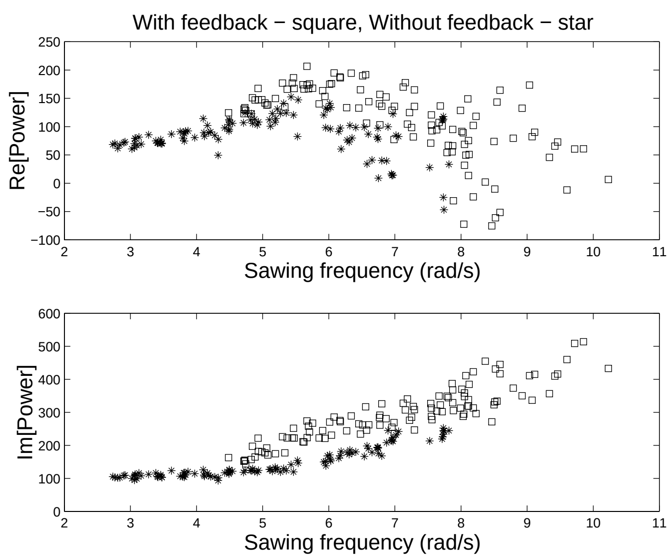

Figure 5-11: $\boxed{\text{REAL}}$ Plot of energy expended in sawing versus sawing frequency. The top plot shows the real energy required to perform the sawing. The real energy is slightly higher with feedback. The lower plot shows the energy stored and released in the arm springs, which is greater for 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 though 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.

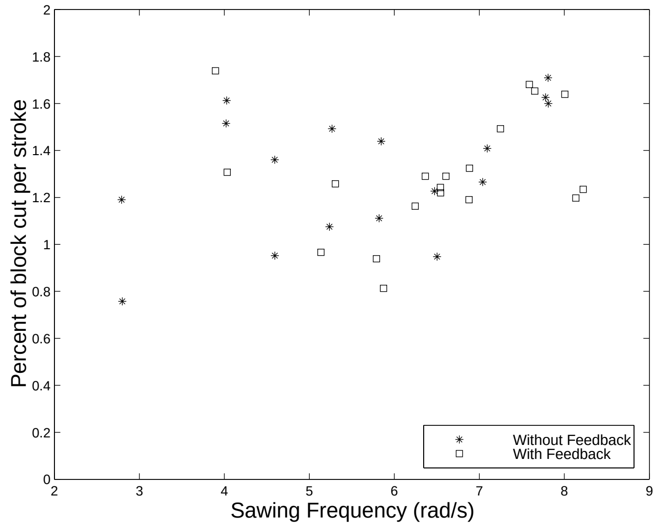

Figure 5-12: REAL Percentage of wood cut per stroke plotted against average sawing frequency. The time and number of strokes to cut a 0.5"by 1.5" block of wood was measured, and the amount of wood cut per stroke calculated as 1/number-of-strokes. The plot shows results using feedback (∗), and without feedback (□). 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:

$$\tau_{potential} = -\frac{\partial V}{\partial \theta_i} \tag{5.5}$$

This can be easily added to the usual joint torque:

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

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:

$$r = \sqrt{\theta_1^2 + \theta_2^2} \tag{5.7}$$

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

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

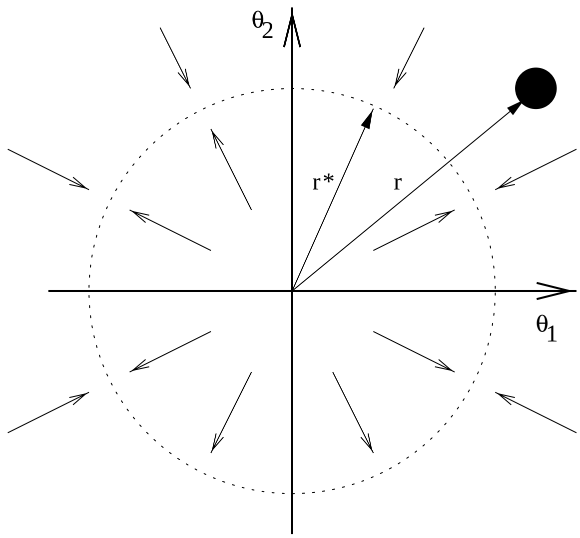

Figure 5-13: Application of potential field. The desired coordination between the two joints is that they describe a circle. The potential field required to create this applies forces to push the arm (represented by the black spot) out from the center, and in towards the dashed circle, with no forces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

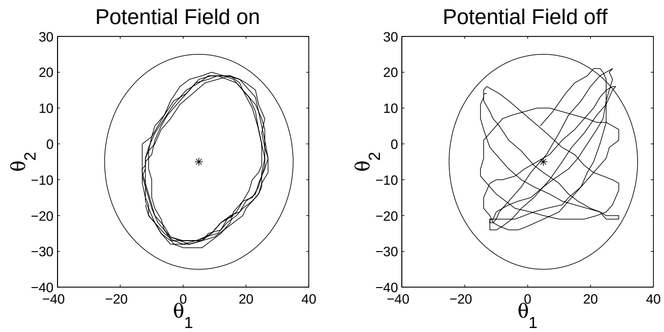

Figure 5-14: $\boxed{\text{REAL}}$ Effect of potential field. The graphs show the arm position of the real robot arm in terms of the two joint angles θ_1 and θ_2, with and without the potential field. When the field is on, it makes a resonant mode of the system which the oscillators find. When the field 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.

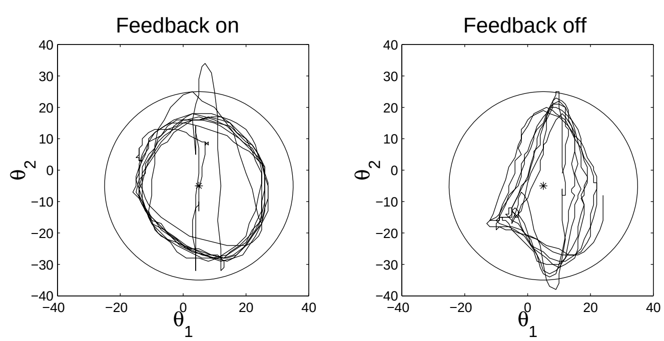

Figure 5-15: $\boxed{\text{REAL}}$ Effect of oscillator feedback. The graphs show transients of the robot arm motion from a fixed position with the oscillator feedback on (left graph), and off (right graph). When the feedback is on, the oscillators use the feedback to converge on the circular motion. 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

$$\begin{aligned}
\tau_1 \dot{x}_1 &= c - x_1 - \beta v_1 - \gamma [x_2]^+ - h[g]^+ && (6.1)\\
\tau_2 \dot{v}_1 &= [x_1]^+ - v_1 && (6.2)\\
\tau_1 \dot{x}_2 &= c - x_2 - \beta v_2 - \gamma [x_1]^+ - h[g]^- && (6.3)\\
\tau_2 \dot{v}_2 &= [x_2]^+ - v_2 && (6.4)\\
[u]^+ &= \max(u, 0) && (6.5)\\
[u]^- &= -\min(u, 0) && (6.6)\\
y_{out} &= [x_1]^+ - [x_2]^+ && (6.7)
\end{aligned}$$

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

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

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

$$v_1(t) \geq -|v_1(0)| \tag{6.9}$$

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

$$v_2(t) \geq -|v_2(0)| \tag{6.10}$$

Integrating (6.1) gives

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

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

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

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

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

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

$$x_2(t) \leq |x_2(0)| + c + \beta|v_2(0)| \tag{6.14}$$

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

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

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

$$v_2(t) \leq (1+\beta)|v_2(0)| + |x_2(0)| + c \tag{6.16}$$

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

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

or

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

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

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

(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:

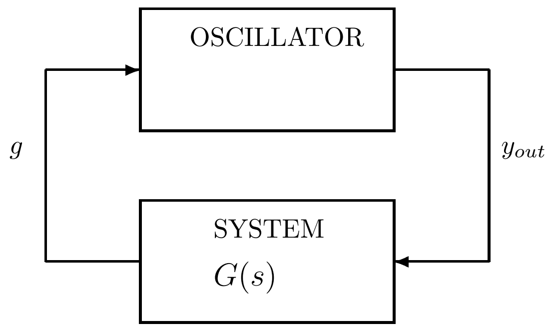

Figure 6-1: Oscillator coupled to a linear system $G(s)$.

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 $\alpha = 1, 2, \ldots, 8$, the system can be written in the form:

$$\dot{\mathbf{x}} = A_\alpha \mathbf{x} + b + H_\alpha g \tag{6.20}$$

$$y_{out} = c_\alpha \mathbf{x} \tag{6.21}$$

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.

x_1

(2)

$$A_2 = \begin{pmatrix} -10 & -20 & 0 & 0 \\ 5 & -5 & 0 & 0 \\ -20 & 0 & -10 & -20 \\ 0 & 0 & 0 & -5 \end{pmatrix}$$

$$c_2 = \begin{pmatrix} 1 & 0 & 0 & 0 \end{pmatrix}$$

(1)

$$A_1 = \begin{pmatrix} -10 & -20 & -20 & 0 \\ 5 & -5 & 0 & 0 \\ -20 & 0 & -10 & -20 \\ 0 & 0 & 5 & -5 \end{pmatrix}$$

$$c_1 = \begin{pmatrix} 1 & 0 & -1 & 0 \end{pmatrix}$$

x_2

(4)

$$A_4 = \begin{pmatrix} -10 & -20 & 0 & 0 \\ 0 & -5 & 0 & 0 \\ 0 & 0 & -10 & -20 \\ 0 & 0 & 0 & -5 \end{pmatrix}$$

$$c_4 = \begin{pmatrix} 0 & 0 & 0 & 0 \end{pmatrix}$$

(3)

$$A_3 = \begin{pmatrix} -10 & -20 & -20 & 0 \\ 0 & -5 & 0 & 0 \\ 0 & 0 & -10 & -20 \\ 0 & 0 & 5 & -5 \end{pmatrix}$$

$$c_3 = \begin{pmatrix} 0 & 0 & -1 & 0 \end{pmatrix}$$

Figure 6-2: Figure showing the different values of A_α and c_α is 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

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

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

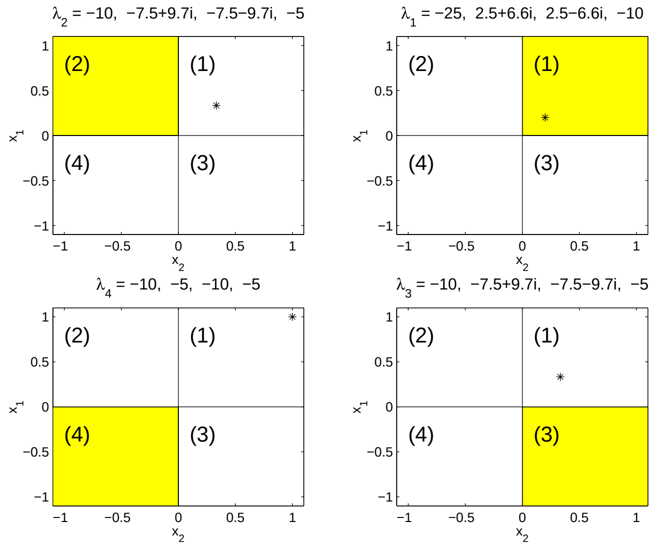

Figure 6-3: $\boxed{\text{SIM}}$ Linear regions, fixed points and eigenvalues. Each sub-plot corresponds to one of the four linear subregions of the oscillator. The titles of each sub-plot give the eigenvalues, the $*$'s mark the fixed points, and the patches indicate the regions for each subsystem. 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)

$$\mathbf{x} = -A_\alpha^{-1} b \tag{6.23}$$

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.

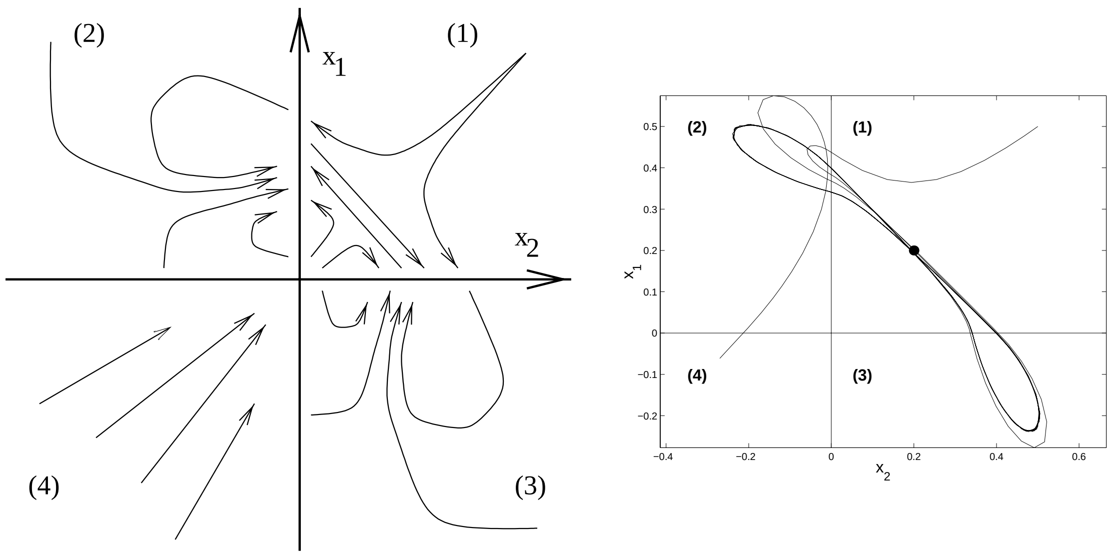

Figure 6-4: $\boxed{\text{SIM}}$ The left hand graph shows the approximate shape of trajectories projected onto the $x_1 - x_2$ plane. The right hand graph shows two trajectories plotted on the same projection, with the unstable fixed point marked with a •. If the trajectory starts in regions (2), (3), 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.

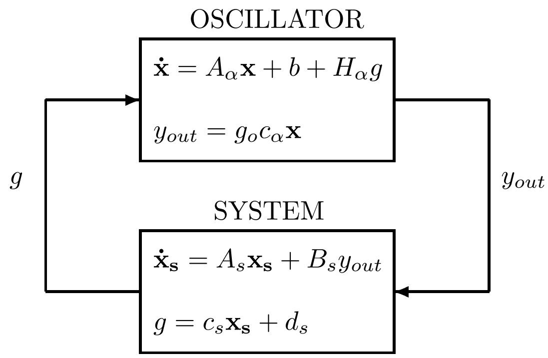

Figure 6-5: Oscillator coupled to a linear system. The individual state equations shown in this figure can be combined to create a single piecewise-linear system as shown in equation (6.28).

(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

$$\dot{\mathbf{x}} = A_\alpha \mathbf{x} + b + H_\alpha g \tag{6.24}$$

$$y_{out} = h_o c_\alpha \mathbf{x} \tag{6.25}$$

$$\dot{\mathbf{x}}_{\mathbf{s}} = A_s \mathbf{x}_{\mathbf{s}} + B_s y_{out} \tag{6.26}$$

$$g = c_s \mathbf{x}_{\mathbf{s}} + d_s \tag{6.27}$$

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

$$\begin{pmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{x}}_\mathbf{s} \end{pmatrix} = \begin{pmatrix} A_\alpha & H_\alpha c_s \\ h_o B_s c_\alpha & A_s \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{x}_\mathbf{s} \end{pmatrix} + \begin{pmatrix} b + H_\alpha d_s \\ 0 \end{pmatrix} \tag{6.28}$$

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

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

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:

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

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

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

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

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

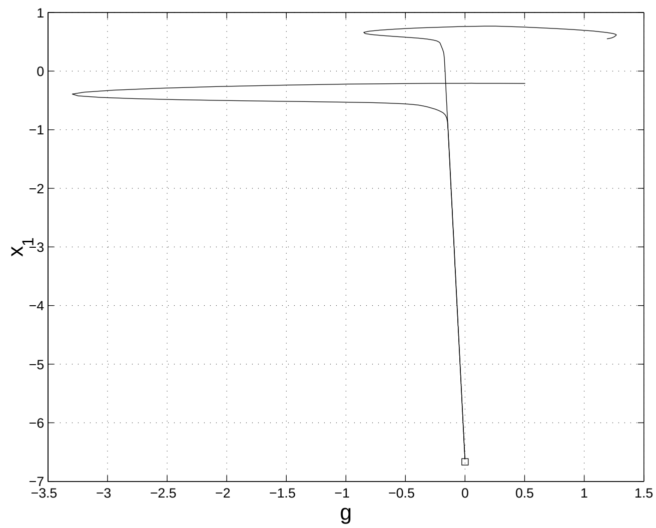

Figure 6-6: $\boxed{\text{SIM}}$ Two transients of the oscillator connected to the system in equation (6.29). The plot shows the trajectories plotted on the $x_1 - g$ plane, and illustrates convergence to the stable fixed point at $(0, 0, 0.33, 0.33, -6.67)'$, marked with a $\square$. The other fixed 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 constants 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

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

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

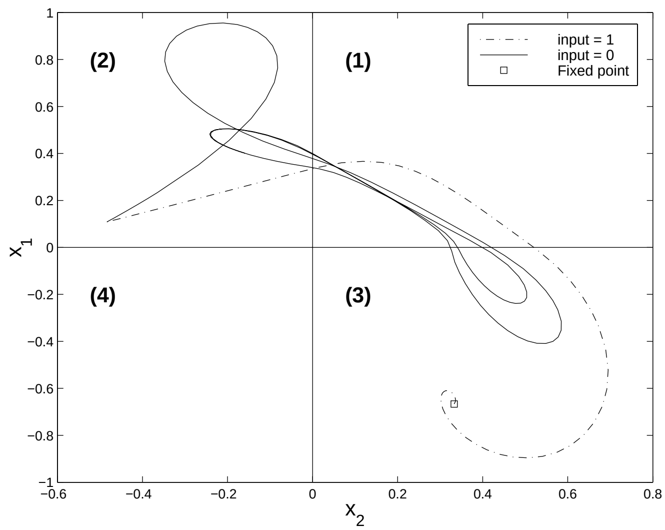

Figure 6-7: $\boxed{\text{SIM}}$ Transients of the system with a constant input. The input changes the system so there is no longer a unstable pole in region (1), but a stable fixed point at $(-0, 67, 0, 0.33, 0.33)$, to which the trajectory (dash-dot) converges. Without the input (solid line) the trajectory 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.

$$N(j\omega, A)G(j\omega) = 1 \tag{6.34}$$

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:

$$|N(j\omega_f, A_f)||G(j\omega_f)| = 1 \tag{6.35}$$
$$\angle N(j\omega_f, A_f) + \angle G(j\omega_f) = 0 \tag{6.36}$$

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

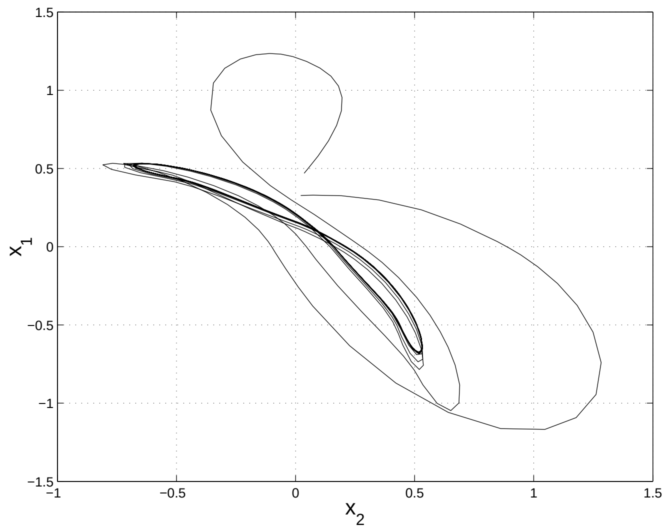

Figure 6-8: $\boxed{\text{SIM}}$ Two transients of the oscillator driving the mass spring system in equation (6.33). The plot shows the 6-dimensional state space projected onto the $x_1 - x_2$ plane. Both transients converge to the limit cycle.

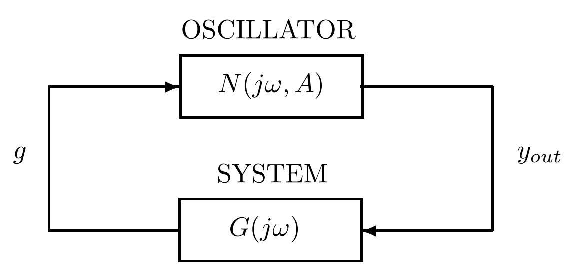

Figure 6-9: Figure showing the oscillator tightly coupled to a driven system. Describing function analysis is used to compute the transfer function $N(j\omega, A)$ for the oscillator as a function of frequency ω and input amplitude A. The transfer function for the system $G(j\omega)$ is easily calculated 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)

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

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

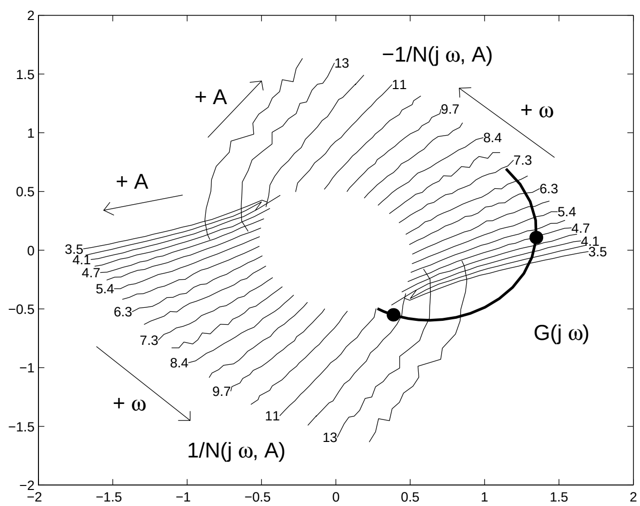

Figure 6-10: $\boxed{\text{SIM}}$ Plot of $1/N(j\omega, A), -1/N(j\omega, A)$ and $G(j\omega$ on the complex plane. $G(j\omega)$ is the frequency response of the non-minimum phase system in equation (6.37), and is indicated with the thick line. The plot predicts limit cycle motions at two different frequencies. 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 $\omega_f = 4.67, A_f = 0.67$.

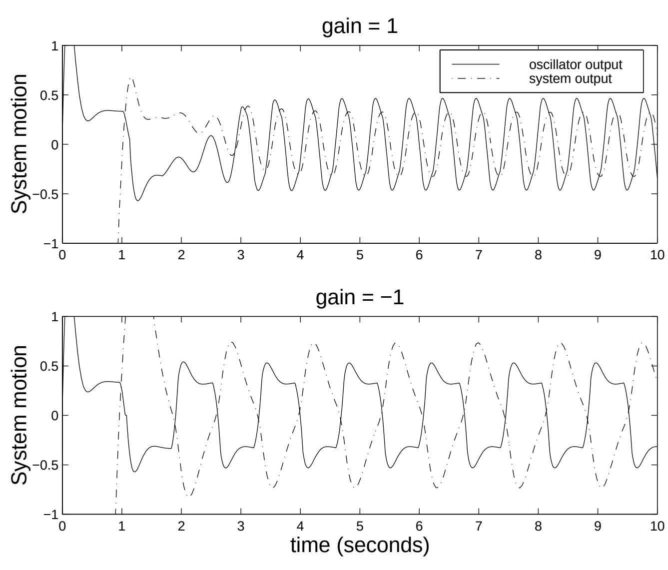

Figure 6-11: $\boxed{\text{SIM}}$ Two transients of the non-minimum phase system, with different values of oscillator output gain h_o. When the gain is positive (top graph), the oscillation is fast, with $\omega_f = 11.4, A_f = 0.32$, compared to the prediction (figure 6-10) of $\omega_f = 11.8, A_f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

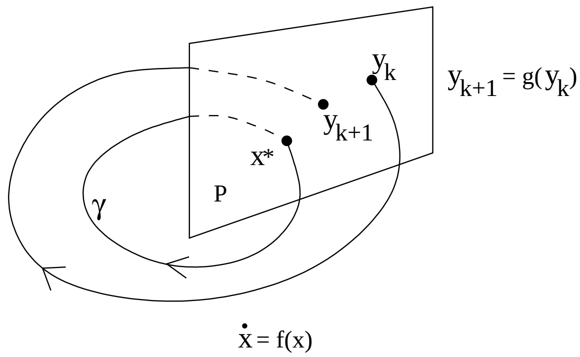

Figure 6-12: Poincaré map. The map turns the n dimensional dynamics $\dot{x} = f(x)$ into the $n-1$ discrete system $y_{k+1} = g(y_k)$ by cutting the dynamics with a $n-1$ dimensional hyper plane P. The limit cycle γ in the full dynamics appears as a fixed point of g, i.e. x^*. The local stability of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

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

and so

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

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

$$S_i \equiv c_i x + d_i = 0 \tag{6.40}$$

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.

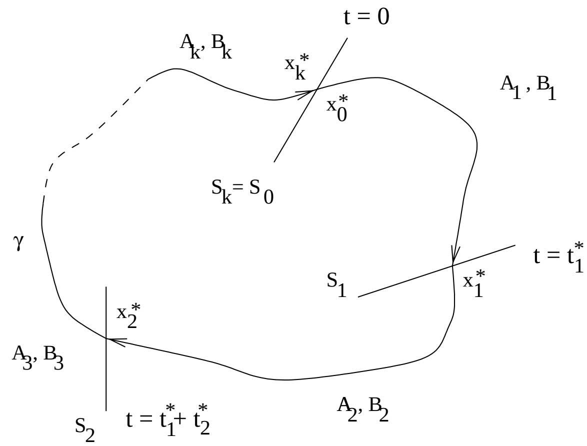

Figure 6-13: The limit cycle trajectory γ starts in state x_0^*, and travels through the first linear system ($\dot{\mathbf{x}} = A_1\mathbf{x} + B_1$) until it crosses the first switching surface S_1 at time t_1^*. It continues into region 2, 3, etc. The limit cycle is complete after k switches, so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

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

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 $c_1 x_i^* + d_1 = 0$,

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

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

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

with times

$$t_0^* = 0.2884, \quad t_1^* = 0.1604, \quad t_2^* = 0.2884, \quad t_3^* = 0.1604 \tag{6.43}$$

These values result in the eigenvalues of W:

$$\begin{pmatrix} 0.1265 \\ -0.1196 + 0.0412j \\ -0.1196 - 0.0412j \\ 0 \end{pmatrix} . 10^{-3} \tag{6.44}$$

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.

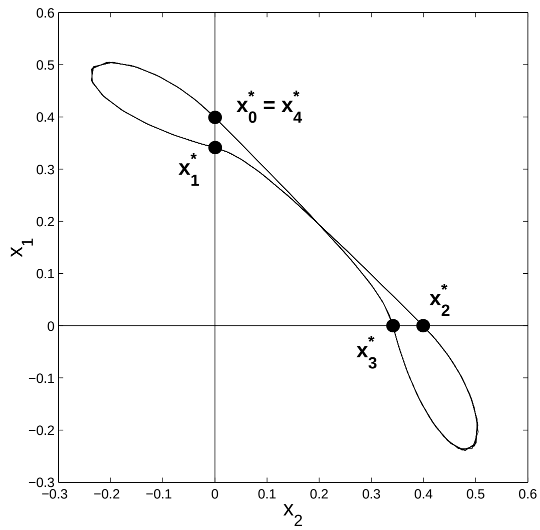

Figure 6-14: $\boxed{\text{SIM}}$ Oscillator limit cycle plotted in the $x_1 - x_2$ plane. The points where the cycle crosses the switching surfaces $x_1 = 0$ (horizontal line), and $x_2 = 0$ (vertical line) are marked with •'s. Using these points, and the times for the trajectory to go between them, 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

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

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

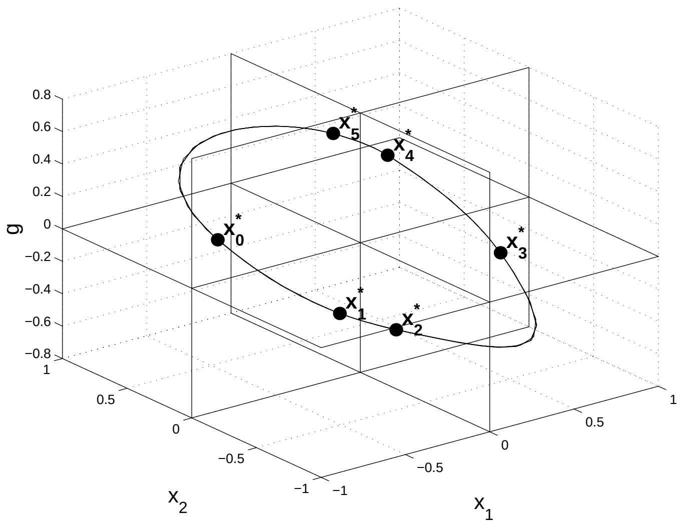

Figure 6-15: $\boxed{\text{SIM}}$ Plot of limit cycle for mass spring system and oscillator. The cycle is plotted in the three switching parameters x_1, x_2, g, with the switching states marked.

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.