A-Level Electronics ELEC4 Notes

Milo Noblet

2 June 2017

1 Control systems

- Open loop — simply carries out instructions given
- Closed loop — monitors its own progress in how well it is carrying out the instructions
  - A generalised closed loop system has an output sensor and an error detector added (the error detector is often a part of the processor)
- Negative feedback — the processor adjusts the output device such that the difference between the output and instructions is reduced/eliminated.
- Positive feedback increases the difference between the output and instructions — sending the system into saturation.

The output sensor can only sense the instantaneous output — due to latency in the control system, by the time the system has reacted, a mechanical system will have moved on. This can result in the control system *'hunting'*, causing considerable wear to mechanical parts. This issue can be solved through *hysteresis*, where there are two separate switching levels.

1.1 Schmitt triggers

A Schmitt trigger exhibits hyteresis — it has a different 'switch on' voltage to its 'switch off' voltage. They use positive feedback to increase their switching speed and define specific voltage levels at which switching will occur.

An inverting Schmitt trigger will switch low when the input signal is greater than the upper switching level, and then switch high when the input signal is less than the lower switching level.

1.1.1 Two-resistor inverting Schmitt trigger

Treat R_f and R_1 as a potential divider, and consider the minimum and maximum outputs that the op-amp can have (looking at its power supply). Apply these voltages to the potential divider formula and this will find the two switching levels — naturally, the upper switching level will occur when the op-amp outputs is at the maximum.

The two-resistor inverting Schmitt trigger is only really suitable when the op-amp has a positive and negative supply.

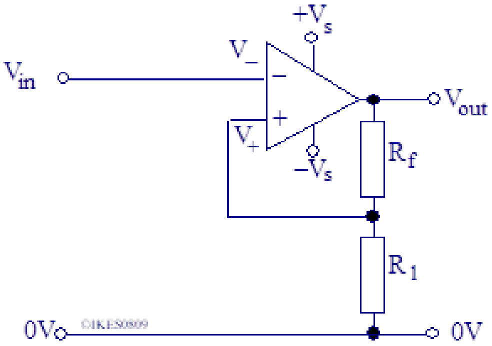

Figure 1: Inverting Schmitt trigger with two resistors

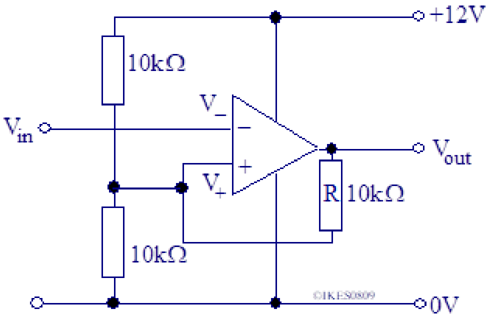

Figure 2: Inverting Schmitt trigger with three resistors

1.1.2 Three-resistor inverting Schmitt trigger

The three-resistor trigger is used when the op-amp has just a positive supply and 0 V. The switching levels can be found by considering when the feedback resistor R is effectively in parallel with either the top or bottom resistor of the potential divider.

The upper switching level occurs when the feedback resistor is in parallel with the top resistor, and the lower switching level when it is in parallel with the bottom.

2 Microprocessors

2.1 Merits

- Cost — microprocessors allow much simpler and smaller circuits, with a reduced component count
- Reliability is generally higher with a microprocessor-based system as there are fewer discrete components that could fail
- Testing is usually more expensive with a hard-wired system than software, though this may not be the case for very complex software systems
- It is much easier to modify the operation of a microprocessor system — one can simply re-flash the ROM rather than having to design a new circuit board. Often the same circuit board can be used for a range of applications, further reducing costs.

2.2 Disadvantages

- Speed of operation — a hard-wired system will usually be faster, and does not have the initialisation delay of a microprocessor
- More electrical noise is generated as a result of microprocessors having many internal gates switching continuously

2.3 Architecture

- von Neumann — uses a single bus along which both program instructions and data are fetched. Used in microprocessors.
- Harvard — used in system-on-chip microcontroller systems, has a separate bus for instructions rather than using the data bus, which allows for ‘pipelining’ — speeding up operation.

2.3.1 Subsystems

- Read Only Memory (ROM) contains the instructions the processor is to follow. It is non-volatile, meaning the information stored remains once power is lost.
- Random Access Memory (RAM) is where the processor can store/write information & read information while working — it is temporary and so volatile: information is lost when power is removed.
- The processor executes instructions stored in the ROM and makes decisions

- The clock keeps the processor and the rest of the system synchronised — it is essential for reliability and ensuring data is not corrupted as it is sent along buses
- Buses are sets of connections between the subsystems of a control system, and are used to prevent a proliferation of connections:
    - Data bus: carries information around the system, is bi-directional
    - Address bus: for the sending of memory and/or port addresses. It is one-way.
    - An instruction bus is only found in Harvard architectures
        - note the data, address & instruction buses are multi-bit signals sent in parallel. In some systems the data & address buses are shared.
    - The control bus is often one-bit signals, with a specific line for each control signal.

2.3.2 Tristate buffers

Tristate buffers are used to prevent bus contention, which occurs when more than one device attempts to put information on the data bus at a time. Tristate refers to having three states: input, output and high impedance (ie not connected). When the I/O device is not selected it just floats with the signals on the data bus and does not interfere with the data bus signals. The tristate buffers are controlled by the address bus and control lines — they have an output enable $\overline{\text{OE}}$ pin, which makes the output high impedance when $\overline{\text{OE}}$ is high.

- There is usually also an R/$\overline{\text{W}}$ line for reading/writing from memory or a port
- $\overline{\text{IRQ}}$ — interrupt request, used to get the processor to deal with another task.
    1. $\overline{\text{IRQ}}$ taken to logic 0
    2. processor finishes current execution
    3. saves current values in registers to the 'stack'
    4. sets the stack pointer
    5. identify which device sent the request, and then service the interrupting device
    6. once complete, the processor returns to the task it was doing before the interrupt occurred, loading information back from the stack.

    - $\overline{\text{IRQ}}$s are maskable interrupts because they can be disabled by the programmer
    - Conversely, $\overline{\text{NMI}}$ refers to a non-maskable interrupt

Addresses Every memory location and port must have a unique address — the number of unique addresses, and thus the number of bytes of memory that can be used, is determined by the width of the address bus: 2^n addresses where n is the number of parallel bits in the address bus.

2.3.3 Interfacing ports

- Memory mapping is easy to implement: the ports are treated just like any other memory location. It has the disadvantage of restricting the amount of RAM and ROM which can be connected to the microprocessor as it takes up addresses.
- I/O mapping requires the microprocessor to have the additional M/$\overline{\text{IO}}$ control line which determines whether the address on the address bus refers to a memory location or an I/O port.
    - The M/$\overline{\text{IO}}$ signal has to be included in the decoding for every address so I/O mapping is more complex to implement, but has the advantage that it does not restrict the amount of RAM/ROM that can be connected to the microprocessor.

2.4 Terms

- registers — data stored here before and after processing
- accumulators — special registers for the result of any arithmetic or logic process, eg the 'working register'
- zero flag — set to 1 if the result of the last operation was 0
- sign flag — set to 1 if the contents of the accumulator is negative

- carry flag — set if the maximum number has been exceeded (eg trying to put 256 in an 8-bit register)
- ALU — arithmetic & logic unit
- program counter — contains the address of the next instruction
- instruction decoder — interprets program instructions into a set of electronic conditions which make the data move appropriately
- stack — a temporary store for data, grows in size as data is added, shrinking as data is removed. It is used sequentially: last in, first out.

2.5 AQA Assembler

- Tristate registers are associated with ports eg `TRISA`, `TRISB`. If a bit is set to 1 then it is an input, 0 = output.
- Status register, **SR**:
  - Bit D_0 is the carry flag, `C`
  - D_1 is set to 1 when TMR register is 0
  - D_2 is the zero flag, `Z`
- Clock prescaler, **PRE**. The clock frequency is divided by PRE + 1. If PRE = 0 then the timing function is disabled.
- 8-bit timer register, **TMR**, which is decremented on the falling edge of the PRE pulse and sets D_1 of SR when it is 0.

- `CALL K` — subroutines are a way of making frequently needed blocks of code available anywhere within a program. `CALL` transfers execution to subroutine `K`. Takes 2 clock cycles to execute as has to increment PC (progam counter) & store it into the stack, decrement SP (stack pointer) then load the address of the subroutine into PC.
- `RET` — ends a subroutine and transfers execution back to the main program; takes 2 clock cycles. Last value in the stack is loaded into PC.
- `INC R` or `DEC R` — increment or decrement register `R`
- `ANDW K`, `ADDW K`, `SUBW K`, `ORW K`, `XORW K` are used for arithmetic or logic operations on the working register `W` with literal `K`
- `JMP K` — go to label `K`. Differs from `CALL` in that no return address is stored, but still takes 2 clock cycles as address loaded into PC.
  - `JPC K` — change the location of the next instruction to be executed only if the carry flag, `C`, is set.
  - `JPZ K` — as with `JPC`, but only if the zero flag, `Z`, is set.
- `MOVWR R` — move the contents of working register `W` into register `R`.
- `MOVW K` — move the value `K` into the working register `W`
- `MOVRW R` — move the contents of register `R` into the working register `W`

Masking bits If we need to know the value of bit D_1 only, we would `AND` with $0000\ 0010_2$, turning all bits except D_1 to 0.

Setting a bit `XOR`-ing will set a bit that is 0 to 1, and will flip a bit that is already 1 so that it is 0.
`OR`-ing will set bits, ignoring their existing value (so bits that are already 1 will stay 1)

2.5.1 Generating a time delay

For short delays, a string of `NOP` ('no operation') commands can be used, each causing a delay of one clock pulse. But for time delays longer than a few μs it is not an efficient method.

Instead we make us of `PRE` and `TMR`, the operation of which is independent of the microprocessor's current instruction (so the processor can be performing other tasks while waiting for the delay to conclude). The value in `TMR` decreases by 1 (decrements) on each falling edge of the pulses from `PRE`

- Example: A delay of 2 ms is needed from a processor with clock speed 1 MHz (so one pulse = 1 μs).
  - Loading `PRE` with 7 divides clock frequency by 8, giving period 8 μs.
  - `TMR` of 250 decreases once every 8 μs and $250 \times 8\ \mu\text{s} = 2$ ms. So will reach 0 after 2 ms then set bit D_1 of `SR`
  - To detect `TMR` reaching zero, need to poll D_1 of `SR` — achieved by moving `SR` into `W`, `AND`-ing with 0x02 to mask all but D_1, then `JPZ` to start. When the delay has finished, $D_1 = 1$ and so the jump will not occur.

2.5.2 Dealing with inputs

The most efficient way of monitoring inputs is to use hardware interrupts: when an I/O device needs attention it takes $\overline{\text{IRQ}}$ low, causing the processor to call a subroutine to deal with the change. To prevent the processor being overwhelmed with requests each I/O device is usually given a priority.

The other way of dealing with devices is polling, where the processor continually checks the state of the port to which the device is connected at regular intervals. This is much simpler to implement, but does not make efficient use of the processor's time.

2.6 System on Chip

Both PIC & AVR microcontrollers are based on the Harvard architecture and use RISC (reduced instruction set computer) processors. Separate buses for data and instructions allow 'pipelining' — while one instruction is being executed the next is pre-fetched from the program memory, allowing instructions to be executed every clock cycle. These microcontrollers have very low power consumption, as well as sleep modes, making them suitable for use in battery-powered equipment.

SoC gives socioeconomic benefits — it has led to the use of engine management systems and safety features such as ABS, although repair and maintenance has become more specialised as software proliferates. Cheaper computer peripherals are now available as less precise mechanisms can be used to reduce cost with a microcontroller ensuring the same output quality.

It can increasingly be difficult to ensure software is reliable due to its increasing complexity.

3 Handling analogue & digital signals

3.1 Ramp ADC

- Uses a digitally generated linear ramp from a microcontroller or binary counter which is converting to an analogue voltage by a digital—analogue converter (DAC).
  - Analogue ramp voltage is then compared with the input signal using a comparator.
  - When $V_{\text{ramp}} \geq V_{\text{in}}$ the comparator switches and through logic the counter is stopped.
  - The digital value of V_{in} can then be read from the counter output.

The issue with ramp ADCs is that they are slow — they are limited by the conversion time of the DAC. For a voltage close to the maximum, $2^n - 1$ operations may be needed (where n is the number of bits).

During the time taken for conversion, the input signal may be changing, which could lead to serious conversion errors — a 'sample and hold' circuit could be used to store the instantaneous voltage of the analogue signal, keeping the input voltage constant during conversion. It is for this reason that a data latch may be seen in some ramp ADCs.

3.2 Flash ADC

Flash ADCs can be used for very fast conversion but are expensive — they require one comparator for every possible binary state eg for an 8-bit flash ADC 255 identical comparators would be needed ($2^8 - 1$).

Multi-bit flash ADCs are very expensive because of the requirement for many comparators and identical resistors to give voltage references. Electrical noise is also a major issue in ADCs since the logic circuit to encode to binary generates transient signals as the gates switch — need to keep the input away from the output and make correct use of separate analogue & digital grounds.

Flash ADCs also have a higher power consumption than ramp ADCs due to the many comparators.

3.3 Digital—analogue conversion

The simplest DAC is based on a summing amplifier, but the inputs are weighted based on the binary digital they represent ie for the most significant bit the amp may have unity gain, so $R_{\text{in}} = R_{\text{f}}$, but for the next bit the weighting needs to be half, so $R_{\text{in}} = 2R_{\text{f}}$. This setup continues, with R_{in} doubling with each bit closer to the least significant bit.

The summing amplifier circuit is inverting, so a unity gain inverting amplifier may be used to make the output positive again.

4 Displays

- It is generally more efficient to multiplex 7-segment displays where more than one is needed, to reduce the component count.
- Each display is powered in turn and the rest of the time it is switched off.
- By switching between displays quickly enough, the persistence of vision of the eye ($\approx$ 50 ms) make sit appear as if all the displays are lit.
- If 4 displays are multiplexed, then each display is only powered for 25% of the time and so appears 25% as bright as it would usually.
  - It is excess heat that damages displays, so if the overall power dissipation of each segment over one complete display cycle remains the same the segments will not be damaged.
  - This means that the brightness issue can be overcome by passing 4 times the current through the display.
    - Note that this is why multiplexing would not give a power saving — although the reduction in driver/decoder ICs needed may help.
- LED displays have two main disadvantages: they consume significant power and are difficult to read in direct sunlight. However, they can be seen in the dark, switch quickly and are physically robust.
  - Liquid crystal displays (LCDs) are the complete opposite of LED displays, although they do offer greater choice in the characters that can be displayed — custom symbols can be added in the manufacturing process.
  - More complex to drive an LCD due to the need to have a square driving signal in antiphase with the backplane
- Another form of display is the LED dot matrix, which offers considerable flexibility in the size and shape of characters displayed, but suffers from similar disadvantages to the multiplexed LED display.
  - The scanning process by which each column (or row) is activated in turn needs to be fast and continuous, which can take up significant processing time from the control system.
  - For a five column display, the LEDs will only be lit for $\frac{1}{5}$ the time, so will appear $\frac{1}{5}$ as bright — pass $5\times$ the current to fix.
  - When choosing whether to put protecting resistors in the columns or rows, consider how the display is going to be switched. If the columns need to be taken low and are scanned, then only one LED per row is ever on — so putting the resistors in the rows means there will not be issues such as varying brightness as more LEDs are lit.

5 Motors

5.1 Shaft encoders

Shaft encoders are used with opto-switches (consisting of a light emitting diode and a light sensitive device such as a photodiode, phototransistor or light-dependent resistor).

5.1.1 Photodiodes

A photodiode takes advantage of the fact that the P—N junctions in diodes & transistors are light sensitive. Incident photons cause electron—hole pairs to be formed, producing a current proportional to the light intensity. A photodiode has a transparent window and is usually used in reverse bias, giving a leakage current which increases with light intensity.

The current produced by a photodiode is very small, so to get voltages great enough to measure the interfacing circuits must have very high resistance — when used with an op-amp the feedback resistor may by around 10 MΩ. In photodiode interfacing circuits, if the diode is connected between 0 V and the inverting input of an op-amp, current flows in the direction of the diode and through the feedback resistor to the op-amp output. If the photodiode is connected between −V and the inverting input, current flows in reverse through the photodiode.

5.1.2 Slotted disk shaft encoder

A slotted disk shaft encoder can be used with either slotted or reflective opto-switches (though if marks on a solid disk are used then only reflective opto-switches are an option).

Such a shaft encoder can only be used to determine how far the disk has rotated, but not the direction of rotation — to determine direction it is necessary to use two optical switches slightly offset from one another.

5.1.3 Optical shaft encoders

A binary coded optical shaft encoder can be used to determine the position of the disk at any point as well as the direction of movement. The least significant bit is given by the outer ring (changes most often), and the most significant bit by the inner.

A major issue with binary coded encoders is that where more than one bit changes simultaneously errors may occur if all the opto-switches are not perfectly aligned and do not operate in sync. This leads to transient false information as the disk rotates.

Gray code Using a Gray code on an optical shaft encoder overcomes this issue as only one bit changes at a time.

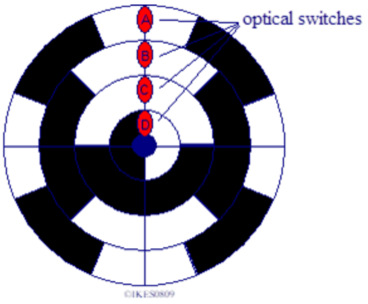

Figure 3: Gray-coded shaft encoder

5.2 Stepper motors

A stepper motor is a DC motor whose armature moves through a discrete angle (‘step’) each time power is applied to the relevant coils. By energising each coil or pair of in sequence, the stepper motor will rotate, one step at a time.

Apart from the accurate positioning of the shaft, stepper motors have the advantage that they hold the shaft firmly in place while stopped. However, stepper motors are far less efficient than conventional motors.

5.2.1 Variable reluctance stepper motor

This motor consists of a soft iron multi-tooth rotor and three coils (stator). When the stator windings are energised the stator poles become magnetised, and rotation occurs when the rotor teeth are attracted to the energised stator poles.

The variable reluctance stepper is an example of a unipolar motor, since the current only flows in one direction through the coils.

5.2.2 Permanent magnet stepper motor

The permanent magnet stepper is a low cost and low resolution motor. The rotor is magnetised with alternating north and south poles in a straight line parallel to the shaft — this provides an increased magnetic flux density and so the motor exhibits improved torque compared to the variable reluctance motor.

The coils in such a motor are centre-tapped, but the motor is still unipolar as current only flows in one direction through each coil (the centre-tapping effectively gives two sets of windings).

5.2.3 Hybrid stepper motor

A hybrid stepper motor is more expensive but provides better performance with respect to resolution, torque and speed. The rotor is multi-toothed as in a variable reluctance motor and contains an axially magnetised concentric magnet around its shaft (giving the increased flux density similar to a permanent magnet motor).

Here, motor coils are also centre tapped, and the motor is still said to be unipolar.

5.2.4 Bipolar stepper motors

Bipolar permanent magnet and hybrid motors have the same mechanism as their unipolar counterparts, but the two windings are simpler with no centre tapping.

This means the motor itself is simpler, but the drive circuitry is more complex as it needs to be able to reverse the current through each coil. As current flows through the coils in both directions, a bipolar motor is also two-phase.

The drive circuitry for a bipolar motor requires a H-bridge for each coil which allows the current through each winding to be reversed independently.

5.2.5 Comparison of conventional and stepper motors

- Speed:
  - Up to 30 000 rpm for conventional motors
  - Up to a few hundred rpm for stepper motors
- Torque:
  - Conventional motors have a high starting torque which is reasonable constant over a range of speeds, but have no stationary holding torque.
    - Note that there is some braking effect if a conventional motor is shorted out due to Lenz's law
  - Stepper motors start with a high torque but it drops rapidly with speed of rotation. Will hold the shaft in place when stopped.
- Connections:
  - Conventional motors usually have two connections, sometimes four if the field coils are separate
  - Four to six connections for stepper motors, can be more depending on number of field coils
- Efficiency:
  - Conventional motors — as high as 90%
  - Stepper — low, 10-40% as most power dissipated as heat.
- Control
  - Much easier to control conventional motor — on/off.
  - Stepper motor needs complex electronic circuit.
- Accuracy
  - No accuracy with a conventional motor unless used with a servo arrangement or a shaft encoder.
  - Stepper motors have an accuracy of 3-5% of a step on the overall angle turned through.

5.3 H-bridges

The H-bridge driver allows the current flow to be reversed through a device. It consists of p-channel MOSFETs on the top, whose gates need to be taken low for them to conduct, and n-channel MOSFETs on the bottom, whose gates are taken high to allow for conduction.

The drains of the MOSFETs are connected together in the centre rather than the sources so that they function as switches instead of source-follower amplifiers.

Note that protection diodes are needed so that the high 'back' emf which comes when switching an inductive load does not damage the MOSFETs.

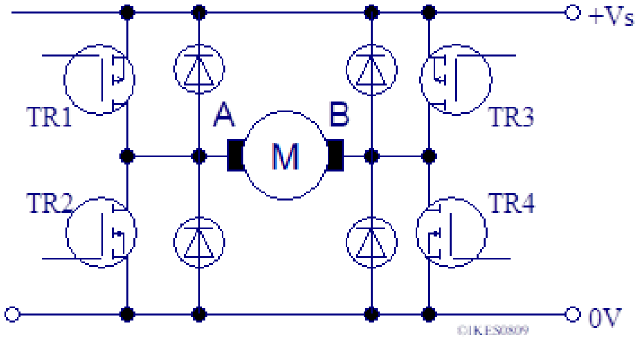

Figure 4: MOSFET H-bridge driver

6 Artificial Neural Networks

Software and hardware can be designed to act more like the human brain, with its neurones and connections. Neural networks are characterised by:

- Local processing in small processing elements
- Massively parallel processing, implemented by a high number of connection patterns between the small processing elements
- The ability to acquire knowledge by learning from experience
- Knowledge storage in distributed memory ie in the weighting of connections and patterns between the processing elements.

6.1 Differences between a neural network and a conventional computer

Most modern computers are based on a von Neumann or Harvard architecture so are sequential. They use a small number of complex processors with multi-threading (so can carry out multiple tasks simultaneously). In comparison, neural networks use large numbers of simple processing elements (which are parallel and so faster).

In a conventional computer, information is stored in central memory components, whereas in a neural network information is held in the weightings of connections between layers.

For a neural network:

- Elementary processors are highly non-linear
- Neurones are highly interconnected, allowing for a higher degree parallel processing
- There is no idle memory containing data or instructions — each neurone is pre-programmed and always active

In a conventional computer, a program must be written which will control the processor and store information at specific locations in memory. A neural network is 'trained' by giving it examples of inputs and desired outputs and allowing it to learn relationships.

Many neural networks learn through 'back propagation' — when the ANN is initially presented with data it makes a 'guess' as to a potential output, hen sees how far its output was from a valid output and adjusts its connection weightings appropriately. The transfer function within a neurone is non-linear, and only gives a significant output above a threshold — this helps to stabilise the ANN.

Back propagation performs error correction aimed at producing an overall weighting for the neurones that provides the lowest possible global error. To achieve this, ANNs often require a large number of iteration to achieve the best solution. This makes ANNs unsuitable for applications such as air traffic control, where one cannot afford to make mistakes — but ANNs are highly suited to activities such as classification and recognition.

It is possible to over-train a neural network — it will respond to only one type of input. If this happens then the network can no longer learn and it has been 'grandmothered'.

7 Robotics

ISO 8373 defines an industrial robot as *'an automatically controlled, reprogrammable, multipurpose manipulator programmable in three or more axes, which may be either fixed in place or mobile for use in industrial automation applications'*

7.1 Sensors

Sensors are required for robotic systems to receive information on their environment and surroundings such that they may react.

- Microphones can be used to detect both sound and mechanical vibrations, with detectors based on piezoelectric crystals used for ultrasound.
- Multiple antennae can be used to find the bearing of a radio signal
- IR photodiodes or imaging systems can be used to react to infra-red
- Photodiodes or camera systems give information on visible light. Two cameras allow for 3D capability — judging distances
- Photoelectric detectors, ionising chambers or G—M tubes etc can be used for UV and/or ionising radiation
- Microswitches can give information on contact (on or off)
- To measure forces, there are conductive polymers whose resistance changes when a pressure is applied, as well as strain gauges (resistance of thin wire changes as stretched)
- Accelerometers can be housed in a small IC, giving three axes of information.
- Global Navigation Satellite System (GNSS) receivers such as GPS, GLONASS, Galileo, provide position data (and thus speed) to a potentially high degree of accuracy
- Gyroscopes can detect changes in orientation
- Ultrasound, RADAR and LIDAR systems can provide distance information for avoiding obstacles — though a microswitch can be used to react to a collision.
- A 'Hall effect' sensor provides information on magnetic fields
- Two or more air-spaced plates which form a capacitor allow a robot to react to changing electric fields.

7.2 Actuators

- Electromechanical actuators include:
  - Conventional motors
  - Stepper motors
  - Solenoids — which are either on or off, so the armature is in the coil or not. Can be used for small movements between two fixed points.
- Pneumatic actuators:
  - operated by compressed gas, usually air
  - used where high speed and large forces are needed in a confined space
  - often controlled using conventional electronic systems (motors, solenoids, optical sensors)
  - power supply is usually a compressor driven by an electric motor, but a compressed air bottle could be used

- Hydraulics:
    - operated by compressed liquids
    - used where very large forces and pressures are needed
    - power supply is a pump driven by electric motor or diesel engine
- Electroactive polymers
- Piezoelectric materials
    - When an electric field is applied to a piezoelectric material the shape of the material changes with the change in applied potential difference.
    - High acceleration rates and short reaction times — so suitable for control of fast processes

7.3 Power sources

- Super-capacitors
    - Electrochemical double-layer capacitors can be as high as 5 kF with energy densities up to 100 kJ kg^{-1}
    - Capacitor, so can be recharged very quickly
    - Disadvantage that only operates at low voltages (2-5 V)
- Lead-acid battery
    - Energy density approx. 150 kJ kg^{-1}
    - Take a long while to charge
    - Very cheap, can deliver high currents (low internal resistance)
    - Tolerant of overcharging. Can be left on trickle or float charge for a long period.
- Nickel Metal Hydride (NiMH)
    - Unlike NiCad, contain no toxic poisonous materials
    - Higher energy density approx. 300 kJ kg^{-1}
    - 'Memory effect' less of an issue — do not need to be completely discharged before they can be recharged
- Lithium ion (Li-Ion)
    - Much higher energy density approx. 500 kJ kg^{-1}
    - Need for greater care with charging and discharging — can undergo 'thermal runaway' and cause fires. Consequently, greater restrictions on shipping.
    - Increased cost
    - Suitable for high power consumption applications.
    - Slow self-discharge rate
- Lithium polymer (Li-Po)
    - Very high energy density approx. 1 MJ kg^{-1}
    - Reduced risk of fire or explosion compared to Li-Ion
- Fuel cells are also an option (hydrogen and methanol) — they offer even greater energy densities, with methanol having the higher energy density and being easier to deal with.
- A petrol engine has the highest energy density of them all, even when operating at low efficiency.