Neural Control of Rhythmic Arm Movements

Matthew M. Williamson
MIT AI Lab,
545 Technology Square Rm 937,
Cambridge, MA 02139
Phone: +1 617 253 7471
Fax +1 617 253 0039

E-mail: matt@ai.mit.edu

May 1, 1998

Abstract

In this paper we present an approach to robot arm control based on exploiting the dynamical properties of a simple neural network oscillator circuit coupled to the joints of an arm. The entrainment and input/output properties of the oscillators are used to perform a variety of tasks with the same architecture, without any modeling of the arm or its environment. The approach is implemented on two real robot arms, and has been used to tune into the resonant frequency of pendulums, perform multi-joint coordinated motion by turning cranks, and exploit the dynamics of a 'Slinky' toy to coordinate the motion of two arms. By exploiting the coupling between the physical arm and the neural oscillator, a range of complex behaviors can be achieved with a very simple system.

Keywords: Oscillator, Neural control, Neural network, Robot Manipulator, Rhythmic movement.

1 Introduction

This paper describes the properties of a set of simple neural network oscillators actuating the joints of a real robot arm. By exploiting the coupling between the oscillators and the physical arm of the robot, many qualitatively different tasks have been performed. These include tuning into the resonant frequencies of driven pendulums, turning cranks, and coordinated operation of a 'Slinky' toy. The tasks have all been performed using the same neural architecture, with minimal parameter changes, and without any modeling of the arm or its environment.

The arms used are mounted on the humanoid robot Cog (Brooks and Stein, 1994), shown in Figure 1. The arms are compliant, each joint having a variable stiffness, damping and equilibrium point. Each joint of the robot is actuated by an independent oscillator, consisting of two simulated neurons in mutual inhibition (Matsuoka, 1985). The output of the neural circuit controls the equilibrium point of the joints, and the input to the neural circuit is either the force, or the position at the actuated joint. When the arm moves, the oscillators use the joint level information to adapt their local behavior, giving coordinated motion of the whole arm. The entrainment properties of the oscillators allow them to do this over a range of frequencies, and in a variety of situations, without requiring any parameter changes.

By coupling these oscillators with the dynamics of the arm, the system can exploit the natural dynamics of the combined system: the entrainment properties of the neurons, the inertias and stiffnesses of the arm, the dynamical interaction forces between the limb segments, and the dynamical loads from interactions with objects. The idea of exploiting and working with the arm natural dynamics was originally presented by Bernstein (1935). Later Greene (1982) suggested methods which would result in simple and naturally stable control. By using parts of the arm dynamics which are suited to the task, the system can avoid the accurate calibration and modeling that is required in more recent traditional robotic approaches (An et al., 1988, Craig, 1989).

Humans certainly learn to exploit the dynamics of their limbs for rhythmic tasks (Schneider et al., 1989), during development (Thelen et al., 1992) and also to perform certain tasks like overarm throwing (Bingham et al., 1989). Robotic examples of this idea include "open-loop stable" systems (Schaal and Atkeson, 1993), where the dynamics are exploited giving systems which require little or no active control for stable operation. Indeed McGeer (1990) has built machines that walk stably and are completely passive. Mason and Lynch (1993) presented a more model based approach to dynamic manipulation, and stressed the complexity of model based approaches in dynamic tasks.

The coupling between the oscillators and the arms is also motivated by the work of Kay et al. (1987), Bingham et al. (1991), and Kugler and Turvey (1987) who have suggested that human rhythmic movement is self-organized through the interaction of the various non-linear components of the physical and neural systems. Hatsopoulos (1996) suggests that it is the coupling rather than the individual systems that is important, an idea supported by the work of Taga et al. (1991), whose simulated biped controlled by neural oscillators shows stable walking and stability against perturbations, through the "mutual entrainment" of the dynamics of the legs and the oscillators. Work on invertebrate central pattern generators (Grillner et al., 1991, Getting, 1988), has also pointed out the importance of coupling between the neural and physical systems.

Section 2 of the paper describes the behavior of the oscillators, which can lock onto the frequency of an input over a wide range of input frequencies. The section includes a detailed

analysis of the oscillator input/output behavior in this entrained condition. The following sections show how the behavior of the oscillators can be used to tune into the resonant frequencies of pendulums (section 3), perform coordinated motion without any kinematic modeling to turn cranks (section 4), and a variety of other tasks including exploiting the dynamics of a slinky toy to coordinate the motion of two arms (section 5). Section 6 concludes the paper with a discussion and suggestions for further work.

2 The arms and the oscillators

The two arms used in this experiment are mounted on the humanoid robot Cog (Brooks and Stein, 1994). The arms have been specially designed to interact stably and robustly with unstructured environments. They have six degrees of freedom (dof) arranged in an anthropomorphic manner, each joint being actuated by a series elastic actuator (Pratt and Williamson, 1995, Williamson, 1995). These actuators give low noise force control, shock tolerance, and are stable when interacting with passive environments (Colgate and Hogan, 1989).

At the joints of the arm, a simple proportional-derivative position control loop is used, making the torque at the ith joint

$$u_i = k_i(\theta_{vi} - \theta_i) - b_i\dot{\theta}_i \tag{1}$$

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 dynamical characteristics of the arm can be changed by altering the stiffness and damping of the arm, and the posture of the arm can be changed by altering the equilibrium points (Williamson, 1996). This type of control preserves stability of motion, and since the inner torque control is provided by the series elastic actuators, the overall system is both compliant and shock resistant, making it easy to operate the arm in unstructured environments.

The oscillator model consists of two simulated neurons arranged in mutual inhibition, as shown in figure 2. The model for the neuron is taken from Matsuoka (1985), and describes the firing rate of a real biological neuron with self-inhibition. The firing rate is governed by the following equations.

$$\begin{aligned}
\tau_1\dot{x_1} &= -x_1 - \beta v_1 - \omega\,[x_2]^+ - \Sigma_{j=1}^{j=n} h_j\,[g_j]^+ + c && (2)\\
\tau_2\dot{v_1} &= -v_1 + [x_1]^+ && (3)\\
\tau_1\dot{x_2} &= -x_2 - \beta v_2 - \omega\,[x_1]^+ - \Sigma_{j=1}^{j=n} h_j\,[g_j]^- + c && (4)\\
\tau_2\dot{v_2} &= -v_2 + [x_2]^+ && (5)\\
y_i &= [x_i]^+ = \max(x_i, 0) && (6)\\
y_{out} &= y_1 - y_2 && (7)
\end{aligned}$$

where x_i is the firing rate, v_i is a variable representing the self-inhibition of the neuron (modulated by the adaptation constant β), and the mutual inhibition is controlled by the parameter ω. The output of each neuron y_i is taken as the positive part of x_i, and the output of the whole oscillator as y_{out}. Any number of inputs g_j can be applied to the oscillator, which can either be proprioceptive signals, or signals from other neurons. The input is arranged to excite one neuron and inhibit the other, by applying the positive part ($[g_j]^+$) to one neuron and the negative part to the other. The inputs are scaled by gains h_j.

The tonic excitation c determines the amplitude of the oscillation, with amplitude proportional to c, as shown in figure 3. There is no oscillation if $c = 0$. The two time constants τ_1 and τ_2 determine the speed and shape of the oscillator output. For stable oscillations, τ_1/τ_2 should be in the range 0.1–0.5, for which the endogenous or natural frequency of the oscillator w_n is proportional to $1/\tau_1$, as shown in figure 3. The stability and properties of this oscillator system and more complex networks of neurons are analyzed by Matsuoka (1985, 1987). Figure 4 shows a typical output from the oscillator.

The oscillator is connected to the robot joints by using the output y_{out} to move the equilibrium point θ_v. One neuron flexes the joint and the other extends it about a fixed posture θ_p, making the equilibrium point

$$\theta_v = y_1 - y_2 + \theta_p = y_{out} + \theta_p \tag{8}$$

For the examples in this paper, the inputs to the oscillators are taken to be either the force (u_i) or the position (θ_i) of the joint. These signals in general have an offset (due to gravity loading, or oscillation about a non-zero posture), so when the positive and negative parts are extracted to be applied to the oscillators, a high pass filter is used to remove the DC component.

2.1 Input/output oscillator behavior

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. However, when an oscillatory input is applied, the oscillator can entrain the input, locking onto the input frequency. This is illustrated in Figure 5 which shows the output of the oscillator as the size of the input signal is increased. The oscillator can lock onto input frequencies over a wide range of frequencies and sizes of inputs. This is illustrated in Figure 6 which 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.

The oscillator is a non-linear system, but given its strong entrainment property, conventional non-linear systems representations (e.g., phase plane plots (Vidyasagar, 1978)) are not very expressive of the system behavior as the input frequency and size is varied. Better intuition is gained by evaluating the behavior of the system at different frequencies, and presenting the analysis in a linear system format (e.g., a bode plot). This approach is known as describing function analysis (Slotine and Li, 1991). Throughout the rest of the paper, the non-linear behavior of the system is plotted and displayed in a linear-like manner, as a convenient way of displaying the data, and a way of garnering intuition about the system behavior as a whole. The strong entrainment property of the system bolsters this intuition.

Figure 7 is such a plot, where the magnitude of the output and the phase between the input and output of an entrained oscillator is displayed as a function of frequency. Since the oscillator output is entrained, both the input and output oscillate at the same frequency, and the magnitude and phase can be computed using a single frequency Fourier transform. The plot shows that the oscillator output magnitude decreases at frequencies away from the oscillator natural frequency. The horizontal line in the graph corresponds to the oscillator amplitude without any input, which is slightly smaller than the driven response at w_n. The lower graph shows the phase relationship, with phases of 180° at low frequencies moving to 90° at high frequencies.

The same shape of plot is obtained independent of the size of the input signal, assuming that it is large enough to cause frequency-locking. This means that the sensitivity of the input gain parameter h_j is low. When a different endogenous frequency is used (different τ_1), the same shape is obtained, but over a different frequency range. If the tonic excitation c is varied, the same phase plot is produced, but the magnitude plot, and the size of the input required to cause frequency locking are scaled.

The ability of the oscillator to entrain the frequency of the input over a wide range of frequencies, using the same set of parameters, is exploited for control of the robot. The input/output phase properties are also exploited, since they provide interesting and appropriate behavior when coupled to the physical arm system. In the following three sections, the behavior of the coupled oscillator-arm system is analyzed when the robot is moving freely in space, when its motion is constrained, and when there are perturbing forces on the system.

3 Interaction with a free mass

Human arms (and Cogs arms) can be thought of as masses connected by springs, whose frequency response makes the energy and control required to move the arm vary with frequency. At the resonant frequency, the control need only inject a small amount of energy to maintain the vibration of the mass of the arm segment on the spring of the muscles and tendons. The frequency response of the system thus determines speeds and frequencies that efficiently move the arm. Finding and using that frequency is a desirable property of a robot arm controller.

It appears that humans exploit the natural frequencies of their arms, swinging pendulums at "comfortable" frequencies equal to the natural frequency (Hatsopoulos and Warren, 1996), a condition that corresponds to a minimum in metabolic cost (Herr, 1993).

When a mass or inertia I is driven by the robot, the joint torque relation given in equation 1 makes the equation of motion of the whole system

$$I\ddot{\theta} + b\dot{\theta} + k\theta = k\theta_v \tag{9}$$

which is the standard equation for a mass I vibrating on a spring with stiffness k, damping b and a forcing function $k\theta_v$. The system has a resonant frequency $w_{sys} = \sqrt{k/I}$, and a damping ratio $\zeta_{sys} = b/2\sqrt{kI}$.

When the oscillator is connected to this system, driving θ_v with the oscillator output (equation 8), and connecting either the position θ or the force on the mass τ to the oscillator input, the two dynamical systems are coupled, and the final frequency and amplitude is determined by how they interact. The system behavior for a typical set of parameters is a driven oscillation of the mass, as illustrated in Figure 8. The plot shows the result of starting the mass in 3 different states, which all converge to the limit cycle behavior within one cycle.

Since the two systems are tightly coupled, the oscillator with input θ and output θ_v and the mass-spring system with input θ_v and output θ, the phase difference between θ_v and θ must be the same for both individual systems. Looking at the phase behavior of the two systems independently cannot give information about the transient behavior of the coupled system, but can bolster intuition about the final frequency of the mass motion. The frequency response of the two systems is illustrated in figure 9, which shows the behavior for one choice of system and oscillator resonant frequencies, looking at position feedback for the oscillator. The top graph shows the magnitude of the mass motion θ, with a clear resonance peak at w_{sys}. The lower graph shows the phase difference between θ_v and θ for both the oscillator and mass system. The mass will oscillate at a frequency where the two phase differences are identical, where the two lines intersect. This frequency is in this case is almost exactly w_n, or 7 rad/s.

When the resonant properties of the mass system are altered, the final frequency depends on the exact shape of the two phase plots. This is illustrated in Figure 10. The top plot gives examples of four different systems, and where the final frequency is the intersection of the mass and oscillator phase plots. The lower graph shows the results of simulations of the actuator driving the mass, plotting the final frequency of the mass against the system resonant frequency. The two plots agree; for systems with $w_{sys} < w_n$, the system moves at approximately w_n, and as w_{sys} is increased, the final frequency tends towards w_{sys}. In the lower plot the result of moving the mass under torque feedback is also shown. Although the preceding argument was developed using position feedback, a similar argument can be developed for the oscillator under torque feedback. The different phase profiles under torque feedback combine to give the low frequency behavior observed in figure 10.

This behavior was verified experimentally using a single joint of the robot to drive pendulums of different lengths with different joint stiffnesses[1]. This more complex scenario was chosen to further show that the oscillators are tuning into a property of the overall system. Figure 11 shows the results, indicating that over a range of system frequencies (5 to 9 rad/s with $w_n = 7$ rad/s) the behavior is to drive the system at the resonant frequency. Above 9 rad/s the oscillator drives the motion at a lower frequency than resonance. This result is most likely explained by the motion exceeding the actuator bandwidth limits.

These results show that an oscillator with a constant set of parameters has the ability to tune automatically into the natural frequency of the system that it drives. It thus finds the most efficient driving frequency for the system. The behavior is robust, driving a variety of different systems with the same set of parameters. The tight coupling of the oscillator and the actuated system results in quick entrainment, and stability to perturbations of the system. In addition, since there is no modeling or system identification, the system is very computationally efficient. These results are similar to those of Hatsopoulos et al. (1992) and Hatsopoulos (1996), who showed similar entrainment using a simulated pendulum, using a variety of different oscillator types.

4 Interaction with a position constraint

Whenever a robot contacts an object in the world, there are position constraints imposed on the motion. The constraint can restrict the motion in just one direction (e.g., touching a wall) or in a more complex manner (e.g., opening a door, turning a crank). Traditional robotic solutions to these problems have been to model the constraint, and determine ways to control the force exerted by the robot (e.g., controlling position and force separately (Raibert and Craig, 1981), or together (Hogan, 1985a)). In recent work by Niemeyer and Slotine (1997) and Deacon (1997), the constraint is used to reactively define the movement, without using explicit modeling. The approach taken in this paper is to use the constraint as much as possible, exploiting it to perform multi-joint coordinated motion without any kinematic modeling. The coordination is achieved through the coupling of the oscillators, the arm and the position constraint, and works because of the entrainment and phase properties of the oscillators. The local behavior of the oscillators interacts through the physical structure of the arm to give global coordinated motion.

[1]For this case, the natural frequency depends on the stiffness k, length l and inertia I, $w_{sys} = \sqrt{g/l + k/I}$, where g is the gravitational constant.

The action of a single joint of the robot under a position constraint can be determined by measuring the response of the oscillator when the angle of the joint is forced to move in a sine wave:

$$\theta = A\sin(\omega t) \tag{10}$$

The oscillator responds either to the angle directly or to the force generated by the motion, depending on what type of feedback is used. In a typical case, the oscillators entrain the frequency of the applied motion, and move the equilibrium point at the same speed. Since the torque applied is given by

$$u = k(\theta_v - \theta) - b\dot{\theta} \tag{11}$$

the effect of the oscillator will be to either reduce or increase the torque, depending on the phase angle between θ and θ_v. This torque must be overcome to maintain the output motion in equation 10.

The relationship between u and θ defines the output impedance Z of the system, which in the linearized case is a complex number with a magnitude and phase. For example, the impedance of the robot joint when the oscillator is not connected ($\theta_v = 0$), is

$$Z(\omega) = u(\omega)/\theta(\omega) = -k - j\omega b \tag{12}$$

where $j^2 = -1$. For a general non-linear system, $Z(\omega)$ cannot be calculated in this way, because u and θ will not scale linearly, or necessarily have the same frequency content. Since the oscillator system has the entrainment property, both input and output are at the same frequency, and a measurement of output impedance can be made. The impedance is only for that particular input magnitude (A in equation 10), as different input magnitudes will not have the same effect on the system, e.g., very small inputs will not cause entrainment. This limited measurement of impedance still provides a good way to interpret the oscillator behavior.

Figure 12 shows the impedance of the oscillator, measured by calculating the size and phase of τ using a single frequency Fourier transform at the imposed motion frequency ω. The motion imposed was the same size as the oscillator amplitude with no inputs.

Under position feedback, at low frequencies the oscillator opposes the motion, creating large forces at the joint (high impedance at $-180°$ phase), while at high frequencies the equilibrium point leads the motion by $90°$, so reducing the impedance. Under torque feedback, at low frequencies the equilibrium point tracks the input resulting in low impedance, while towards w_n the equilibrium point again leads by $90°$. At high frequencies the oscillator cannot entrain and the impedance reverts to that of the spring and damper alone.

If different amplitudes of motion are used, the oscillator behavior is more complicated, since there is an amplitude difference between the imposed motion and the oscillator natural amplitude. The general shape of the response is similar to that described above, however for small input amplitudes, the force generated tends to push the input towards larger amplitudes, and for larger inputs the opposite is the case, the oscillator acting to oppose the large motion. For very small inputs, there is no entrainment, and the impedance measure loses its meaning.

The local impedance property of the oscillator system is useful as the joint level behavior in multi-joint position constrained tasks. When a motion is imposed, the oscillators independently track the motion and reduce the forces at the joints. This reduces the force required at the endpoint to generate the motion. They can do this over a range of frequencies, without requiring parameter

changes. The oscillators do require the motion to be imposed, since there is no guarantee that if the oscillators drive the system, the overall behavior will be correct.

On the robot these behaviors have been used for a number of position constrained tasks, including pumping a bicycle pump and turning a crank, as shown in figure 13. For the results shown in Figure 14, the shoulder was driven at a constant frequency, and position feedback at the elbow allowed the oscillator there to entrain the motion. When the feedback is on, the elbow motion is coordinated with the shoulder and the crank is turned.

As mentioned above, the oscillators are better suited for following an imposed motion than creating their own. However, in the crank turning task, they are used to create the motion. To make the task easier, some inertia was added to the crank. Without the extra inertia the arms were still coordinated, although the motion was just back and forth, not turning the crank in full circles. In addition it was found that during the period of exploration before the oscillators settle into a pattern of coordination, internal forces caused by unsynchronized oscillators were disrupting the system. This was particularly true under force feedback, where the arm could get stuck but still be exerting and tuning into the internal forces. Under position feedback the situation is slightly better since when the crank gets stuck the input falls to zero, causing the oscillator to return to its natural frequency, so moving the arm out of the stuck condition. A torque limit implemented on the shoulder further reduced the effect of the internal forces, which helped the entrainment of the system onto the desired crank turning behavior.

The performance of the crank turning is robust to changes in frequency and to perturbations, returning to the stable crank motion. It is not robust to large changes in crank length or location, mainly because the oscillators operate around a fixed posture, with somewhat fixed amplitudes[2]. In addition, the system cannot handle very large cranks, because the sinusoidal-like outputs of the oscillators are not appropriate as joint commands for large motions. The torque limit mentioned above does allow the oscillators to produce a non-sinusoidal output, although in a rather limited way. Traditional robot approaches to this kind of manipulation (hybrid force/position control (Raibert and Craig, 1981), or impedance control (Hogan, 1985a)), use explicit kinematic knowledge of the arm and the crank location to solve this problem. Like the oscillator system, these methods are sensitive to errors in crank location. The advantage of using the oscillator method is that once the posture and approximate amplitudes have been specified, the other details of the system emerge from the interaction of the oscillators, the arm and the crank, without any further calculation. Another advantage is that no further calculation is needed to handle changes in crank frequency, or to deal with unexpected disturbances.

5 Interaction with external forces

As well as driving free motions, and responding to position constraints, there is another class of tasks for the oscillators, that of driving the limb under the action of perturbation forces. These may come from the internal dynamics of the arm (coriolis, centripetal, inertial forces) or from interaction with objects in the world (e.g., the rhythmic force from a bouncing ball). As in the previous two cases, the entrainment and phase properties of the oscillators lend themselves to this

[2]The actual amplitude of the oscillator output is partly determined by the constant c (equation 2), and partly by the size of the feedback signal.

application.

Figure 15 shows the response of the oscillator to a small disturbing force applied to the oscillator output. Under position feedback, the force is rejected and ignored by the oscillator but under force feedback the frequency of the force is entrained. This occurs at frequencies below w_n, and results in the equilibrium point of the system lagging the force perturbation by a small angle which increases with the frequency of the disturbance. This behavior is somewhat esoteric but has been found to be useful in the following examples.

5.1 Slinky

The force behavior can be used to coordinate two arms when operating a slinky toy, as shown in Figure 1. As the slinky is moved from hand to hand, the weight on each hand provides a force perturbation that can be used to entrain the oscillators. The oscillators, which can be initially unsynchronized, and have different natural frequencies, are coordinated through the physical motion of the slinky. Figure 16 shows the drive to the two hands with and without the force feedback, showing that the motion of the slinky is enough to very quickly lock the phase and the frequency of the two oscillators. If one of the hands is moved faster, the other speeds up to match it. If both hands are stopped and released, within one cycle the anti-phase motion is established. Interestingly, as well as the stable anti-phase motion, the system exhibits a less stable solution of moving both hands in phase. It is difficult to get this solution, since only a small asymmetry in the weight on the two hands results in the anti-phase motion. The slinky behavior has been achieved using a variety of different joints, the only proviso being that the oscillators require a minimum input for entrainment, so the mass of the slinky has to act fairly directly onto the actuating joints.

This behavior is interesting as it opens up the possibility of performing a number of rhythmic tasks, such as ball bouncing, yo-yos, and perhaps even throwing and catching using this simple architecture as a base. It also shows clearly how the properties of the environment (in this case the mass of the slinky) can be used to give simple control and robust behavior.

5.2 Flailing

As well as perturbations from external objects, the arm when it moves is subject to perturbations from the arm dynamics (Hollerbach and Flash, 1982). By relying on these internal forces, the oscillators can be used to coordinate the arms during motion. For example, if the shoulder is driven at a constant frequency, the elbow oscillator can sense the torques induced at the elbow due to the shoulder motion, and entrain into the frequency of the arm motion. An example of this entrainment is shown in Figure 17. The frequency of the shoulder motion can be varied, and the elbow will remain entrained without any changes in parameters. The motion is also stable, and will return to the motion when perturbed.

As in the last example, there is some sensitivity in this application to the posture of the arm, as it has to provide a minimum force to allow the oscillators to entrain. Some postures provide more stable flails than others, which is not surprising given the complex dynamics of the arm itself. Buchanan and Kelso (1993) showed some results on the stability of different flailing patterns with changing posture, although they found a change in pattern with frequency which one might not expect in this particular system.

5.3 Force impulses

As well as entraining the effect of rhythmic force applied from the environment, the oscillator system can also entrain forces applied at the joints. This is useful as it gives a way to internally influence the behavior of the oscillators while still retaining and exploiting their entrainment properties. The extra force u_p can be applied by changing the joint force control to be

$$u_i = k_i(\theta_{vi} - \theta_i) - b_i\dot{\theta}_i + u_p \tag{13}$$

If the new force is fed back to the oscillator, the system does not entrain to the frequency of u_p. Entrainment is achieved if the effect of the perturbing force is hidden, by feeding back $k_i(\theta_{vi} - \theta_i) - b_i\dot{\theta}_i$.

The extra force can be any rhythmic signal; the effect of a series of impulses is shown in Figure 18 for the wrist joint of the robot. The torque feedback synchronizes the motion of the joint with the torque impulse. This property shows not only another way to control the joints, but also demonstrates the oscillators ability to respond to different types of rhythmic signals.

6 Discussion

This paper has presented in detail the behavior of a simple neural oscillator coupled to a real robot arm. The oscillator's entrainment property over a range of frequencies and its output phase relationship have been reported to be useful in a variety of situations. Under different conditions of the arm and its environment, the properties of the oscillators can be exploited to give simple as well as stable control. The system can give coordinated motion without any global synchronization or control, the motion emerging from the local interaction of the oscillators and the physical structure of the system.

One limitation of the present scheme is that the oscillator output waveform has a constant shape that is scaled only by changes in frequency. If the oscillators could output more complex waveforms, a wider range of tasks should be possible. For example, to turn a large crank requires commands to the arm joints which are not simple sinusoids. A more complex waveform can be achieved by changing the type of oscillator, to perhaps that suggested by Elias and Grossberg (1975). Oscillators which use different basis functions to generate their output signals might also be appropriate. Whatever the particular oscillator chosen, an important issue is maintaining the coupling between the arm and the oscillator system. The coupling is crucial to the work presented in this paper, as well as to the work of Hatsopoulos (1996). Work on gait transitions in legged locomotion also indicates the importance of coupling: researchers have achieved similar gait transitions independent of the particular choice of oscillator (Collins and Richmond, 1994, Pribe et al., 1997).

An alternative to increasing the complexity of the oscillators is to increase the complexity of the oscillator network, i.e., allow connections between the oscillators at different joints. This would then allow coordination based on explicit connections rather than implicit coupling with the world. In the slinky example, coupling through the physical slinky was used to coordinate the two arms, but the force from the slinky needed to act directly on the driving joint. Other joints that could also generate the motion, but due to the arm configuration might not have a

strong physical coupling with the slinky force, could then be coordinated using extra software connections.

A related idea is to use a single oscillator to control a number of joints. Complete arm motions could then be achieved, and would respond to the dynamics of the situation in a similar way to the joint-level oscillators. Operating at arm-level as opposed to joint-level should allow more complex tasks to be performed. As well as having muscles that span one joint, humans have muscles which span two or more joints (Kahle et al., 1992). This redundant actuation architecture gives a rich repertoire of arm behaviors (Hogan, 1985b). By applying oscillators both at the joint and spanning various joints, the human system could be mimicked. Coordination of this redundant control system could then be achieved using the entrainment properties of the oscillators.

The stable motions exhibited by the oscillator system particularly during flailing (section 5.2), suggest that the system could be used for general oscillatory arm motion. This could be achieved by using these motions as primitives which are combined to create more complex motions. The primitives could be implemented as single oscillators driving a number of joints, or with oscillators at the joints coordinated through the arm structure. The details of the combination method (superposition, winner-take-all etc.) could be adapted to provide coordinated behavior, for example giving accurate visual-arm coordination. There is evidence for this type of hierarchical organization of movement in the spinal cord of frogs and rats, the movement of their legs being achieved by combining a small number of primitives (Mussa-Ivaldi et al., 1994, Bizzi et al., 1991). Using a set of stable motion primitives to span the possible motions for the arm should confer stability to the combined motions. The reduced number of degrees of freedom of the system (since full arm motions rather than single joint motions are combined) should also facilitate learning. In the work of Marjanović et al. (1996), the use of static primitives was found to reduce the complexity of learning in a visual-motor coordination task.

Alternatively, the oscillator structure could be tuned directly to achieve a desired motion. The amplitudes of the joint motions, the type of feedback, the time constants, and maybe even the connections between different oscillators could be altered by a learning algorithm, using information about the system performance. The use of the oscillator structure would then confer on the final motion all the advantageous properties of the oscillator system: robustness to changes in frequency, low sensitivity to parameters values, and stability against perturbations. One method to achieve this may be to use feedback error learning (Kawato, 1990). The combination of a feedforward oscillator controller and a feedback system could be used to generate the arm motion. The output of the feedback system (the feedback error) could then be used to alter parameters in the oscillator circuit, which would gradually assume responsibility for control of the entire motion.

The ability of the oscillators to respond to imposed motions (for example during the crank turning task) opens up the possibility of putting a human in the teaching loop, guiding and modifying the movement while it occurs. This is a natural way to teach complex motor tasks—it is extensively used by humans when teaching tennis, golf, etc. The arm is also subjected to external forces when it touches objects during motion, for example while painting or wiping a surface. If it is possible to capture the effect of these external influences on the arm, then the underlying control can be altered to reproduce the correct motion. While clearly teaching the robot in this way is appealing, questions remains as to how the guiding influence can be measured, represented, remembered, and maybe even generalized to other situations.

It is also appealing to suggest that this approach can be extended to discrete motions, perhaps

by using the approach of Schöner (1990), although evidence from human motion suggests that rhythmic and discrete motions may be planned in separate systems (Adamovich et al., 1994).

To conclude, using the approach presented in this paper may allow general oscillatory motion, and more complex rhythmic tasks to be achieved by exploiting the coupled dynamics of an oscillator system and the arm dynamics. The success of the simple oscillator system demonstrated in this paper suggests that more complex systems based on the same ideas will exhibit the same properties: low sensitivity to parameter values, robustness to changes in frequency, and stability against perturbations.

References

Adamovich, S. V., Levin, M. F., and Feldman, A. G. (1994). Merging different motor patterns: Coordination between rhythmical and discrete single-joint movements. *Experimental Brain Research*, **99**, 325–337.

An, C. H., Atkeson, C. G., and Hollerbach, J. M. (1988). *Model-based control of a robot manipulator*. Cambridge, MA, MIT Press.

Bernstein, N. A. (1935). The problem of the interrelation of co-ordination and localization. *Archives of the Biological Sciences*, **38**.

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.

Bizzi, E., Mussa-Ivaldi, F. A., and Giszter, S. F. (1991). Computations underlying the execution of movement: A biological perspective. *Science*, **253**, 287–291.

Brooks, R. A. and Stein, L. A. (1994). Building brains for bodies. *Autonomous Robots*, **1** (**1**), 7–25.

Buchanan, J. J. and Kelso, J. A. S. (1993). Order parameters for the neural organization of single, multijoint limb movement patterns. *Experimental Brain Research*, **94**, 131–142.

Colgate, 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.

Craig, J. J. (1989). *Introduction to Robotics: Mechanics and Control*. Reading, Massachusetts, Addison-Wesley, second edition.

Deacon, G. E. (April 1997). *Accomplishing Task-Invariant Assembly Strategies by Means of an Inherently Accommodating Robot Arm*. PhD thesis, Department of Artificial Intelligence, University of Edinburgh, Scotland.

Elias, S. and Grossberg, S. (1975). Pattern formation, contrast control and oscillations in the short term memory of shunting on-center off-surround networks. *Biological Cybernetics*, **20**, 69–98.

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*, pp. 101–127. New York, Wiley.

Greene, P. H. (1982). Why is it easy to control your arms? *Journal of Motor Behavior*, **14** (**4**), 260–286.

Grillner, S., Wallén, P., and Brodin, L. (1991). Neuronal network generating locomotor behavior in lamprey: Circuitry, transmitters, membrane properties and simulation. *Annual Review of Neuroscience*, **14**, 169–199.

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. and Flash, T. (1982). Dynamic interactions between limb segments during planar arm movement. *Biological Cybernetics*, **44**, 67–77.

Kahle, W., Leonhardt, H., and Platzer, W. (1992). *Locomotor System*, Vol. 1 of *Color Atlas and Textbook of Human Anatomy*. New York, Thieme Medical Publishers Inc.

Kawato, M. (1990). Feedback-error-learning neural network for supervised motor learning. In R. Eckmiller (Ed.), *Advanced Neural Computers*, pp. 365–372. Elsevier Science Publishers B.V.

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.

Kugler, P. N. and Turvey, M. T. (1987). *Information, natural law and the self-assembly of rhythmic movement*. Hillsdale, NJ, Erlbaum.

Marjanović, M. J., Scassellati, B., and Williamson, M. M. (1996). Self-taught visually-guided pointing for a humanoid robot. In *From Animals to Animats: Proc 1996 Society of Adaptive Behaviour*. Society of Adaptive Behavior.

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.

Matsuoka, K. (1985). Sustained oscillations generated by mutually inhibiting neurons with adaption. *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 walking with knees. In *Proc 1990 IEEE Intl Conf on Robotics and Automation*.

Mussa-Ivaldi, F. A., Giszter, S. F., and Bizzi, E. (1994). Linear combinations of primitives in vertebrate motor control. *Proceedings of the National Academy of Sciences*, **91**, 7534–7538.

Niemeyer, G. and Slotine, J.-J. E. (1997). A simple strategy for opening an unknown door. In *Proc. of 1997 IEEE Int. Conf. on Robotics and Automation*, pp. 1448–1453.

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.

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. and Craig, J. J. (1981). Hybrid position/force control of manipulators. *Journal of Dynamic Systems, Measurement, and Control*, **103**, 126–133.

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.

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. and Li, W. (1991). *Applied nonlinear control*. Englewood Cliffs, N.J., Prentice Hall.

Taga, G., Yamaguchi, Y., and Shimizu, H. (1991). Self-organized control of bipedal locomotion by neural oscillators in unpredictable environment. *Biological Cybernetics*, **65**, 147–159.

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.

Vidyasagar, M. (1978). *Nonlinear systems analysis*. Englewood Cliffs, N.J., Prentice-Hall.

Williamson, 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. (1995). *Series Elastic Actuators*. Master's thesis, Massachusetts Institute of Technology Artificial Intelligence Lab, Cambridge, Massachusetts.

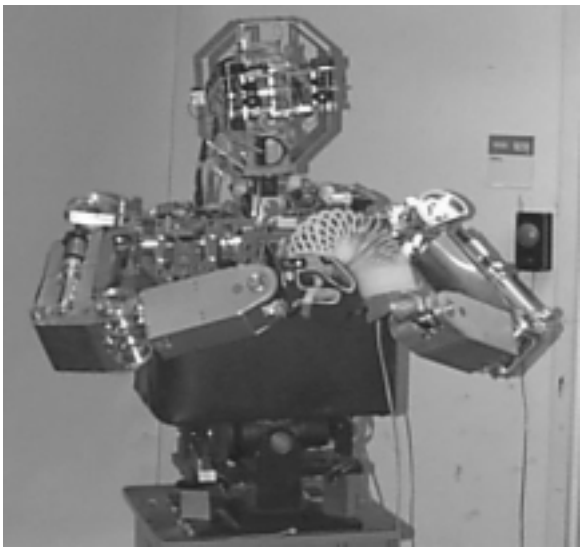

Figure 1: Cog playing with a 'Slinky' toy. This picture shows the humanoid form of the robot, with the two 6 degree of freedom arms used in this paper. The robot is using its elbow joints to move the slinky, exploiting the physical structure of the slinky to coordinate the two arms (see section 5).

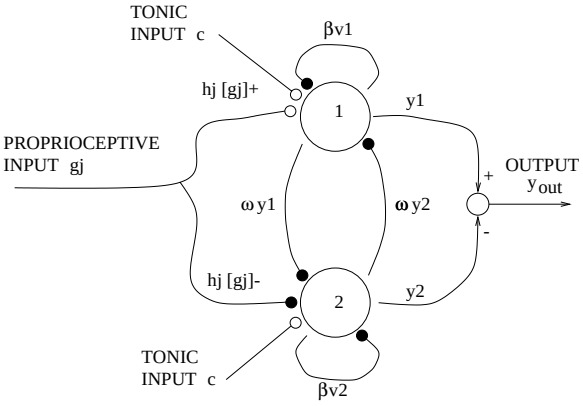

Figure 2: Schematic of the oscillator. Black circles correspond to inhibitory connections, open to excitatory. The βv_i connections correspond to self-inhibition, and the ωy_i connections give the mutual inhibition. The positive and negative parts of the input g_j are weighted by the gain h_j before being applied to the neurons. The two outputs y_i are combined to give the oscillator output y_{out}. See text for more details.

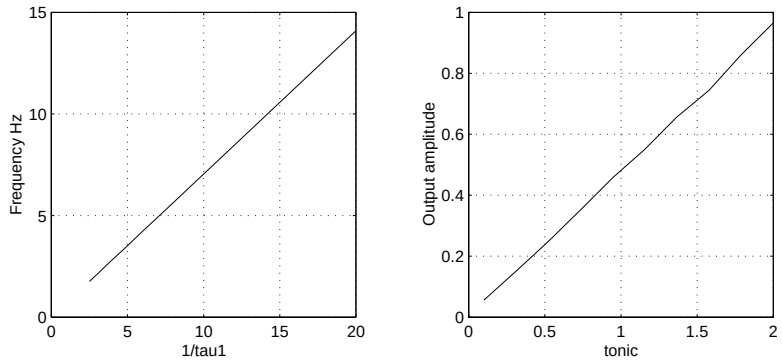

Figure 3: Oscillator behavior under changing tonic and time constant. The left hand figure shows the output frequency of the oscillator plotted against $1/\tau_1$, and the right hand graph shows the amplitude of the oscillator output plotted against the tonic excitation. For this example $\tau_1/\tau_2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

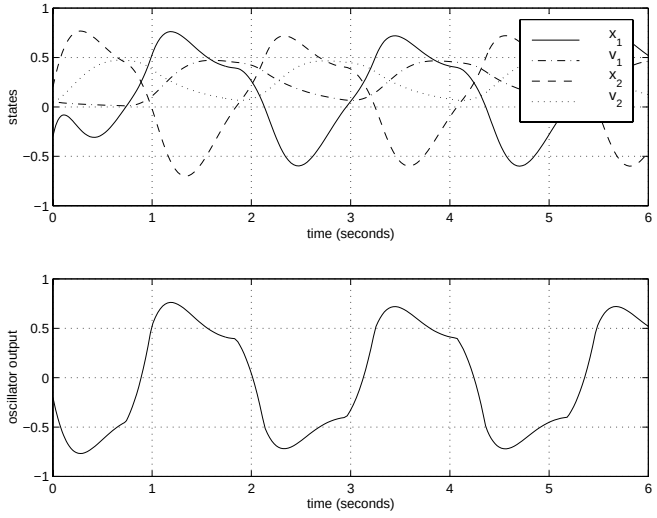

Figure 4: A sample output from the oscillator. The top graph shows the variation with time of the states of the oscillator x_1, x_2, v_1, v_2 during normal operation. The bottom graph shows the output of the oscillator $y_{out} = [x_1]^+ - [x_2]^+$. For this example $\tau_1 = 0.25, \tau_2 = 0.5, c = 1.5, \beta = 2.5, \omega = 2.5$.

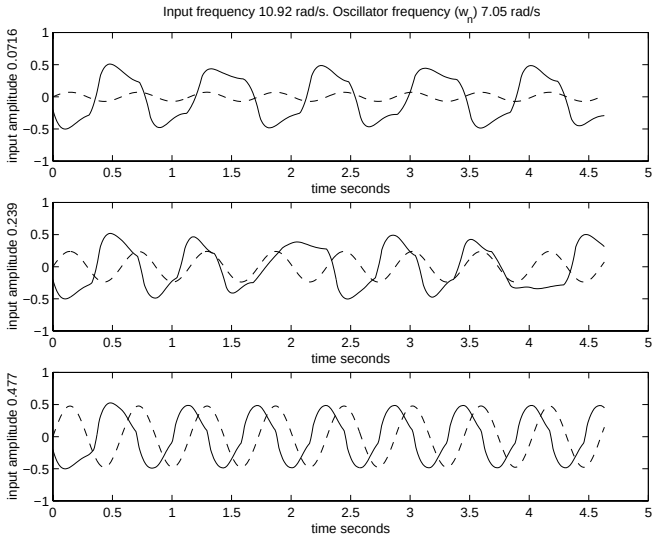

Figure 5: Figure showing effect of increasing input signal. For small input (top graph), the oscillator is not entrained, and oscillates at its endogenous frequency w_n. In the middle graph, the input is larger, and the oscillator is almost entrained, but slips every couple of cycles. The lower graph shows the oscillator locked onto the frequency of the input. For this example $c = 1.0, \tau_1 = 0.1, \tau_2 = 0.2$.

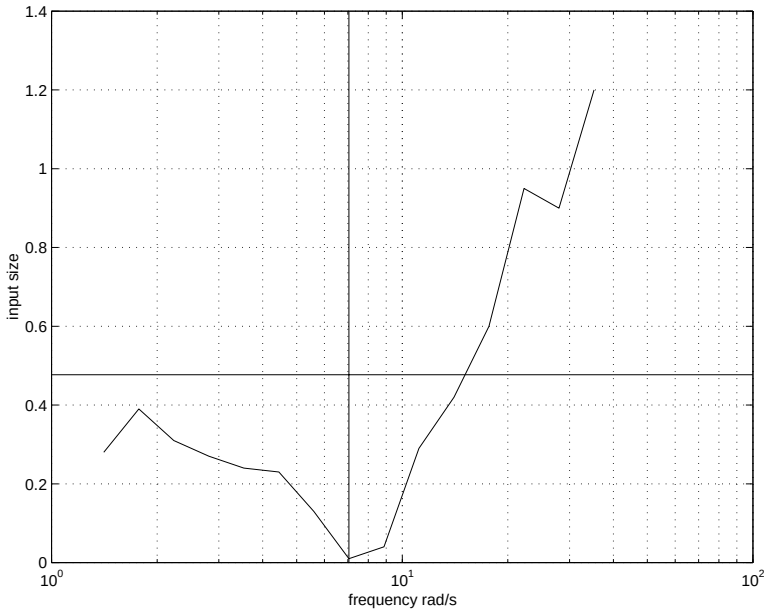

Figure 6: Minimum input for entrainment. The figure shows the minimum input signal required for entrainment of the oscillator. The oscillator natural amplitude and frequency are given by the horizontal and vertical lines respectively. For this example $c = 1.0, \tau_1 = 0.1, \tau_2 = 0.2$.

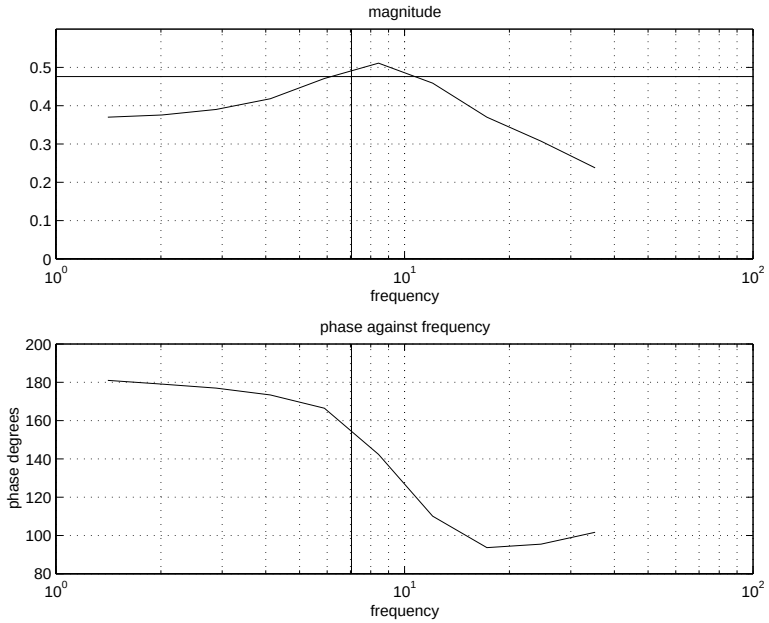

Figure 7: Oscillator frequency response. The top graph shows the magnitude of the oscillator output plotted against frequency. The horizontal line is the magnitude when no input is applied to the oscillator. The endogenous frequency of the oscillator w_n is indicated by the vertical line on both plots. The oscillator can entrain the input over a wide range of frequencies, in this case 1.5 to 35 rad/s with $w_n = 7$ rad/s. The lower graph shows the phase relationship between the output and input, with the output leading the input by an angle in the range 180° to 100°. For this example $c = 1.0, \tau_1 = 0.1, \tau_2 = 0.2$.

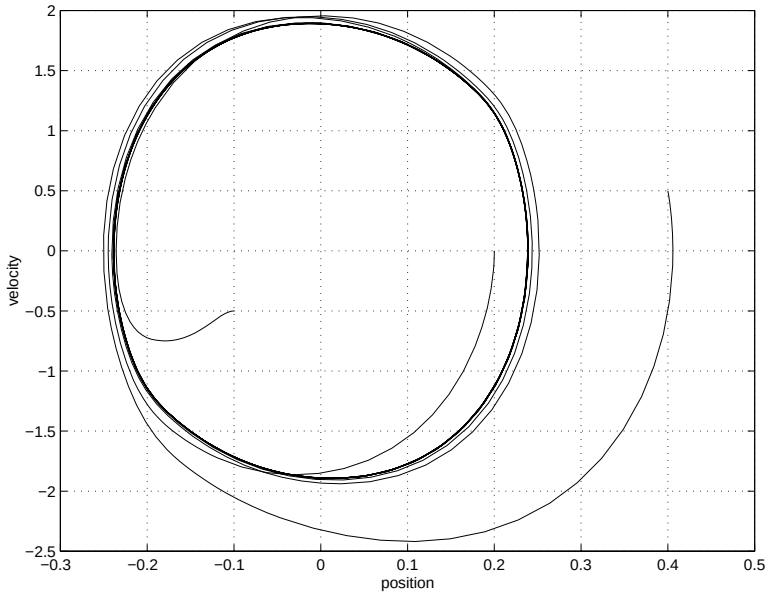

Figure 8: Phase plot of mass motion. This plot shows the velocity plotted against position for a mass actuated by the oscillator under position feedback. The mass is started from three different locations, but the system quickly converges to the limit cycle behavior. For this example $c = 1.0, \tau_1 = 0.1, \tau_2 = 0.2, k = 30, I = 1, b_j = 2$.

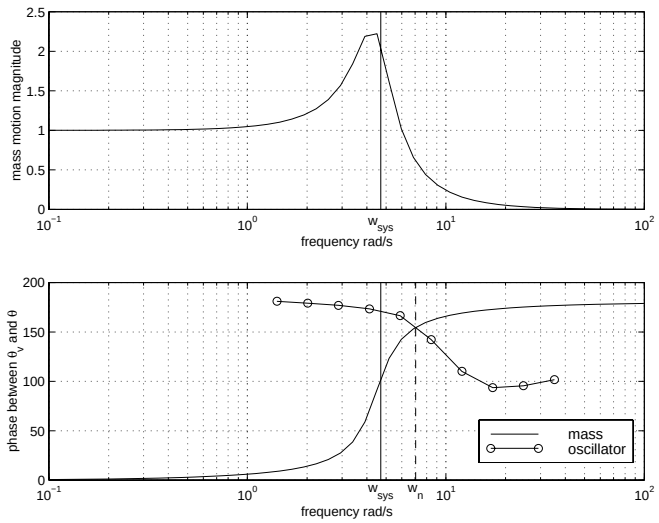

Figure 9: One example of oscillator and mass interaction. The top plot shows the magnitude of the mass motion plotted against frequency, with the obvious resonance peak at w_{sys}. The lower plot shows the phase angle between θ_v and θ for both the mass (solid line) and the oscillator (circles). Since the oscillator and mass systems are tightly coupled, the oscillator with input θ and output θ_v, and the mass system with input θ_v and output θ, the steady state solution to the coupled system occurs when the phase difference across both systems is identical. Given these two systems, the mass would be expected to oscillate at 7 rad/s, the frequency at which the two phase plots intersect. The final frequency is not generally so close to w_n, but varies given different oscillator and system properties. This plot is for the oscillator under position feedback; similar plots could be drawn for torque feedback, which would compare the phase between θ_v and torque u.

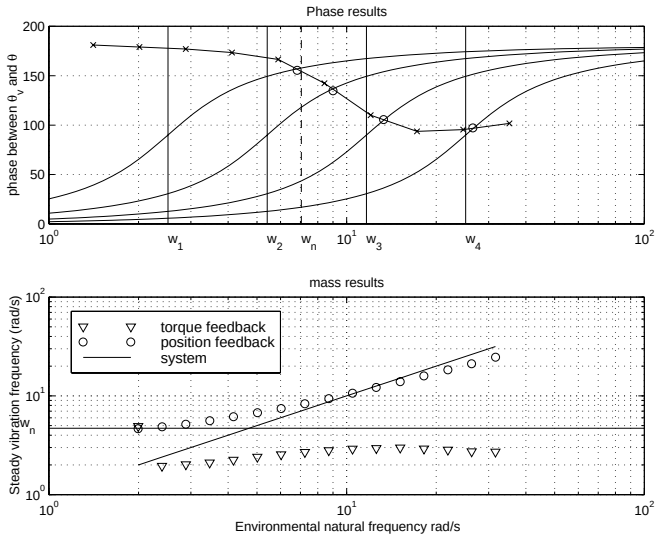

Figure 10: The response of the oscillator when coupled to different systems. The top graph shows a similar scenario to figure 9, plotting the phase behavior of four systems with different resonant frequencies ($w_1, \ldots, w_4$, solid lines), together with the oscillator phase behavior (crosses). The steady state conditions are at the intersections of the lines, indicated by open circles. When the system frequency is less than the oscillator natural frequency (e.g., w_1 and w_2), the steady state frequency is close to w_n. However for w_3 and w_4, both greater than w_n, the final frequency is closer to the resonant frequency of the systems. The lower plot shows this in more detail. Plotted is the steady state oscillation frequency measured by simulating the mass-oscillator system. The plot under position feedback (circles) compares well to the phase information in the upper plot, giving frequencies close to resonance for $w_{sys} > w_n$. Under torque feedback, the system oscillates at a low frequency, the behavior not strongly dependent on the system natural frequency w_{sys}.

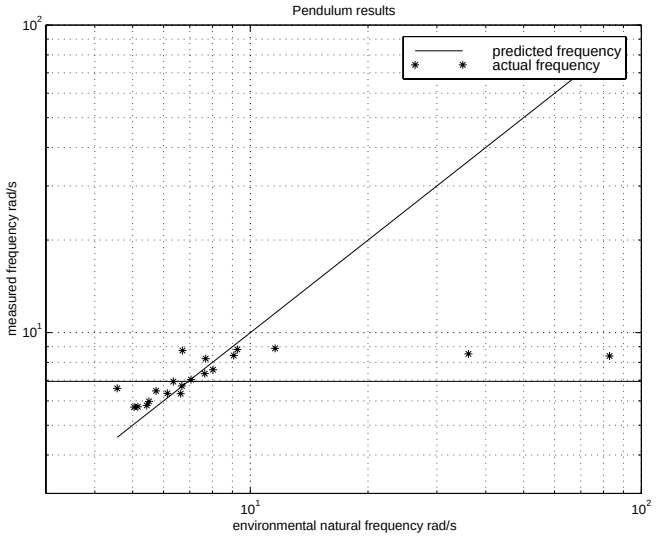

Figure 11: Performance of the robot swinging a pendulum. The graph shows the pendulum frequency when the oscillator is using angle feedback plotted against the natural frequency of the pendulum. The pendulum is swung at its natural frequency over the range 5 to 9 rad/s. The natural frequency of the oscillator (w_n - horizontal solid line) is 7 rad/s making the entrainment range about 60%. The behavior at higher frequencies is most likely due to actuator bandwidth limitations.

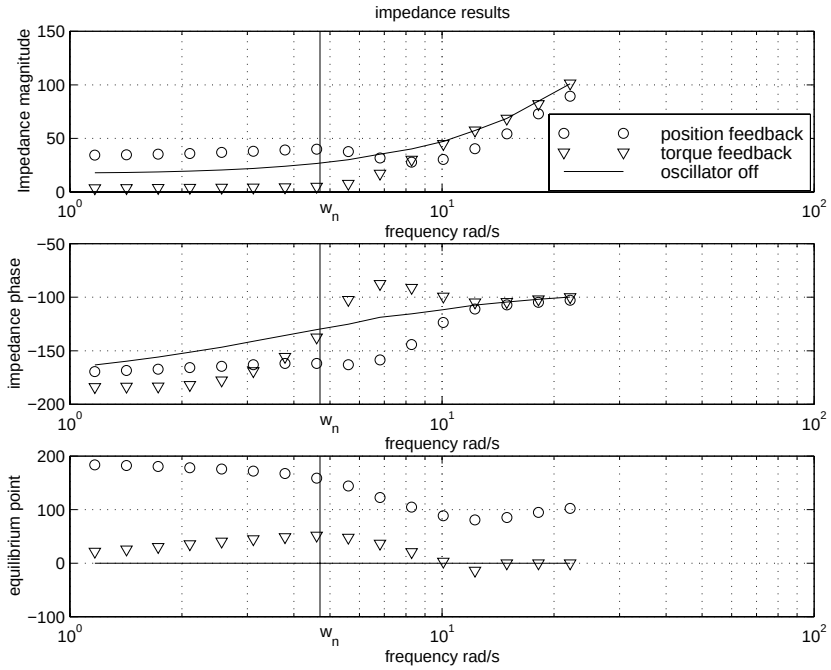

Figure 12: Output impedance of a robot joint driven by an oscillator. The top plot shows the magnitude of the impedance (how much force would be generated at the output, when the input is moved at that frequency), the middle plot the impedance phase, and the lower plot the phase of the joint equilibrium points. The impedance where the oscillator is turned off is indicated by the solid line. The force feedback (triangles) shows much reduced impedance over a range of frequencies less than w_n. The position feedback (circles) opposes the motion at low frequencies (impedance is 180° out of phase with the motion, with large amplitude), while at high frequencies the equilibrium point (lower graph) actively drives the motion (leading by 90°).

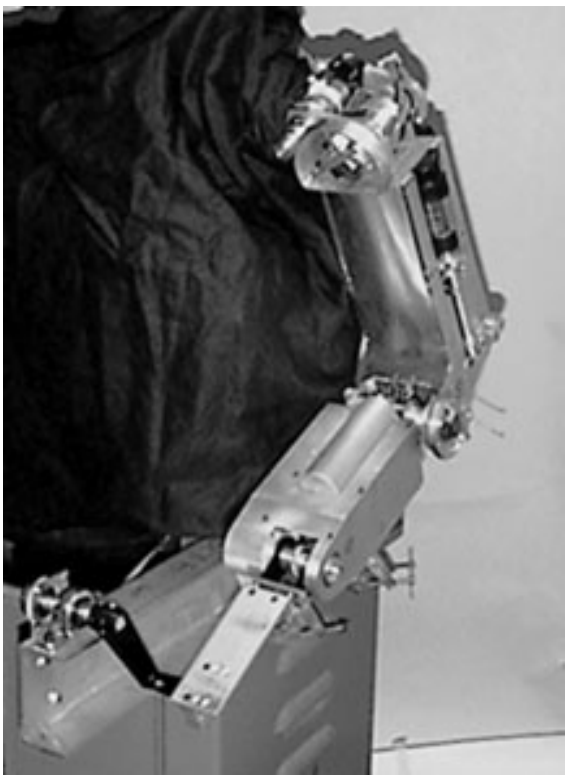

Figure 13: Photograph of the crank turning. The crank is turned in the plane of the arm, using the shoulder and elbow joints. The wrist joint is used for the crank turning, but is passive, not being driven by an oscillator.

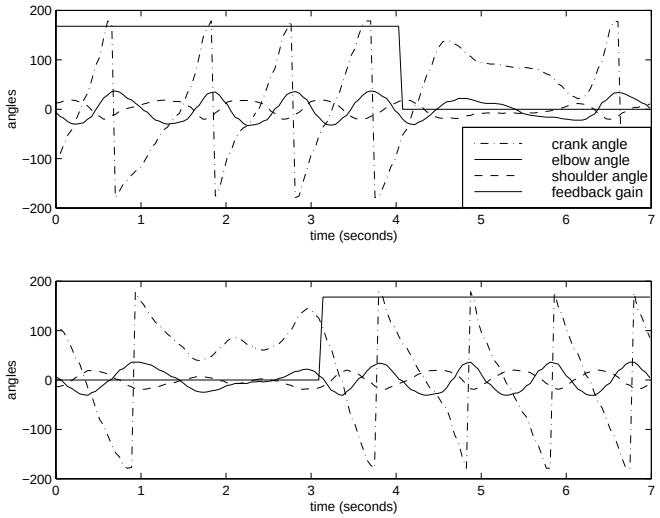

Figure 14: Two examples of shoulder and elbow angles during crank turning. The angle of the crank is also illustrated (dash-dot), the saw tooth shape arises due to the position sensor for the crank wrapping around. The shoulder was driven at a constant frequency, and an oscillator with position feedback was used at the elbow. When the feedback is on, the angles are coordinated and the crank is turned, otherwise there is no coordination.

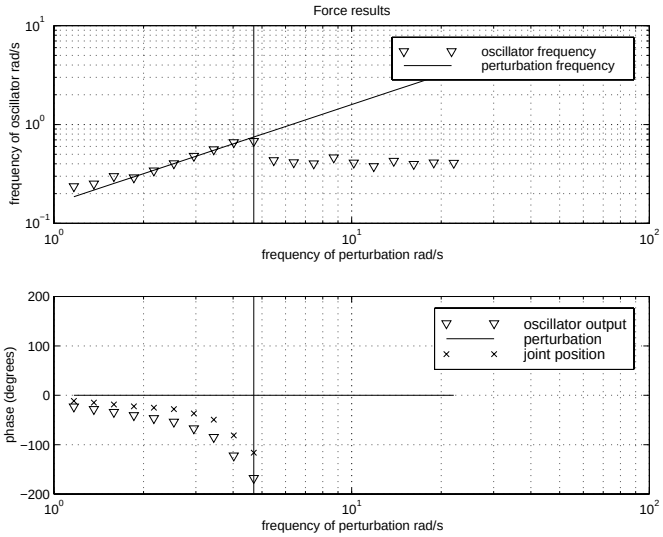

Figure 15: The response of the oscillator under a perturbing force. The top graph shows the oscillator output frequency (triangles) and the perturbation frequency (solid) plotted against the perturbation frequency. It is only at frequencies below w_n (vertical line) that the oscillator is entrained (range 1.2 to 4 rad/s). The lower graph shows the phase of the oscillator output (triangles) and the joint angle (crosses) over the entrainment region. Both variables lag the driving force by small angles which increase with frequency. Since the oscillator is not entrained above w_n it is impossible to calculate phases for those frequencies.

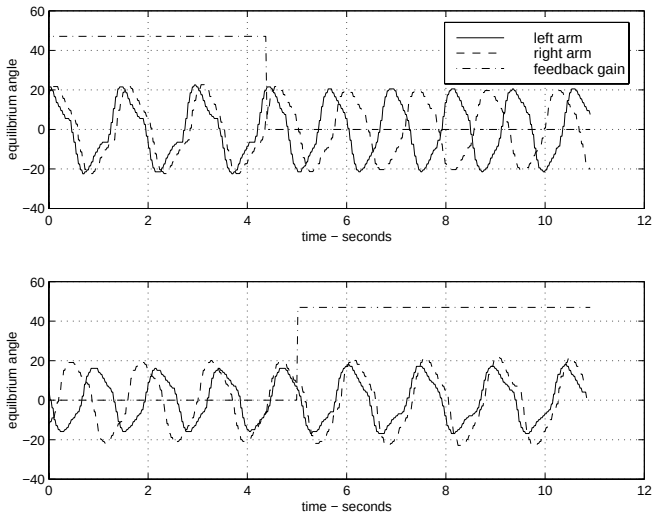

Figure 16: Two examples of slinky operation. Both plots show the outputs from the oscillators as the torque feedback (dashed) is turned on and off. When the traces are in phase, the slinky is moving in anti-phase. When the feedback is on, the two arms are coordinated and the outputs are synchronized, but when off, the oscillators are no longer synchronized. The only connection between the oscillators is through the physical structure of the slinky.

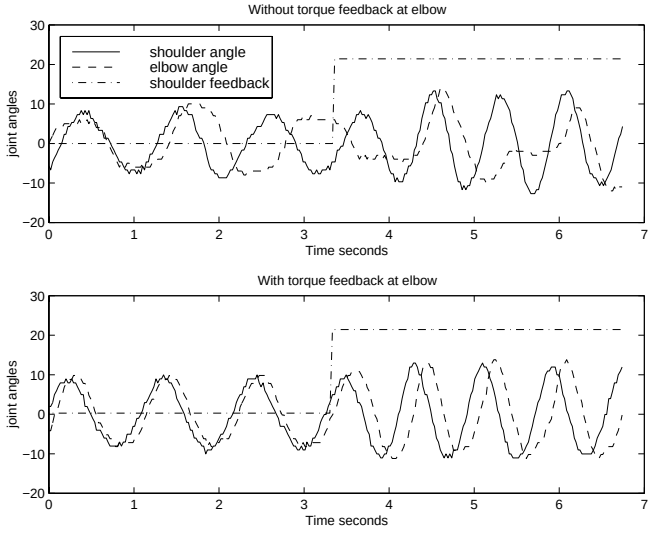

Figure 17: Flailing. Both plots show the angle of the shoulder (solid) and the elbow (dashed) as the frequency of the shoulder is changed. The speed of the shoulder is changed by applying position feedback to the shoulder oscillator (changes the frequency from w_n to a higher frequency—see Figure 10). The top graph shows the response of the arm without torque feedback at the elbow, and the lower graph with force feedback. The synchronization is clear in the lower graph, the only connection between the joints being through the physical structure of the arm.

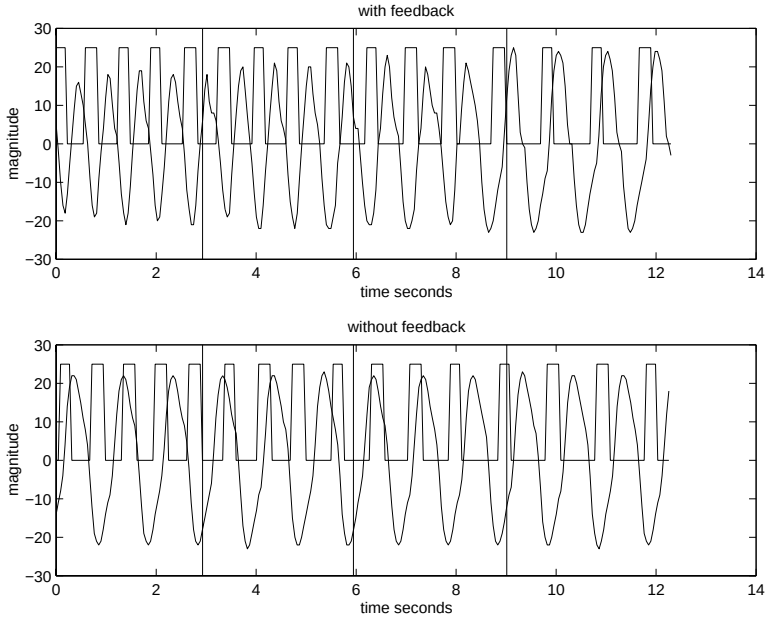

Figure 18: The response of a robot joint to force impulse. The graphs show the perturbing force (the square wave) and the equilibrium point of the joint as the speed of the impulses is changed. The vertical lines indicate where the frequency changed, going from 1.6Hz to 1Hz, in approximately equal steps. When force feedback to the oscillator is provided (upper graph), the output maintains the same frequency as the disturbing force, while when there is no feedback (lower graph), there is no tuning.