Appendix A

Experimental Apparatus

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

A.1 Overall arm design

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

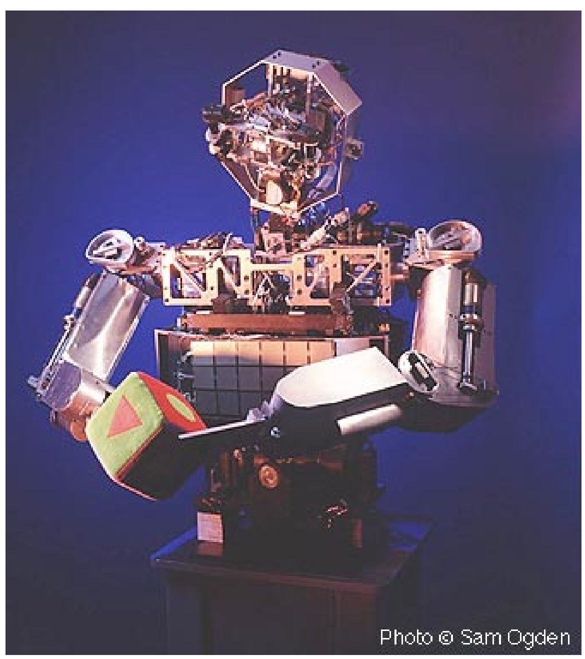

Figure A-1: Picture of the two robot arms. Each arm has 6 degrees of freedom, arranged in a similar manner to a human arm. The arms are approximately the same length as a human arm

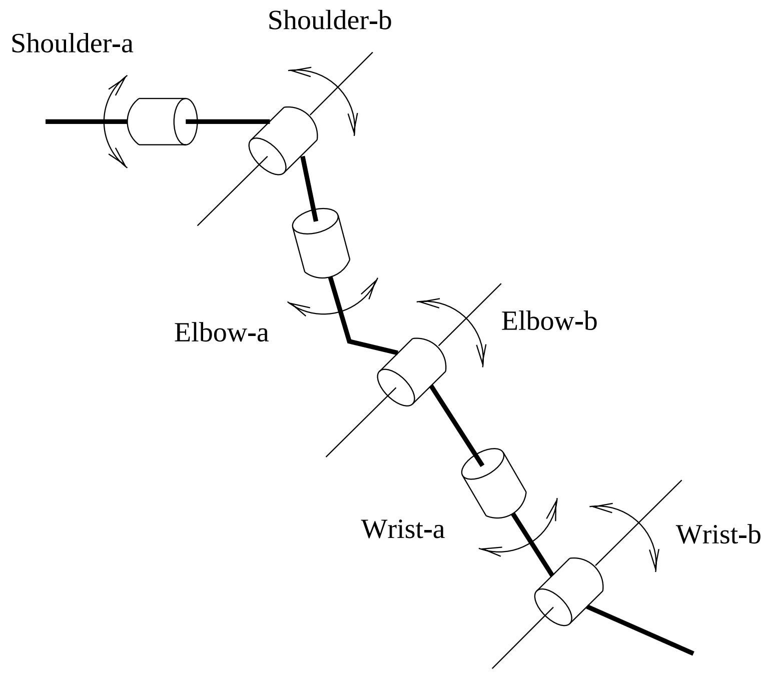

Figure A-2: Arm kinematics. The thick lines indicate the structure of the arm, and the thin lines the axes of joints which do not line up with the arm structure. There are two joints for each of the shoulder, elbow and wrist, making six degrees of freedom in total. The axes of the shoulder-a, shoulder-b,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

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

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

A.2 Actuators

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

Another consideration comes from robot force control. Early researchers found that controlling stiff robot arms to contact hard surfaces was difficult, but that by covering the robot is some compliant material (foam rubber), the task was made somewhat easier (Whitney, 1977).

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

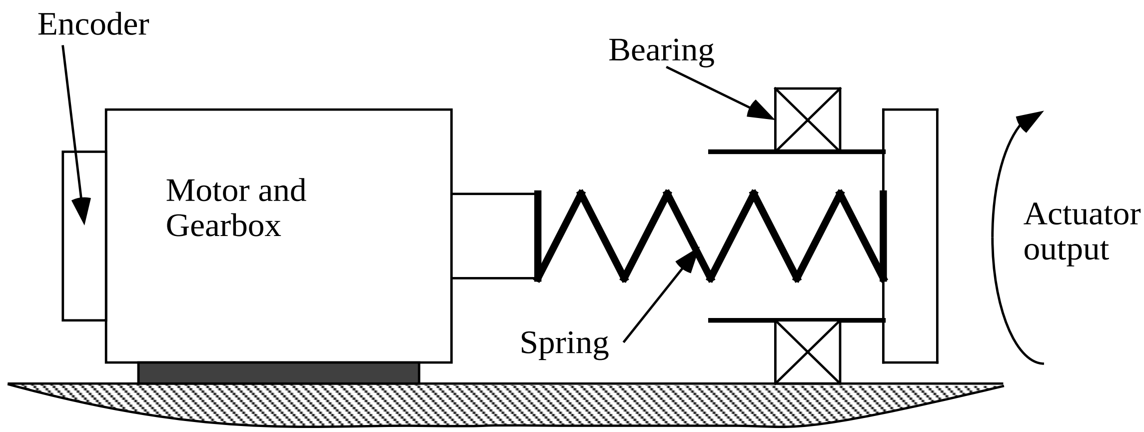

Figure A-3: Schematic of actuator design. The series elastic actuator consists of an ordinary electrical motor and gearbox, with a spring connected between the motor output and the actuator output. By controlling the deflection of the spring, the output force of the actuator can be controlled. The spring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

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

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

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

A.3 Spring design

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

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

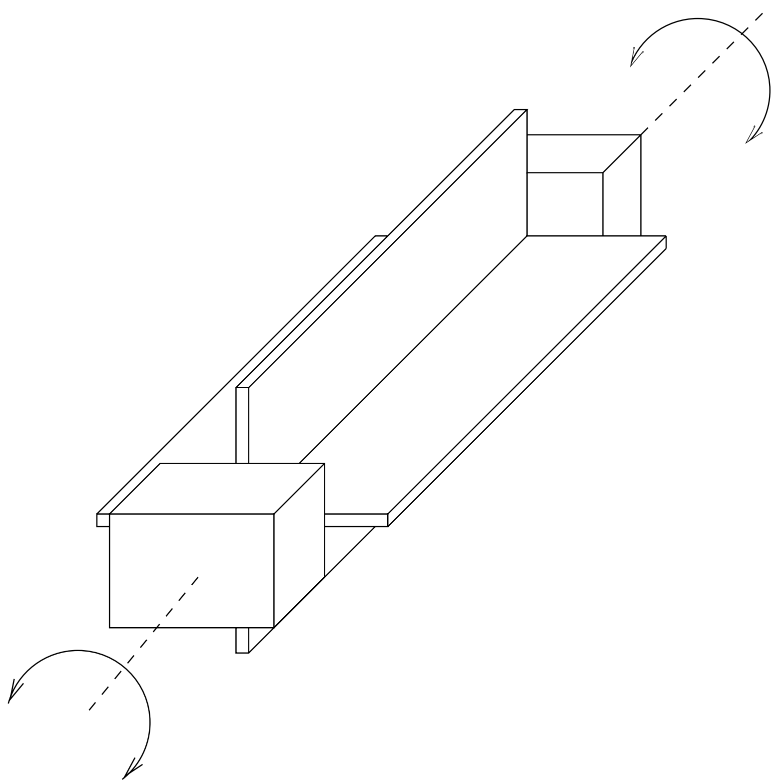

Figure A-4: Spring design. The figure shows a torsional spring with a cross-shaped cross-section. This shape gives a spring which is strong while still being flexible.

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

$$\theta_{max} = \frac{l\tau_{max}}{tG} \tag{A.1}$$

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

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

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

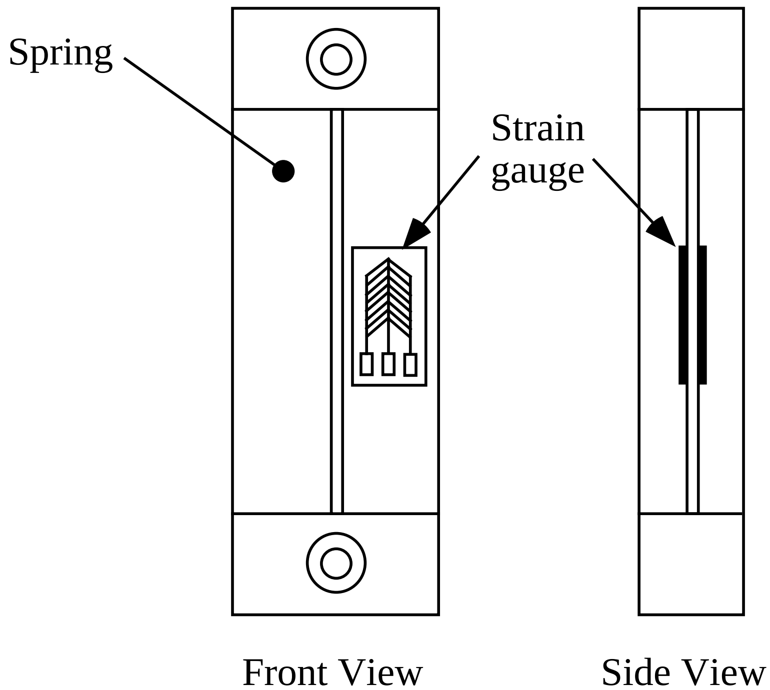

Figure A-5: The deflection of the spring is measured using strain gauges mounted on the flats of the spring, as shown here. Four gauges are used in a Wheatstone Bridge configuration to measure the strain in the spring, and so the deflection. The deflection is proportional to the torque transmitted through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

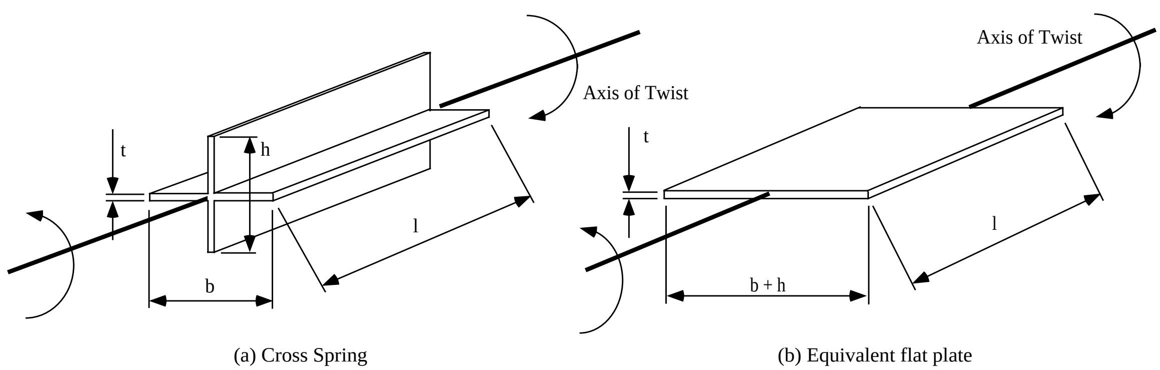

Figure A-6: Figure comparing the cross-shaped spring with a flat plate in torsion. Since the cross-shaped spring is composed of flat plates in torsion, its behavior is related to the behavior of a flat plate. The equivalent flat plate spring has the same length (l), web thickness (t) with a width the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

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

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

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

$$\theta_{max} = = 1.65\frac{l\tau_{max}}{tG} \tag{A.4}$$

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

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

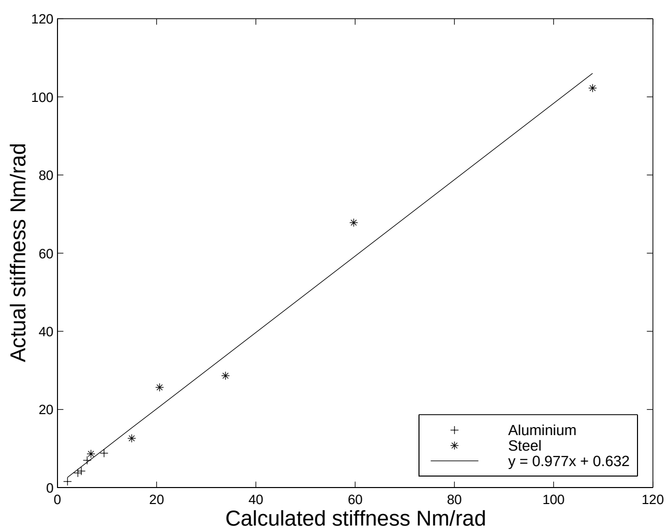

Figure A-7: REAL Comparison of cross-spring stiffness and flat plate stiffness. The figure shows the measured stiffness of a number of springs plotted against the stiffness of an equivalent flat plate of similar dimensions and material. The points lie on a straight line, with gradient 0.98. This suggests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

A.4 Actuator control law

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

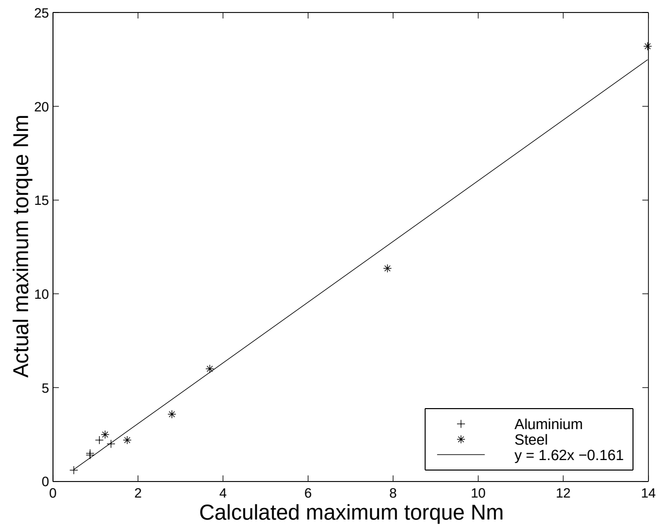

Figure A-8: REAL Comparison of maximum torque before yielding for a cross-shaped spring and the calculated value for an equivalent flat plate. The points lie on a straight line, indicating that the flat plate is a good model of the cross-shaped spring's behavior. The gradient of the line is 1.62, showing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

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

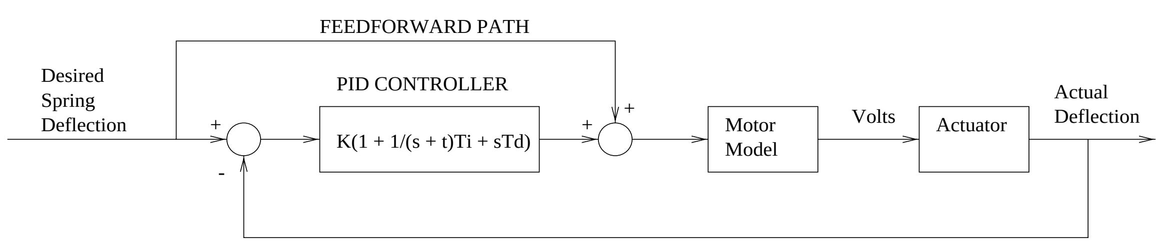

Figure A-9: Figure showing the control law for the actuator. The desired force of spring deflection is the input, and a PID controller is used to control the actual deflection to track the desired deflection. The integral term is rolled off at low frequencies to preserve passivity (see Williamson (1995)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

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

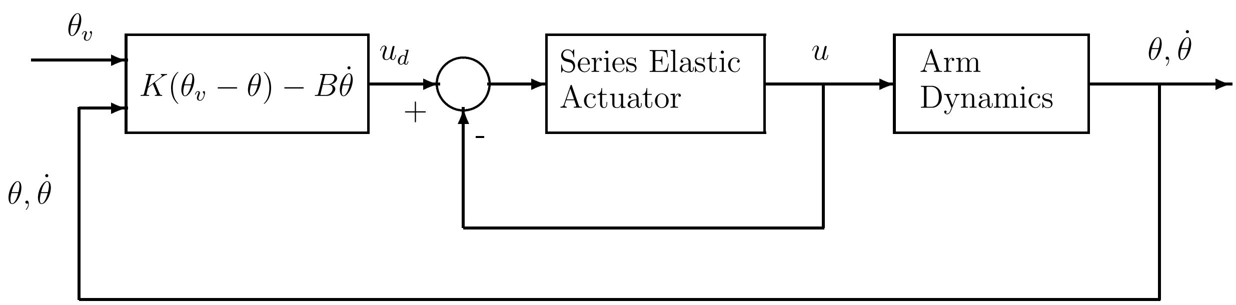

Figure A-10: The joint level control consists of a position control loop, with stiffness K, damping B, and setpoint θ_v organized around an inner torque control loop provided by the series elastic actuators. The inner loop controls the force output of the joint u to accurately track the desired force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

A.5 Arm design

A.5.1 General

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

A.5.2 Joint configurations

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

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

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

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

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

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

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

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

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

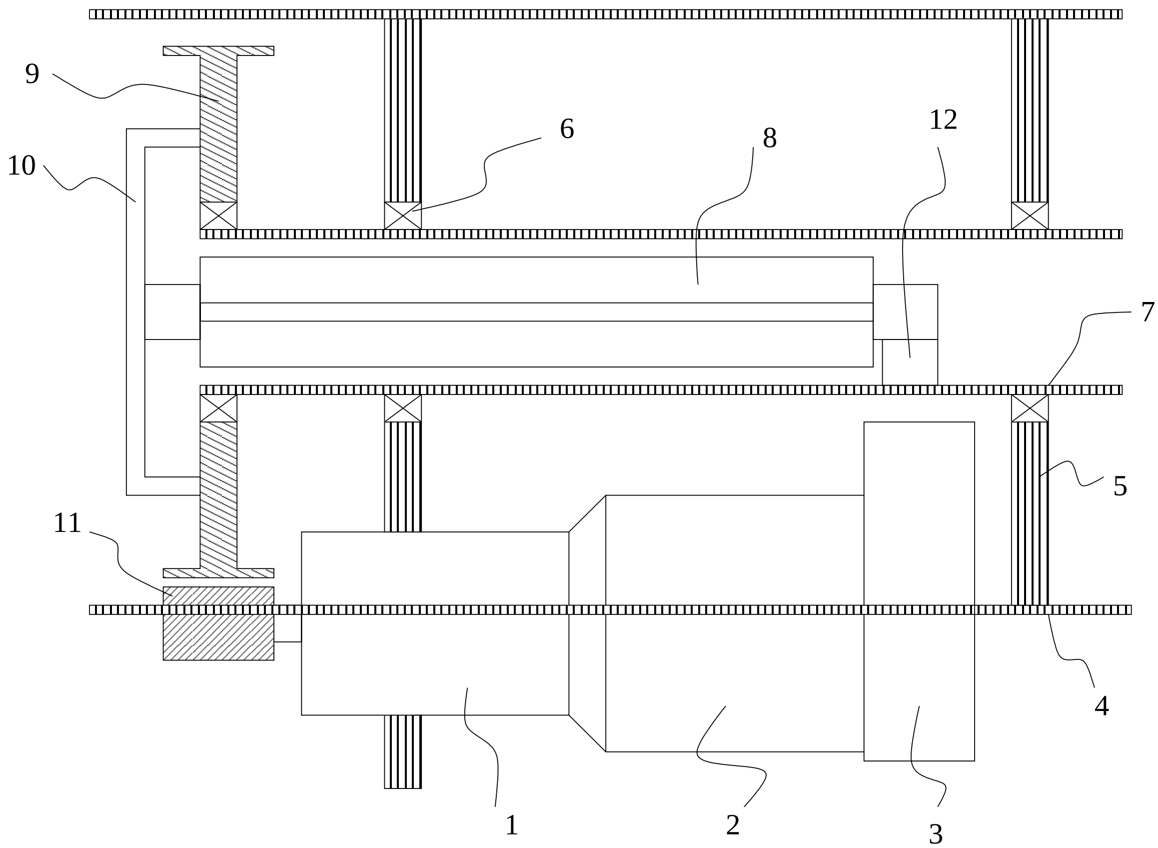

Figure A-11: Arm joint configuration I. The figure shows schematically the layout of the series elastic actuator at the elbow-a and wrist-a joints. The motor (gearbox - 1, motor -2, encoder -3), is arranged to be parallel to the spring (8). The motor drives a small pulley (11), which is connected with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

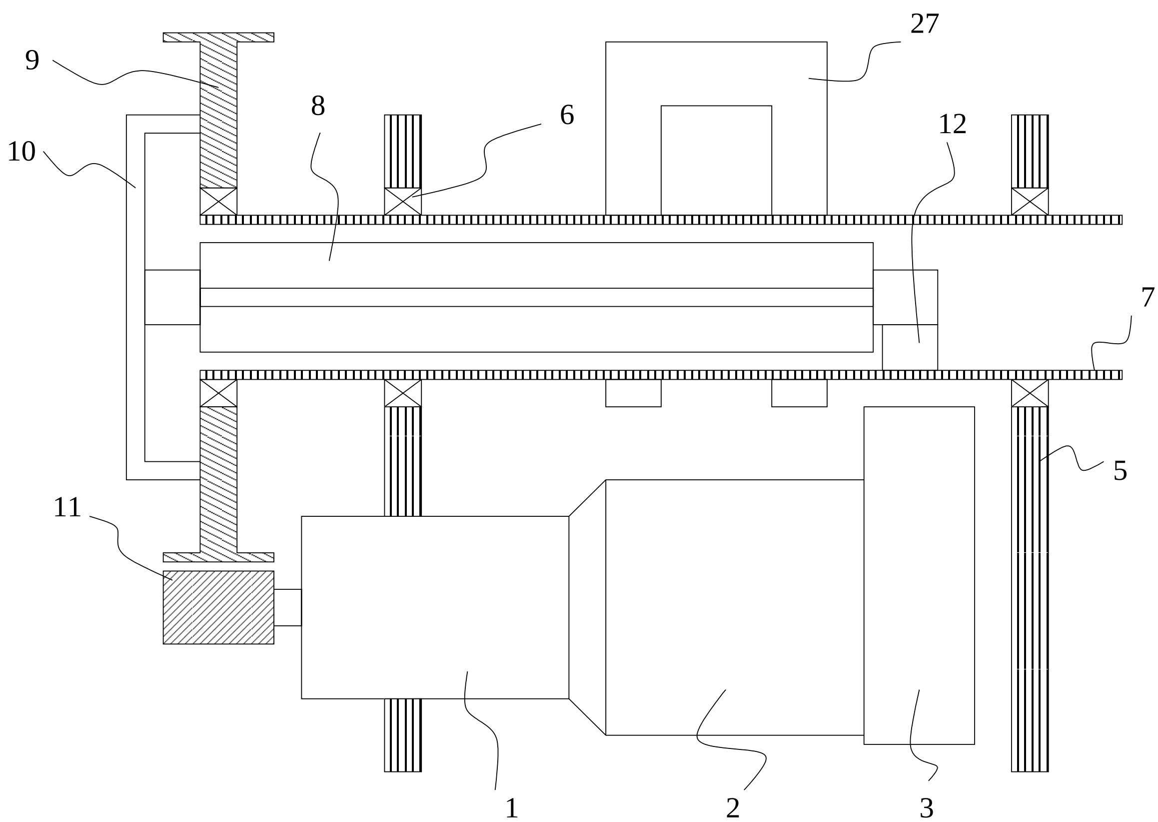

Figure A-12: Arm joint configuration II. This configuration is used for the shoulder-a, shoulder-b, elbow-b and wrist-b joints. The configuration is largely the same as in figure A-11, with the difference that the inputs and outputs to the actuator have been swopped. The motor and supports (1,2,3,5) is fixed, and the tube (7) moves. The output of the actuator can either be the right hand end of the tube (used in shoulder-a), or through the clamp (27) (used in the other joints). The advantage of using a clamp is that it can be easily removed.

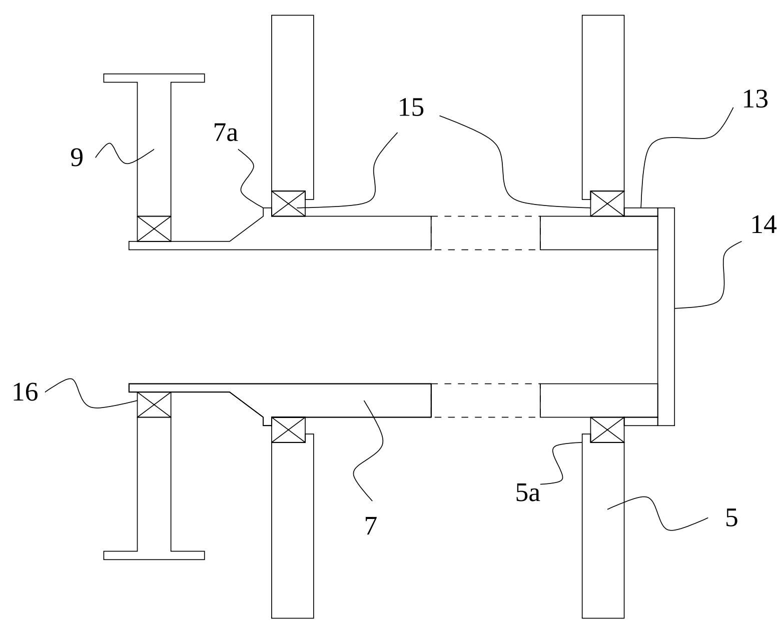

Figure A-13: The same bearing configuration was used for all the joints of the arm, the only difference being the sizes of the various tubes. The tube is supported on two angular contact bearings (15). These are preloaded by shoulders on the tube (7a), shoulders on the supports (5a), a spacer (13) and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

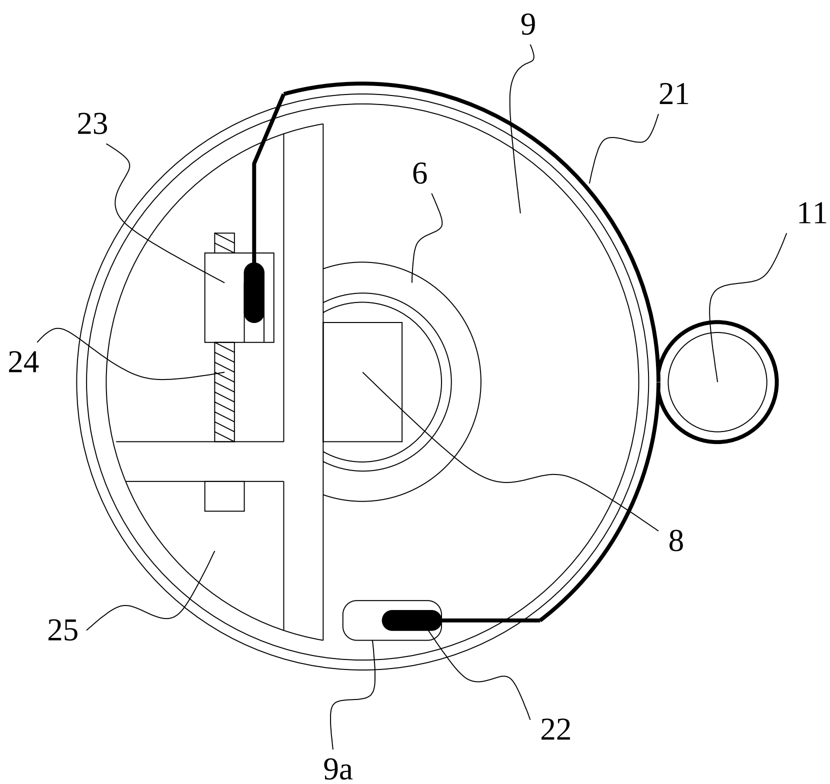

Figure A-14: Cable tensioning mechanism. The same tensioning mechanism was used for all the cables in the arm. One end of the cable (21) is terminated in a slot in the large pulley (9a). It wraps around the large pulley (9), around the motor pulley (11), and terminates in the tensioning block (23). The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

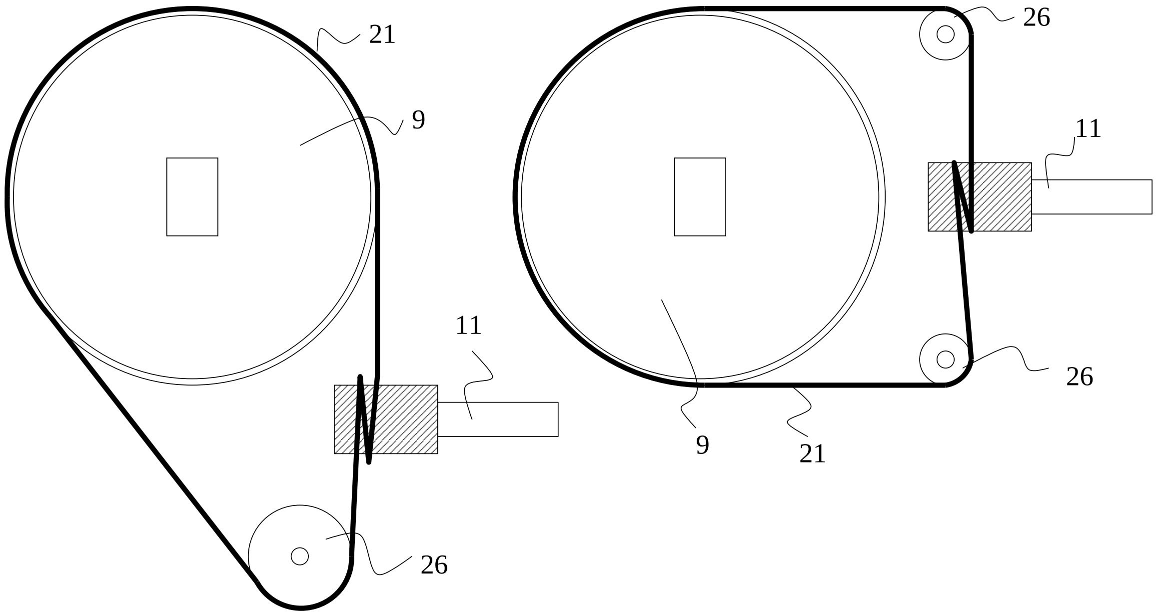

Figure A-15: The left hand figure shows the cable configuration used for the elbow-b joint. An extra pulley (26) is used to transfer the cable drive from the motor pulley (11) to the large pulley (9). This allows the motor shaft to be perpendicular to the axis of the joint, freeing up the motor position.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

A.6 Arm technical data

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

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

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

A.7 Sensors

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

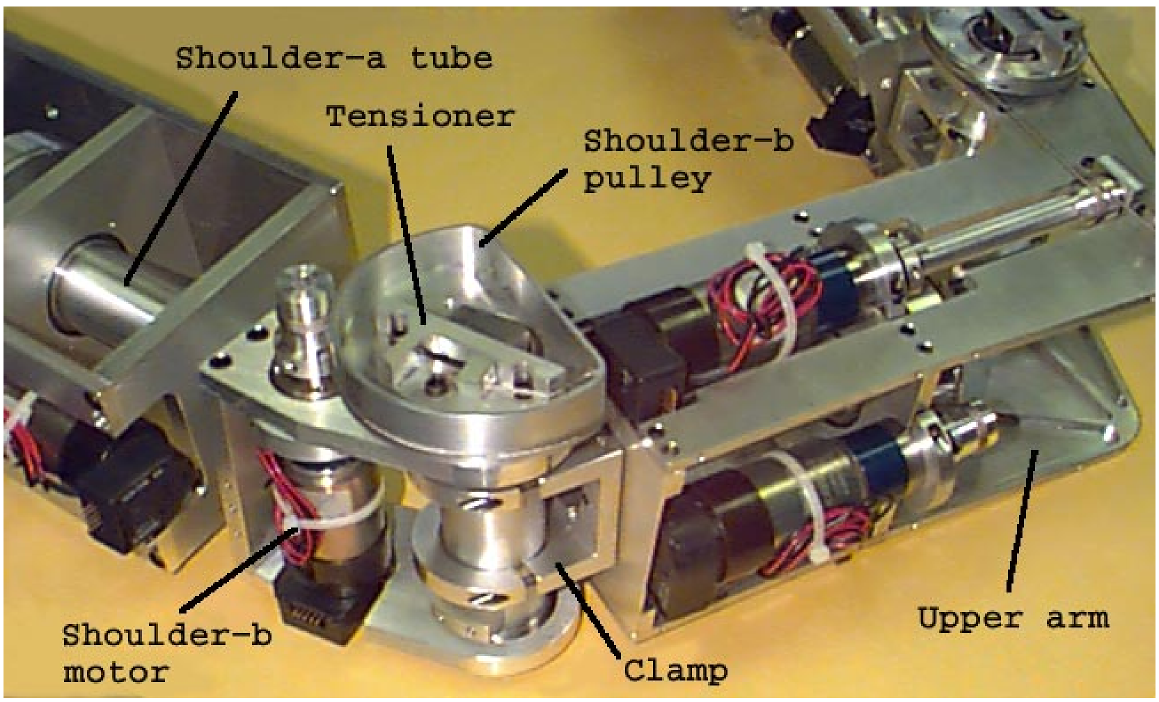

Figure A-16: Robot shoulder joint. The photograph shows the shoulder joint and the upper arm. The shoulder-b joint is shown most clearly with the motor, pulley, tensioning mechanism tube and clamp all in the middle of the picture. The upper arm is clamped onto the shoulder-b tube in order to allow quick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

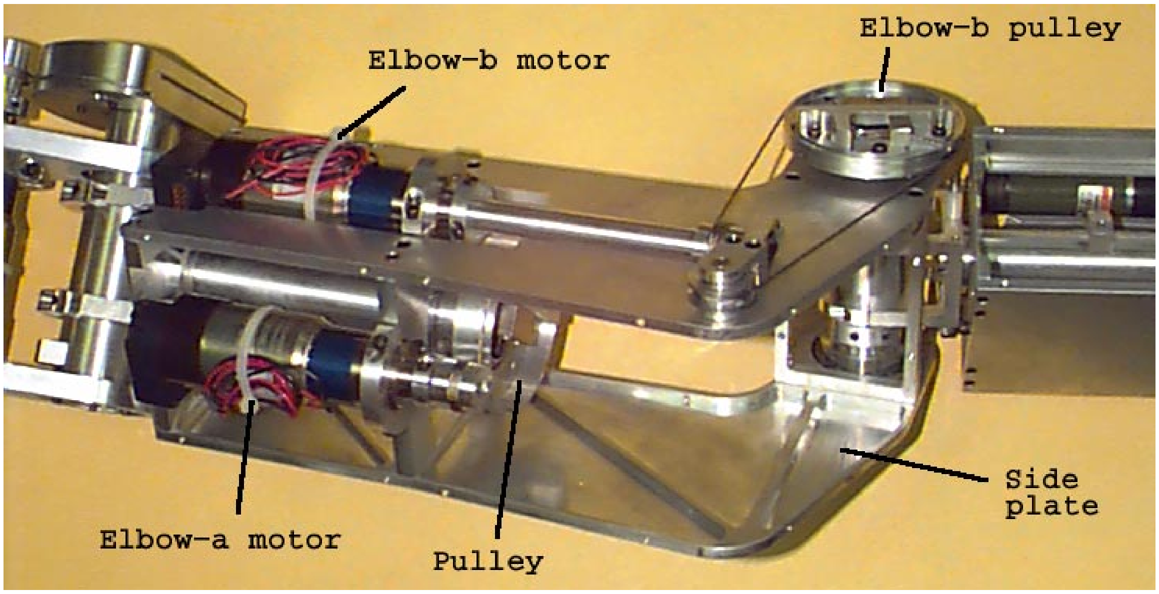

Figure A-17: Photograph of the upper arm of the robot. Both motors are placed high up on the arm, so reducing the inertia of the link. The elbow-a motor is arranged as described in figure A-11 walking around the tube, and driving the whole upper arm around. The elbow-b motor drives the joint through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

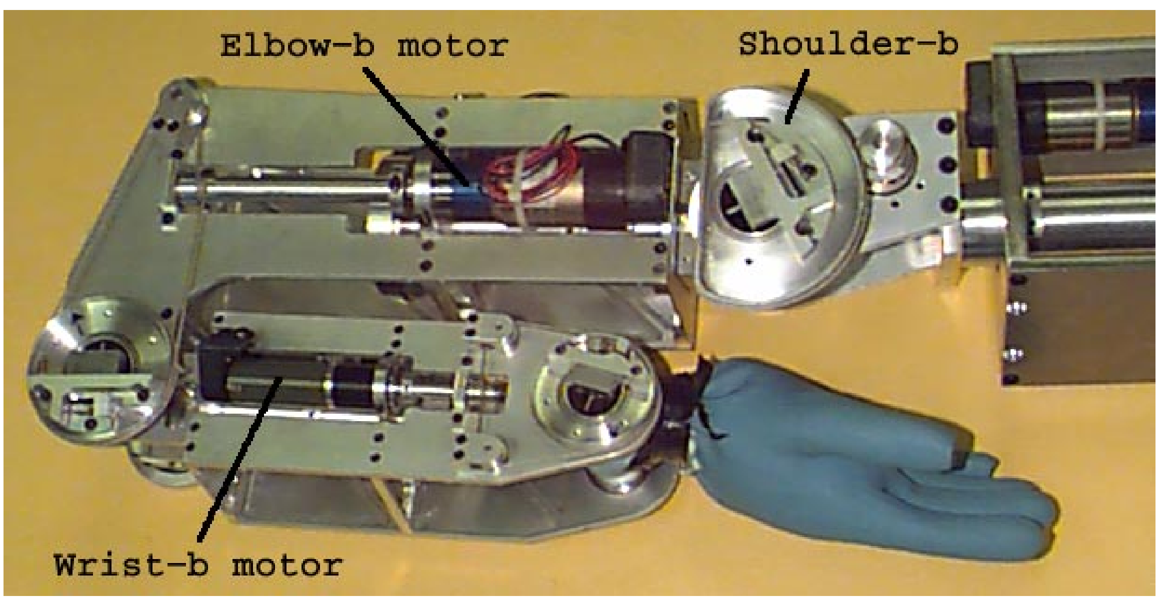

Figure A-18: Figure showing the arm curled up. The elbow-b joint is offset from the axis of the arm to allow the arm to curl up into this position.

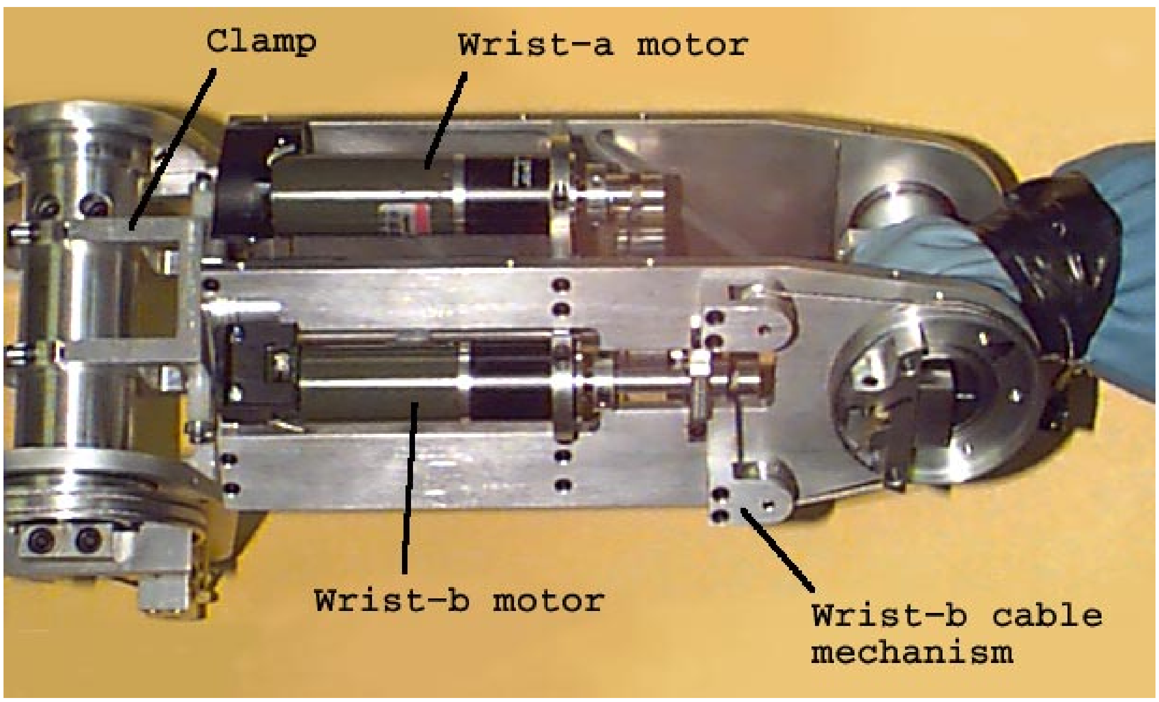

Figure A-19: Lower arm design. The upper motor drives the wrist-a joint, moving the whole lower arm section around its longitudinal axis. The lower motor drives the wrist-b joint using two extra pulleys. The lower arm is attached to the elbow-b joint using a clamp to facilitate assembly.

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

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

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

- Strain gauge processing. The integrated circuit which provided a stable voltage for the strain gauges, and processed the strain gauge output was included. Tuning potentiometers were included to alter the gain of the amplifier, as well as its offset. The Analog-to-Digital (A/D) converter used on the motor controller boards has a voltage range of $0-5$ Volts. This means that input signals must have a zero reading of 2.5v and swing by ± 2.5 v around that zero. The sensor board also included a buffer on the strain gauge chip output in order to drive the signal along the wire to the motor controller.

- Strain velocity. The force control loop requires signals recording the strain as well as the rate of change of strain in order to control the force accurately. The rate of change of strain (Strain velocity) could be calculated digitally, but with only an 8 bit A/D on the motor controller, the errors would be large. The sensor board thus included an analog differentiation circuit and buffer to produce this signal directly.

- Joint angle. The board included an amplifier and level shifter so that the maximum, minimum and zero voltage for the joint angle motion could be set. It also included a differentiating circuit to calculate joint velocity. Setting the joint angle on the board sets the calibration for all the joints. This means that the robot can be powered on and off without having to recalibrate the joint limits. This means that the calibration is relative to the voltage in the system, so a well regulated power supply is required.

- Encoders. The board provided power distribution for the motor encoders.

- Spare channels. The board provided two extra connectors, connected back to A/D channels on the motor controller boards. The connectors for these have power and ground also, which makes adding extra potentiometers, touch sensors etc. easy.

A.8 Electrical wiring

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

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

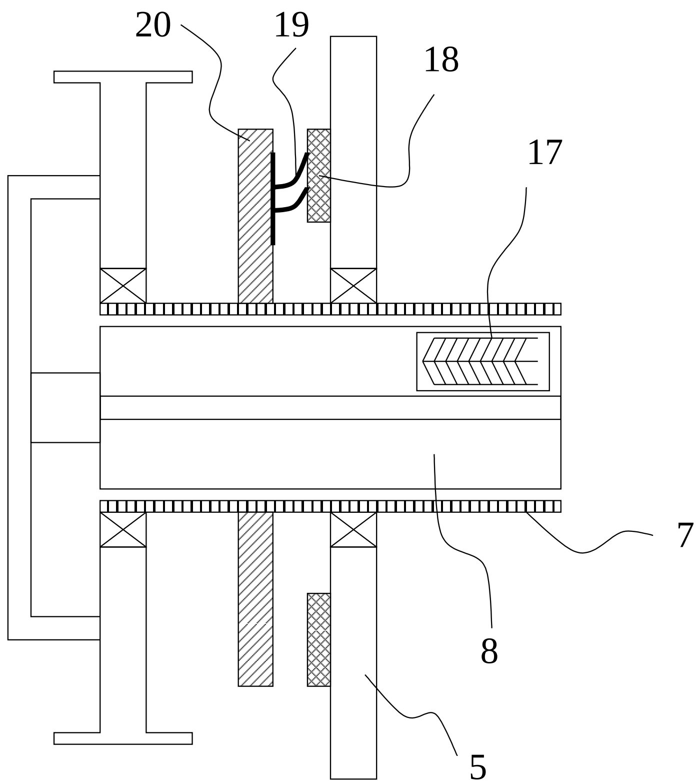

Figure A-20: In addition to the encoders on the backs of all the motors, each joint has a strain gauge (17) mounted on the spring (8) which measures the torque at the joint, as well as a potentiometer. The potentiometer is of a hollow shaft construction, with the resistive element (18) forming a ring mounted on the support (5). A plastic ring (20) clamped to the tube (7) holds a wiper (19), which completes the potentiometer. The measurement is of the angle between the tube and supports, so measuring the angle of the joints.

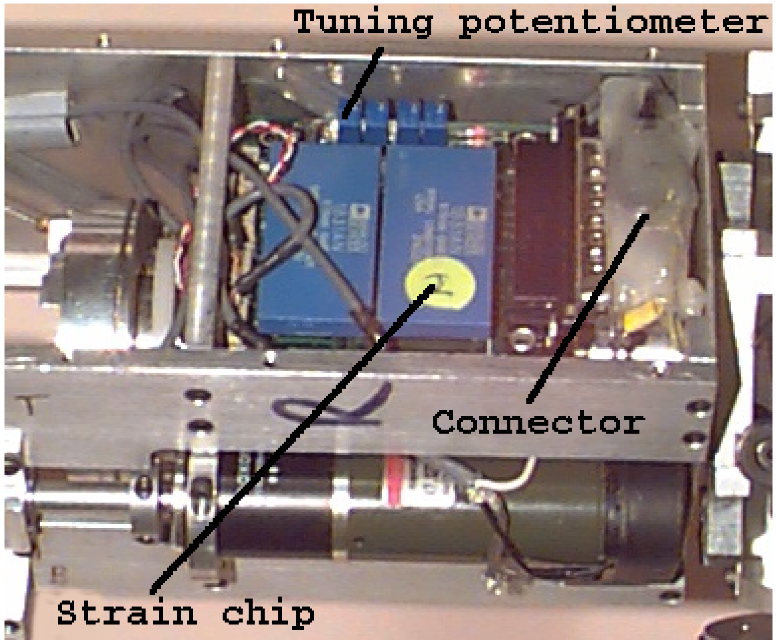

Figure A-21: Photograph of the sensor board. The large chips process the strain gauges, and the tuning potentiometers are used to adjust gains and offsets for all the signals. The power to the board and the outgoing signals are carried by the large connector.

used to test all the signals coming from the sensor board. This board together with the modularity of the whole system proved invaluable for debugging.

Physically routing the cables through the robot is somewhat of an art form. The arm was designed so that the majority of cables could pass through the axes of the joints as shown in figure A-25. This reduces the strain on the cables, and makes them less likely to break. The wires were passed through the joints for the wrist and elbow units, but at the shoulder there were too many wires, so they were passed around the joints.

A.9 Control hardware

The control for the arm joints was implemented using a home made motor board, illustrated in figure A-26. This board contains all that is necessary to control two joints. The heart of the board is a Motorola 6811 microcontroller which runs the control program. This is written in assembly language and runs with a loop sampling rate of 2kHz. The 6811 has A/D converters which are used to sample the analog inputs (joint angle, velocity, strain, rate of change of strain). Extra circuitry processes the signals from the motor encoders. The board uses a PIC to generate the Pulse Width Modulated (PWM) signal which is used to command power to the motors. An H-bridge motor amplifier is also included on board which provides power to the motors.

The board communicates with a network of Motorola 68332 processors which perform all the higher level processing for the arm, as shown in figure A-27. The communication is low bandwidth, with the 68332's and 6811's exchanging 13 byte packets at around 50Hz. This motivated coding the control law for force as well as the spring-like joint behavior in the 6811's. The 68332's thus send commands such as the setpoint of the joint as well as the stiffness and damping at 50Hz. Data from the robot, joint position, torque, etc. is thus also read at 50Hz.

The 68332 processors run a version of Common Lisp called L (Brooks, 1995). L is a multi-

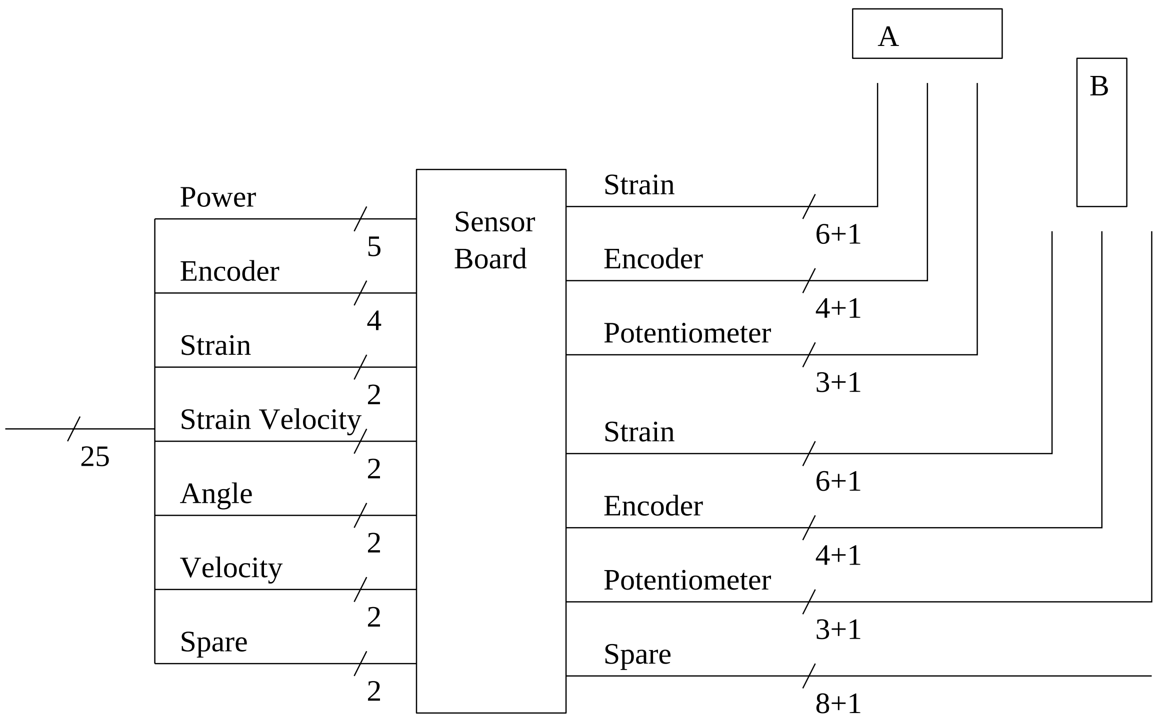

Figure A-22: Each joint has a strain gauge, which requires 6 wires, an encoder (4 four wires) and a potentiometer (3 wires). These all require shields indicated by the $+1$. Rather than pass these cables all the way up the arm, a sensor board was used to process the signals and package the signals into a fewer number of wires and cables. Each board manages two joints worth of electrical signals. Although only 20 wires are required to carry all the signals, the connector had 25 pins, so the extra wires were used to carry power and ground to the board.

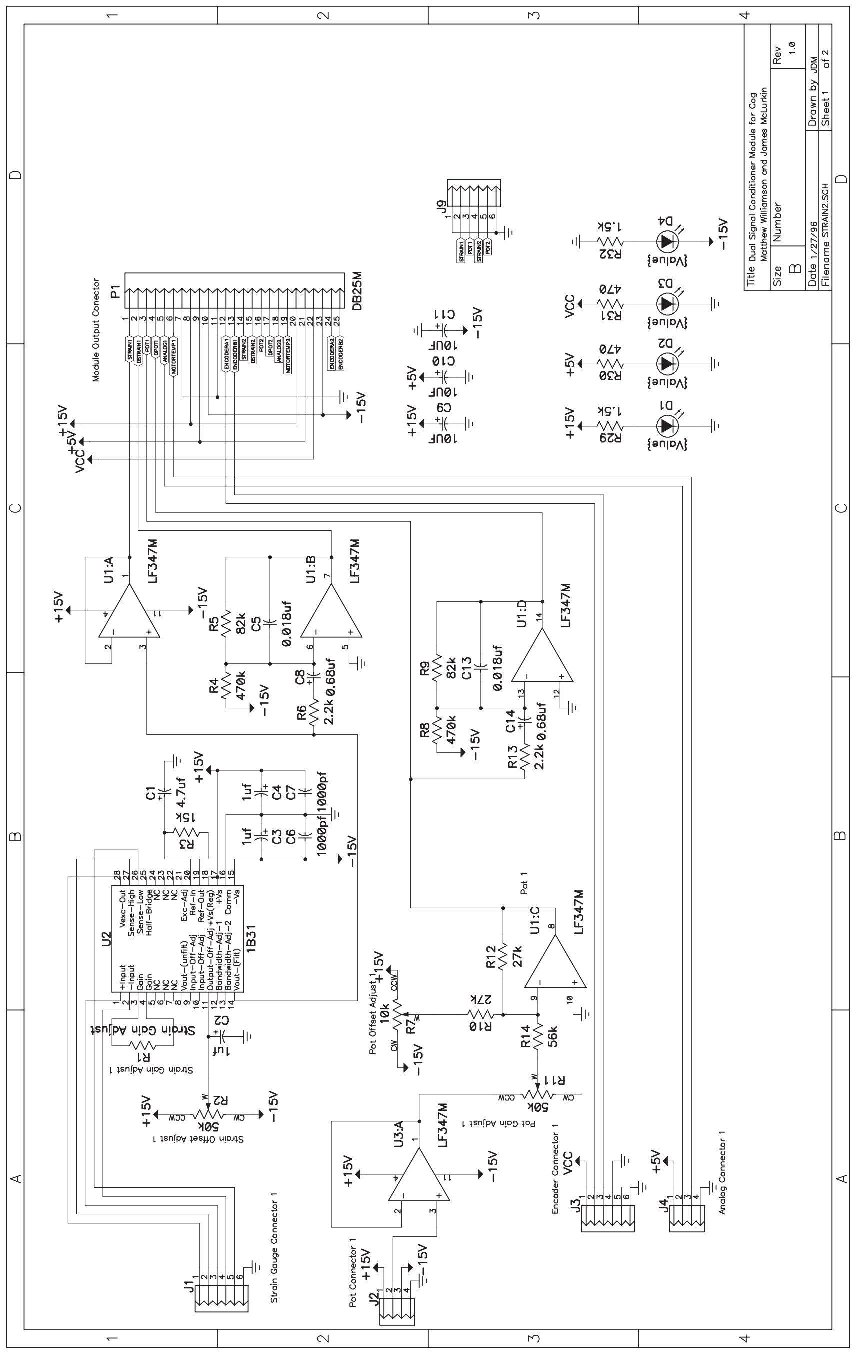

Figure A-23: Schematic of sensor board. The connectors for potentiometer, strain gauges, encoders and spare channels are shown at the bottom, and the output connector at the top. This schematic shows the circuit for one joint. The circuit for the second is identical.

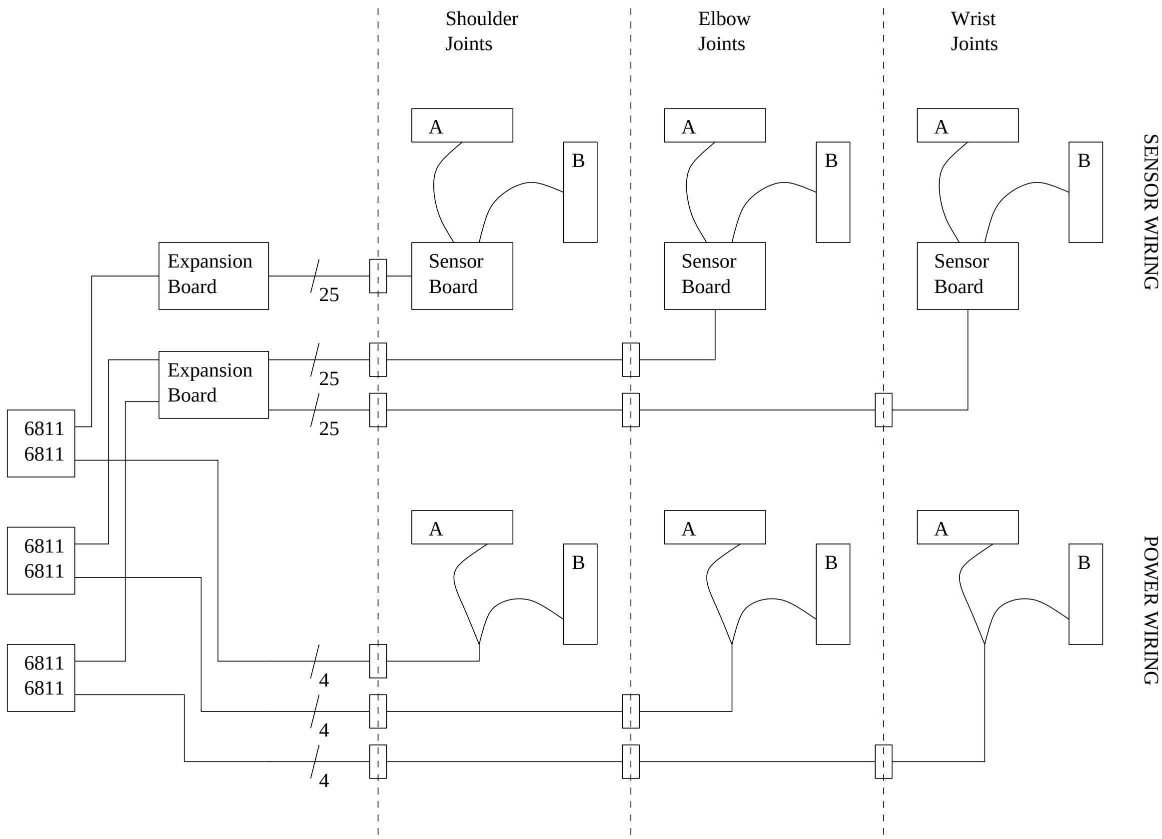

Figure A-24: Wiring diagram. The top half of the figure shows how the sensor cables from each set of joints are passed back along the arm to connect with expansion boards, and then to the motor controller boards (marked "6811"). Each sensor board handles the signals from two joints, and each expansion board handles two sensor boards. Each motor controller can drive two motors. The lower half of the figure shows the cables for the motor power. The dashed lines indicate where the arm can be easily split apart, with the boxes aligned with the lines indicating connectors. Thus the wrist joint can be removed by disconnecting only two connectors.

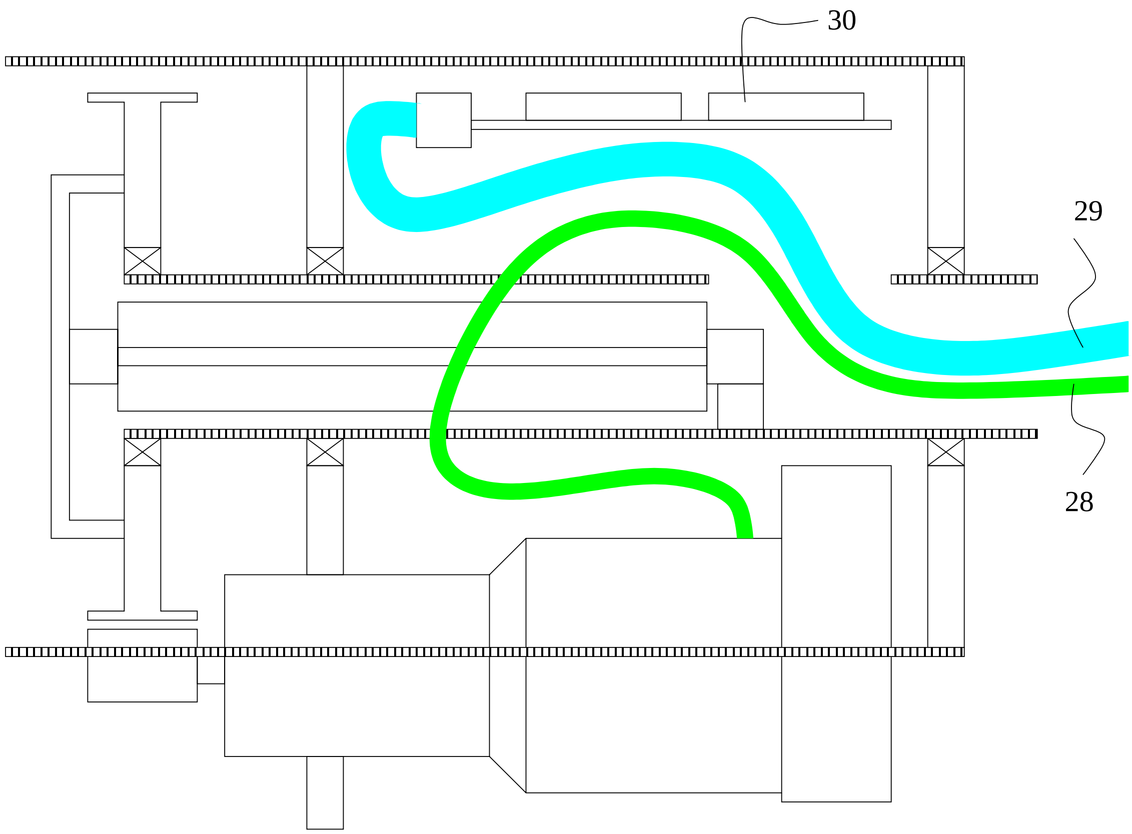

Figure A-25: Figure showing how the sensor board (30) is mounted between the supports. The cables from the sensor board (29), and those carrying power to the motor (28) are passed through the tube and so up the arm. Arranging the cables in this way reduces strain on the wires, and reduces the likelihood of cable breakages.

processing LISP environment, which was convenient as a language to control the arms. The controller for the arms was partitioned into different processes which communicated using shared variables. The L Listener provided a simple way to connect, examine, spawn and kill processes. The 68332's are organized in a network as shown in figure A-27. The connection is also low bandwidth, using dual-ported RAMs (DPRAMs) to share information.

For example, in the configuration shown in figure A-27, the 68332 labeled 'POSTURE' communicates the posture of the arm to the two processors which run the oscillators ('LEFT OSC' and 'RIGHT OSC'). The 'LEFT OSC' and 'RIGHT OSC' processors can communicate to implement connections between the oscillators (as described in chapter 5). The command to the motors is passed to the 'LEFT ARM' and 'RIGHT ARM' processors. Sensory signals from the arms are passed through the 'ARM' processors to the 'OSC' processors.

The network of processors is distributed with low bandwidth communication both within the network, and also with the arm joint controllers. This arrangement motivated some of the solutions found in this thesis. The example has no global clock, or global synchronization. There is also no direct communication between the joints.

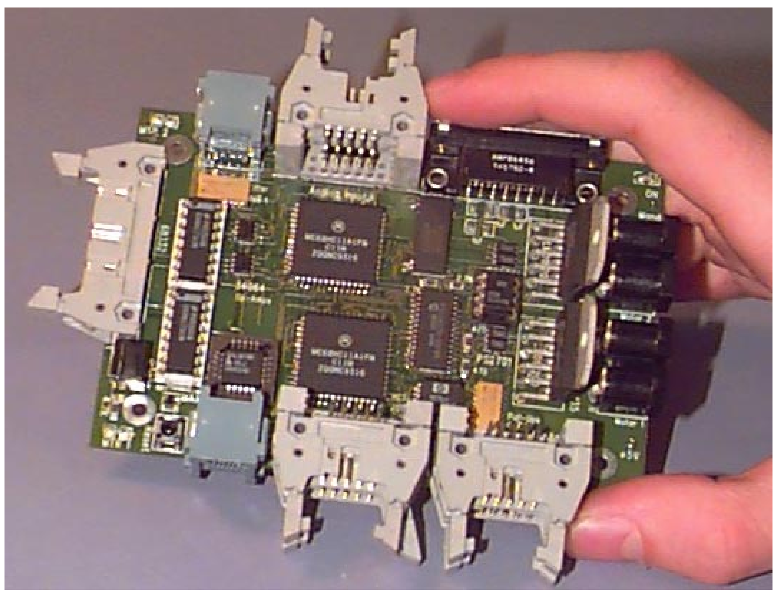

Figure A-26: Photograph of motor controller board. Each board contains all the circuitry to drive two series elastic actuators. Two Motorola 6811 microcontrollers implement the control law for the actuators, with extra circuitry to process the signals from the motor encoders, and generate PWM signals for the motor drives. A power amplifier is also included on board (at right hand end in this picture. The 6811's communicate with the higher level processors using a differentially driven serial line, which connects to the large connector at the left hand end.

The code for the 68332's is written and compiled on a Macintosh PowerPC. This computer acts as a Front End Processor (FEP) for the network. The serial line is used to communicate with the network, the serial communication being mediated by another 68332, dubbed the 'InterFEP'. This configuration is illustrated in figure A-28.

The DPRAM interface is used to connect to other parts of the robot. This is also shown in figure A-28. For example, for the drumming in chapter 7, the auditory signal was detected using a sound card in a PC running QNX, a real time UNIX. The auditory signal was communicated through the DPRAM to the 68332 network. Referring to figure A-27, the auditory signal entered the network through the 'POSTURE' processor, then was passed to the 'OSC' processors.

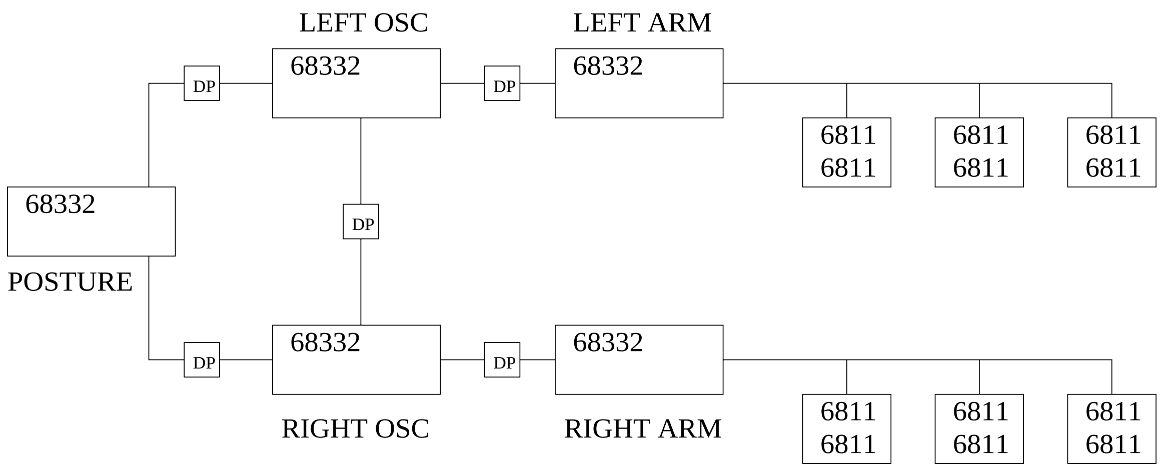

Figure A-27: The figure shows the network of Motorola 68332 processors used for higher level control, connected to the 6811 joint controllers. The connection between the '—ARM' processors and the 6811s is low bandwidth, with 13 byte packages exchanged at around 50Hz. The communication in the network is also slow, with dual ported RAMs (marked DP) used to share information between the processors. The 'POSTURE' processor was used to generate the desired posture of the arm, the '—OSC' processors integrated the oscillator equations, and the '—ARM' processors communicated with the joint level controllers.

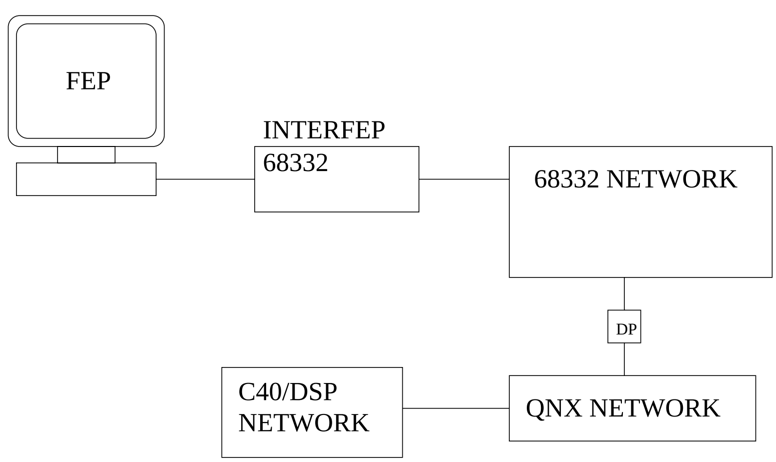

Figure A-28: The code for the 68332 network is written and compiled on the Front End Processor (FEP), which is a Macintosh PowerPC computer. The FEP communicates with the network using its serial port, the communication mediated by another 68332, the 'InterFEP'. The 68332 network can also receive input from other computers associated the robot. The robot vision system uses a network of C40 Digital Signal Processors (Scassellati, 1998), which are attached to a network of PC's running QNX, a real time UNIX operating system. These are patched into the 68332 network using the dual ported RAM interface.

A.10 Hand design

The arm design does not include a design for the hand. In practice it was found to be relatively easy to modify a simple paddle into a special purpose hand to grip cranks, hammers etc. without requiring the extra design, control and weight of a dextrous hand. Velcro was used extensively as a quick and surprisingly strong method of attaching objects to the hand of the robot.

A.11 Conclusion

This appendix has described the design of the arms used in this thesis. They are robust due to their mechanical design and to the choice of actuators. The control system is somewhat slow and low bandwidth, but given the stiffness of the arm this is not a difficulty. The low bandwidth and slow communication forces the control system implemented to be robust, which is what has been demonstrated in this thesis.

Appendix B

Arm and oscillator parameter tuning

This appendix covers some aspects of tuning of the oscillator parameters. Its describes the parameters which need to be set, the automatic tuning methods which were used, and a sample procedure for implementing a new rhythmic task using oscillators.

B.1 Parameters

For two arms with six degrees of freedom there are inevitably a large number of parameters that need to be set. However, most of them can be fixed and not altered in normal operation. The parameters are described below.

- **Force control gains.** These parameters control the performance of the force control in the series elastic actuators, described in detail in appendix A. The parameters can be calibrated once and then fixed. The calibration procedure involves tuning the PD control of force, usually by commanding a square wave of force, and observing the closed loop behavior. The calibration of the joint angles and velocities (described in appendix A) does not need to be particularly accurate, mostly because the control implemented on top (the oscillator system) does not require the arm to be accurately calibrated.

- **Arm stiffness/damping.** These parameters control the arm dynamic properties. There is an upper limit on the arm stiffness due to noise in the sensors, as well as bandwidth limitations of the force control. There is also an upper limit on damping, due to noise in the joint angle sensors. The stiffness and damping for each joint of the arm can be set to reasonable values and effectively fixed. Occasionally it is useful to alter the stiffness and damping of individual joints, in order to reduce or accentuate their motion.

- **Arm posture.** The arm posture is important for the overall system behavior, and can be easily set with one parameter per joint.

- **High pass filters.** There is a high pass filter on each oscillator input which removes the DC component from the signal. In fact a low pass filter is used to calculate the offset and this is subtracted from the input signal. A simple single pole digital filter was used with a cut-off frequency of 0.2 Hz. This was fixed for all the oscillator experiments.

- **Choice of oscillator input.** In the work for this thesis there are two possible choices of input parameter to the oscillators. The two signals used are joint angle θ_i and joint torque u_i. There is no reason to restrict the input signal to these parameters (for the jugging task in chapter 3 the ball trajectory was used, and auditory feedback was used for drumming in chapter 7. For any input signal, the describing function analysis described in chapter 3 is useful to determine the approximate effect of the feedback.

The describing function analysis can also be used to help choose between joint angle and joint torque. Williamson (1998b) and Williamson (1998a) describe the difference in some detail. Appendix D describes the effect of torque feedback on oscillators coupled through the natural dynamics (compared to the model in chapter 4).

The only major difference between the feedback types is that if torque feedback is used, it is possible to get a solution where the arm is not moving, but the oscillators are entrained. The oscillator outputs are entrained with internal forces. When the joint angle is used as input this obviously cannot occur. This means that although the torque feedback could be used for crank turning with one arm, it did not work with two arms. The torque feedback also appears to be more sensitive to mechanical coupling, being appropriate for the Slinky toy described in chapter 4 and appendix D.

- **Oscillator parameters.** The main parameters that need to be set are the tonic c in equation (3.1) which determines the amplitude of the joint level oscillation, and τ_1 which determines the oscillator frequency. The third parameter is the input gain which as discussed in chapter 3 is rather robust. As long as this gain is large enough to cause entrainment, its particular value does not appear to be important.

B.2 Automatic tuning

Since for most of the tasks described in this thesis, the system was set up to drive the resonant mode of the arm-task system, the posture of the arm and the amplitudes of the oscillator motions were the most important parameters to set. These have to be roughly correct to get the required motion, but their exact values are not important. The other oscillator parameters are easy to set, with their values being relatively unimportant.

The initial values of posture and amplitude can be set by hand to match a particular task, say for example a new crank turning configuration. These values can then be modified using an automatic tuning scheme, measuring the performance in the task, and adjusting the parameters accordingly.

The tightly coupled nature of the system makes most supervised learning techniques difficult to apply. It would be difficult to calculate an "error" and then relate that error to the parameters involved, especially given the different time-scales involved. The timing issue, as well as the rather large state space makes reinforcement learning techniques also not appropriate. The best learning techniques for this problem appear to be genetic algorithms, evolutionary strategies (Bäck and Schwefel, 1993) and hill climbing (e.g. Rich and Knight (1991)). Initial experiments on a two degree of freedom simulated arm suggested that the most effective strategy was hill climbing, and this was implemented to tune the parameters of the real arms.

The learning algorithm was then

1. Make a small random change to the parameter vector.
    - A number of strategies were tried here. The most successful was a technique called "run-and-twiddle". In this method the parameter vector is changed, and if the change causes an increase in performance, then further changes to the parameter vector are made in the same direction. This is continued until changing in that direction no longer gives an increase in performance, at which point a new direction is tested.
2. Apply these parameters and test the performance.
    - The performance evaluation generally included terms directly related to the task, as well as terms related to the energy required to produce the motion. This then optimized for

efficient motion. When using two arms the evaluation balanced the work between the two arms by including terms for the total energy, and the difference in energy between the two arms. The performance evaluation was sometimes noisy making the learning less robust. The noise can be reduced by measuring performance a number of times and averaging the results.

3. If the performance is better than ever before, keep the new parameter set and go to step 1. If not, revert to old parameter vector and go to step 1.

   - To prevent the learning being stuck in local minima, the new parameter setting was accepted if the performance was within some small percentage of the maximum performance. Simulated annealing could be used to apply this idea more rigorously, although the simpler technique appeared to work well.

It is indicative of the robustness of the oscillator parameters that a simple learning algorithm like this was successful for tuning. This was even the case when 10 or 12 parameters were tuned simultaneously. In the crank turning case (chapter 4), and also the hitting application (chapter 7) there appeared to be a variety of parameter setting that gave efficient performance, which the hill climbing algorithm was fairly quick to find.

B.3 Procedure for new task

The procedure of tuning the arm for a new task is described below, taking as an example the task of pumping a bicycle pump.

1. Connect pump to arm, either with hand or special attachment.

2. Manouever arm to approximately correct posture for the pumping motion.

3. Adjust oscillator time constants for approximately the correct frequency.

4. Adjust oscillator amplitudes to create motion.

5. Adjust feedback gains to cause entrainment of the system. At this point it is easy to determine which feedback type (angle or torque) works better.

6. If the system does not operate as intended: adjust posture, oscillator amplitudes, maybe stiffen unused joints, adjust arm stiffness and damping.

If the automatic tuning is then required to further optimize the system performance, then

1. Find way to measure performance of the task. In the bicycle pump example this might be a measure of the linear motion of the pump, or the airflow out the pump, as well as an estimate of the energy expended by the arm to perform the task.

2. Select the parameters which should be tuned. The most important parameters are probably the posture and the tonics (controlling the oscillator amplitudes), not only because these dictate whether motion will occur, but also the quality of the final motion (matched to task, too violent etc.).

3. Run hill climbing on the most important parameters.

Appendix C

Oscillator Inputs

This appendix considers the effect of using linear inputs to the oscillator instead of the usual non-linear inputs. Using linear inputs is not a good idea because they reduce the robustness of the system.

The oscillator used throughout this thesis is governed by the following equations:

$$\begin{aligned}
\tau_1 \dot{x}_1 &= c - x_1 - \beta v_1 - \gamma [x_2]^+ - \Sigma_j h_j [g_j]^+ && \text{(C.1)}\\
\tau_2 \dot{v}_1 &= [x_1]^+ - v_1 && \text{(C.2)}\\
\tau_1 \dot{x}_2 &= c - x_2 - \beta v_2 - \gamma [x_1]^+ - \Sigma_j h_j [g_j]^- && \text{(C.3)}\\
\tau_2 \dot{v}_2 &= [x_2]^+ - v_2 && \text{(C.4)}\\
y_i &= [x_i]^+ = \max(x_i, 0) && \text{(C.5)}\\
y_{out} &= [x_1]^+ - [x_2]^+ = y_1 - y_2 && \text{(C.6)}
\end{aligned}$$

The input is applied through a max operator, so that the input inhibits the equation. For example, when the input g_j is positive, its only effect is in the x_1 equation, making $\dot{x}_1$ more negative. The input can be applied without this non-linearity, replacing $\Sigma_j h_j [g_j]^+$ by $\Sigma_j h_j g_j$ and $\Sigma_j h_j [g_j]^-$ by $-\Sigma_j h_j g_j$.

The major effect of not using the non-linearity on the oscillator input is that the oscillator output amplitude becomes dependent on the size of the input. This is illustrated in figure C-1, showing the output amplitude of the oscillator with and without non-linear inputs. This affects the Bode plot for the oscillator (shown in figure C-2), making the gains higher, and with a smaller range in gains than for the non-linear input case. The phases for the two cases are approximately the same.

By reducing the variety of gains possible from the oscillator, using linear inputs reduces the robustness of the system. In addition, it couples together the parameters which control input gain (h_j) with those controlling output amplitude (c), so making the tuning more complicated. Since the output is not dependent on the input, the stability results discussed in chapter 6 do not hold in their present form. A more rigorous determination of the H_∞ norm of the oscillator is required before the behaviors of systems coupled to an oscillator using linear inputs can be predicted.

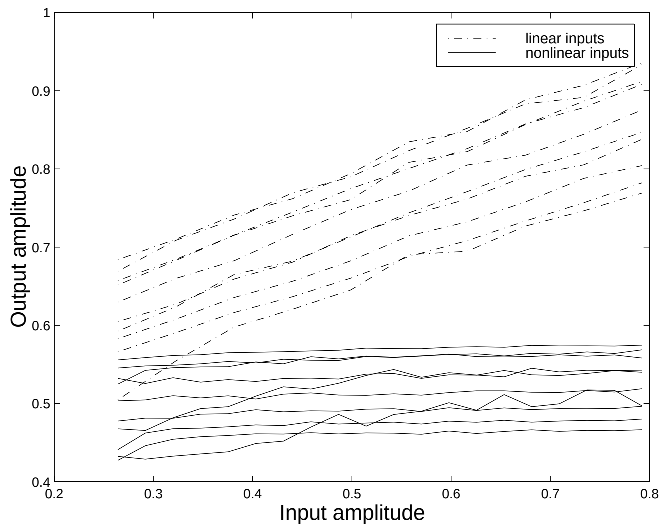

Figure C-1: SIM Plot showing effect of using non-linearity on oscillator input on the oscillator output amplitude. When the non-linearity is used, the oscillator output is approximately constant (solid lines, where the multiple lines correspond to different frequencies). When the inputs are applied in a linear fashion, the output amplitude is approximately proportional to the input amplitude (dash-dot lines).

Linear Inputs

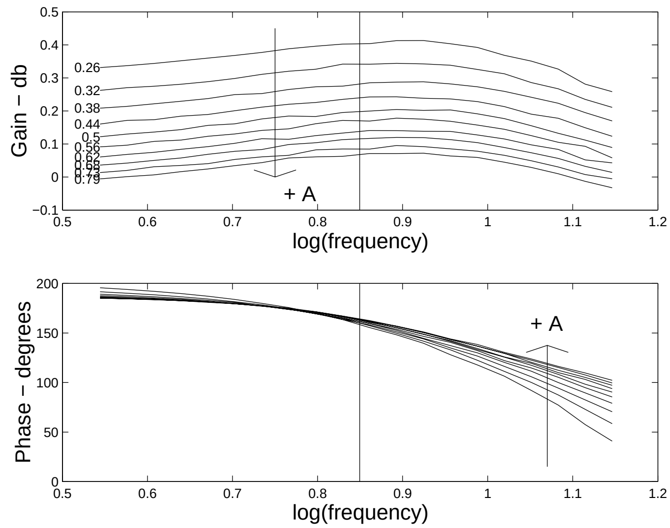

Non-linear Inputs

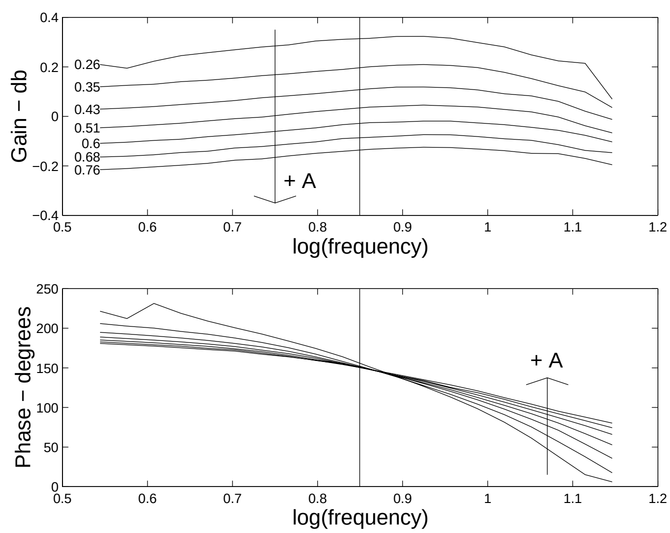

Figure C-2: SIM Describing function of oscillator with and without linear inputs. The top plots show the Bode plot of the oscillator with linear inputs, and the lower plots with non-linear inputs. When linear inputs are used, the gain of the oscillator is higher, and covers a lower range of gains than with the non-linearity. This reduces the robustness of the oscillator to changing gains and system parameters. The phases in both cases are similar.

Appendix D

Oscillators and torque feedback

This appendix describes how the simple model of coupling introduced in chapter 4 is modified when the joint torque, rather than the joint angle is used as the feedback signal to the oscillators. The results are the same with the different feedback, the oscillators finding the resonant mode of the underlying mechanical system. The use of this feedback is illustrated with the example of coordinating two arms of the robot to pass a Slinky toy from hand to hand.

D.1 Torque feedback

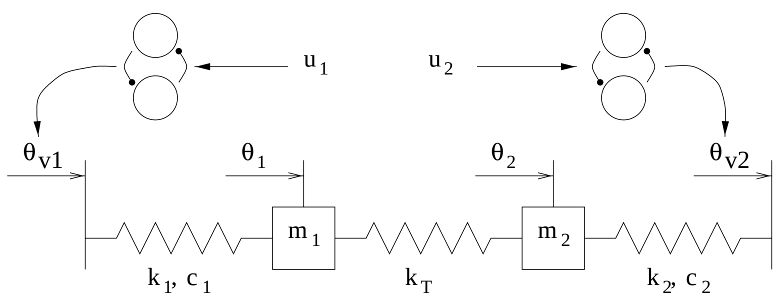

Figure D-1: A simple model of coupling through the natural dynamics. The model consists of two masses driven by oscillators, connected by a coupling spring k_T.

The model of coupling when force or torque feedback is as shown in figure D-1. The only difference is that instead of using the joint angle θ, the oscillator uses the joint torque $u = k(\theta_v - \theta) - b\dot{\theta}$ as the feedback signal. Since each joint is effectively a mass-spring system, the joint torque is related to the joint angle:

$$u = k(\theta_v - \theta) - b\dot{\theta} = m\ddot{\theta}$$

Using the describing function analysis, the effect of the oscillator can be written:

$$\theta_v = ge^{j\phi}u = ge^{j\phi}(-m\omega^2)\theta \tag{D.1}$$

where m is the mass of the system. This alters the equation of the example coupled system to:

$$\begin{pmatrix} k_1 + k_T + j\omega c_1 - m_1\omega^2 & -k_T \\ -k_T & k_2 + k_T + j\omega c_2 - m_1\omega^2 \end{pmatrix} \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} = \begin{pmatrix} -m_1\omega^2 k_1 g_1 e^{j\phi_1}\Theta_1 \\ -m_2\omega^2 k_2 g_2 e^{j\phi_2}\Theta_2 \end{pmatrix} \tag{D.2}$$

For steady state solutions, the damping terms must be zero:

$$\omega c_1 + m_1\omega^2 k_1 \mathrm{Im}[g_1 e^{j\phi_1}] = \omega c_2 + m_2\omega^2 k_2 \mathrm{Im}[g_2 e^{j\phi_2}] = 0 \tag{D.3}$$

which implies that the imaginary parts of $g_i e^{j\phi_i}$ must be negative, and vary with $1/\omega$. Removing the damping terms results in the eigenvalue problem:

$$\begin{pmatrix} \frac{k_1+k_T}{m_1(1-\mathrm{Re}[g_1 e^{j\phi_1}])} & \frac{-k_T}{m_1(1-\mathrm{Re}[g_1 e^{j\phi_1}])} \\ \frac{-k_T}{m_2(1-\mathrm{Re}[g_2 e^{j\phi_2}])} & \frac{k_2+k_T}{m_2(1-\mathrm{Re}[g_2 e^{j\phi_2}])} \end{pmatrix} \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} = \omega^2 \begin{pmatrix} \Theta_1 \\ \Theta_2 \end{pmatrix} \tag{D.4}$$

The torque feedback makes the effect of the oscillator on the coupled system dynamics more apparent than in the angle feedback case. The effect of the oscillator is to alter the apparent inertia in the system, changing the effect of mass m_1 to $m_1(1 - \mathrm{Re}[g_1 e^{j\phi_1}])$, and similarly for mass m_2. Because of the oscillator properties, the values of phase for both oscillators will be similar, so $(1-\mathrm{Re}[g_1 e^{j\phi_1}]) \approx (1-\mathrm{Re}[g_2 e^{j\phi_2}])$. This results in the effect of the oscillator being a common factor in (D.4), so only effecting the eigenvalues (the final frequency), not the eigenvectors (the mode shape). The solutions to the damping constraint (D.3) have imaginary parts which are negative, which due to the oscillator properties implies that the real parts will also be negative. This means that the oscillator increases the apparent inertia of the system i.e. $(1 - \mathrm{Re}[ge^{j\phi}]) > 1$.

The oscillator solutions are also locally stable using this feedback method, which is seen by looking at the effect of amplitude changes on the damping in the system (see figure D-2).

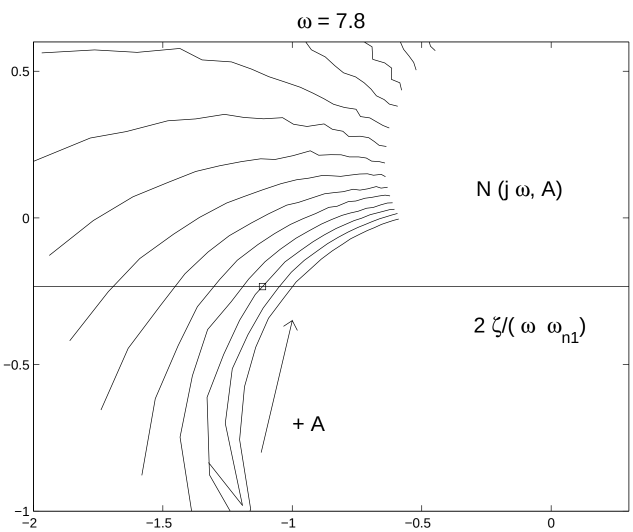

Figure D-2: SIM Local stability under torque feedback. The plot shows $N(j\omega, A)$ and the line $j2\zeta/(\omega\omega_n) = jc_1/(m_1 k_1 \omega)$ plotted on the complex plane, with a sample limit cycle marked with a $\square$. The limit cycle can be shown to be locally stable by examining changes in the system damping as the amplitude of oscillation A is varied. An increase in amplitude (in the direction of the arrow) results in $N(j\omega, A) > 2\zeta/(\omega\omega_n)$, which is equivalent to an increase in damping by (D.3), and so a reduction in the amplitude. A reduction in amplitude has the opposite effect, showing that the limit cycle is locally stable.

D.2 Slinky example

The task chosen is the coordination of two arms of the robot using a Slinky toy. This is a good example because it illustrates that the oscillators are sensitive to weak coupling as opposed to the strong coupling in the crank turning case. As the Slinky is passed from hand to hand the weight on the hand varies, with the difference in weight being greatest when the hands are furthest apart. This behavior could be approximated by a coupling spring which is weak and has a negative spring constant, giving a larger force away from the central position. This can then be approximated by the coupling model.

When the Slinky is placed between the two arms of the robot as shown in figure D-3, and the joint torque is used as the feedback signal to the oscillators, the oscillators quickly find the anti-phase motion. An in-phase motion is sometimes observed, but it does not appear to be a stable solution, as small perturbations result in the out of phase motion. A sample time trace for the motion of the two arms is illustrated in figure D-4, showing the effect of the torque feedback in coordinating the two arms.

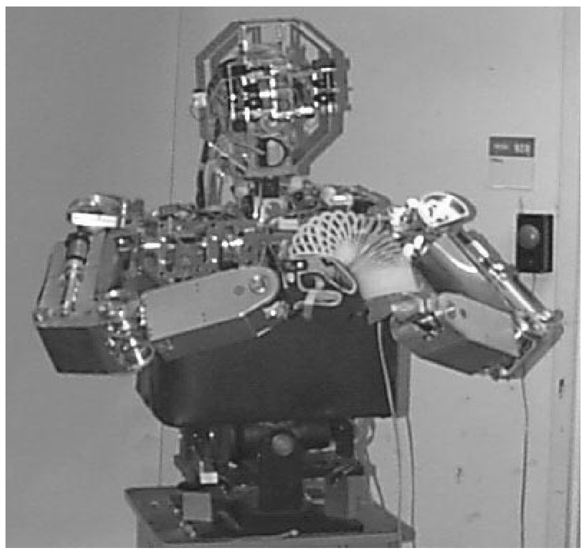

Figure D-3: Picture of Cog passing the Slinky toy from hand to hand. The two elbow joints are used to move the hands up and down, where the coordination between the hands is given by the interaction between the oscillator dynamics and the coupled arm-slinky system.

Data from the Slinky was taken, recording the frequencies, amplitudes and phases of the joint motions as the natural frequency of the oscillator was varied. This data could then be used to directly calculate the gain of the oscillators, and the mode of the Slinky. The process followed is similar to that for the crank turning example described in chapter 4.

The damping constraint in the coupling model predicts that the imaginary part of the oscillator action $\mathrm{Im}[g \exp j\phi] \propto 1/\omega$, which is plotted in figure D-5. The constant of proportionality should be $-c/(km)$, making the slope negative. The data lies approximately on a straight line, with data at very slow speeds not fitting on the line. The slope is negative as predicted by the theory.

A plot of the real parts of the oscillator actions is included in figure D-6. The model predicts that the real parts should be proportional to $m - k_T/\omega^2$. The graph shows the real part of the oscillator action plotted against $1/\omega^2$. The data falls on a positive slope line at high speeds, indicating that the coupling stiffness k_T is negative. At low frequencies the points do not lie on the line but appear to fan out. This may be due to a change in the coupling type of the Slinky. At high speeds the spring-like properties of the Slinky complement the transfer of mass from arm to arm, while at low frequencies the coupling will only come from the mass. This may reduce the strength of the coupling, giving rise to the points not lying on the same line.

These results show that the Slinky operation is a resonant mode. The oscillator quickly responds to the resonant properties of the system and tunes automatically into the mode shape, which is

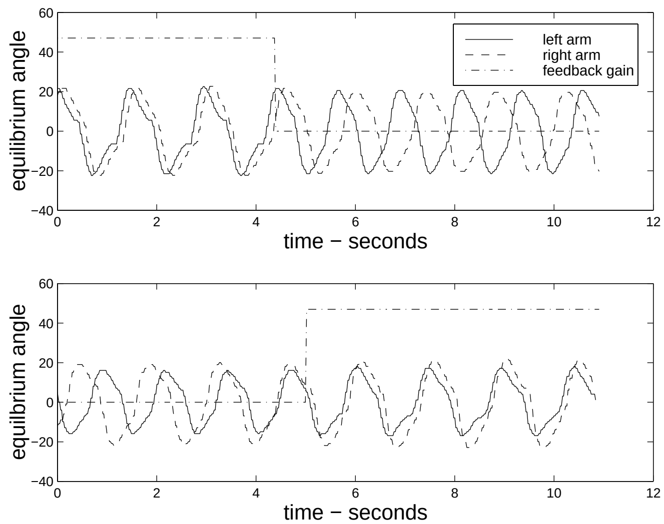

Figure D-4: REAL Two examples of Slinky operation. Both plots show the outputs from the oscillators as the torque feedback (dash-dot) is turned on and off. When the traces are in phase, the Slinky is moving in anti-phase. When the feedback is on, the two arms are coordinated and the outputs are synchronized,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

passing the Slinky from hand to hand. The behavior is also robust as shown in figure D-7, the oscillators finding the correct coordination even when the oscillator time constant is doubled.

D.3 Conclusion

This appendix has shown that the effect of torque feedback on the oscillators is somewhat similar to joint angle feedback in multi-joint motions. In both cases, the oscillator finds the resonant mode of the system. The effect of the torque feedback is to increase the inertia of the system, as opposed to alter the stiffness in the angle case. This means that the overall motion is often slower, and the mode found by the oscillators is often the slower one in these cases. In addition, because the input to the oscillator is the joint torque, it is possible to have a solution where the arm is completely still, and the oscillators are pushing against one another, entrained with the internal forces.

The example with the Slinky toy showed that small forces can be used to entrain the oscillators, and that the analysis from the model is useful in understanding the nature of the coupling. The robustness of the system to parameter changes was also shown. The sensitivity of the oscillators to the resonant mode makes them useful for tasks where the mechanical coupling is weak as well as strong.

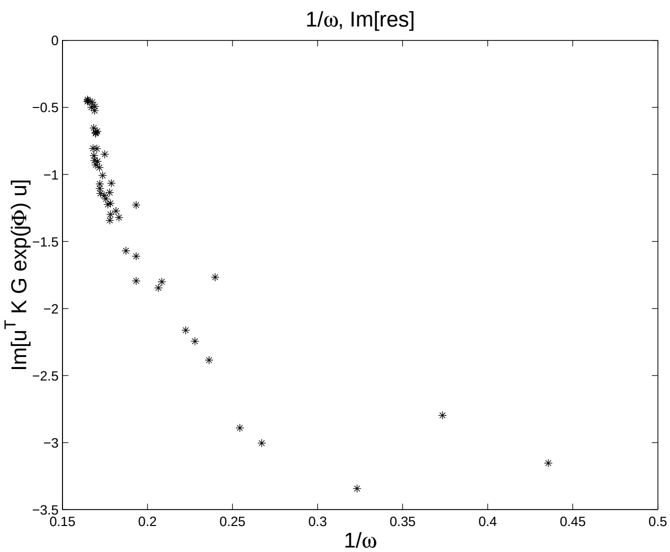

Figure D-5: REAL Plot of $\mathrm{Im}[u_i^T K' G e^{j\Phi} u_i]$ against $1/\omega$ for the Slinky. The model predicts this to be a linear relation with slope $-c/(km)$. The data confirms this prediction, lying approximately on a line of negative slope. At low frequencies (high values of $1/\omega$), the data does not lie on the same line.

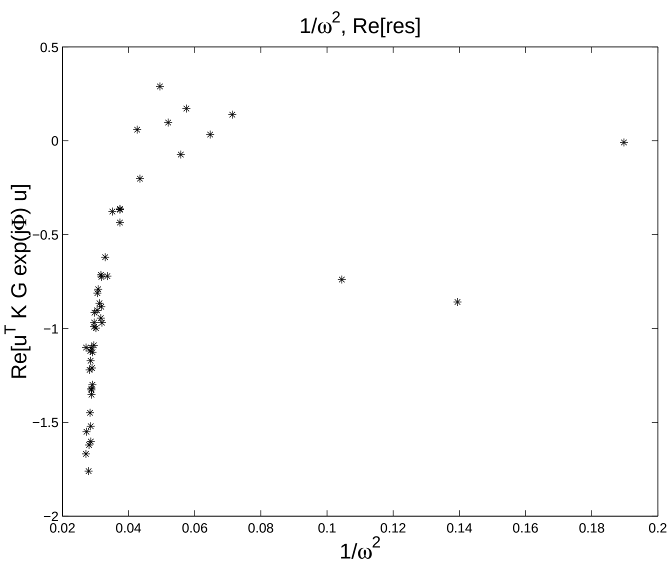

Figure D-6: REAL Plot of $\mathrm{Re}[u_i^T K' G e^{j\Phi} u_i]$ against $1/\omega^2$ for the Slinky. The model predicts that this relation should have a slope $-k_T$. If the coupling from the Slinky is equivalent to a negative spring, then the slopes of the data should be positive. This is shown in the plot. At low frequencies the data does not lie on the same line, perhaps indicating that the type of coupling is changing to one dominated by the transfer of the mass of the Slinky from arm to arm.

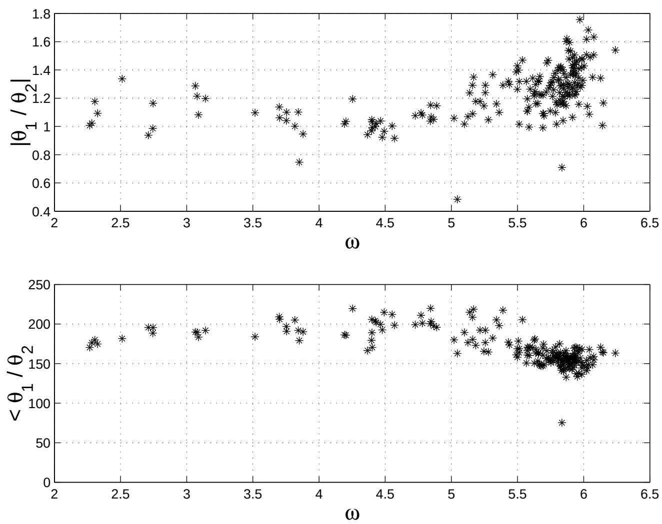

Figure D-7: REAL Mode of Slinky operation. The top plot shows the amplitude θ_1/θ_2 of the Slinky motion versus frequency as the natural frequency of the oscillator is changed. The lower plot shows the phase between the two hands. Even though the frequency of the system approximately doubles over this plot, the mode remains approximately constant. This indicates not only the robustness of the oscillator solution, but also its ability to find the mode of the system as its final periodic solution.