The Technology and Challenges of Mirosot Robot Football

P.Robinson, P.Hall, J.Wolf , R.Phillips, C.Peck P.Culverhouse, R.Bray & AJ.Simpson,
School of Computing, Communications and Electronics, University of Plymouth, Plymouth, PL4 8AA, Devon, U.K.
(Tel : 044 (0) 1752 234078; Fax :044 (0) 1752 232540; E-mail: probinson@plymouth.ac.uk)

Abstract: Robot football combines microrobotics, artificial intelligence, multi agent real time control and dynamic vision sensing in a unique way. This paper discusses developments in each of these technologies resulting from over five years work by the University of Plymouth Mirosot robot football group. Problems encountered along the way are highlighted and various solutions presented. The limitations of present technology are discussed and the challenges of eleven a side team competition identified.

Keywords: Robot football, Mirosot

1. Introduction

The University of Plymouth Mirosot robot football project started in 1997 [1]. Since that time there has been major improvements in the constituent technologies. PC's have more than quadrupled in speed. Motors have reduced in size and increased in power. Camera technology, including the necessary frame grabbing cards, has become much more affordable. FM communication systems have evolved rapidly with higher frequencies now becoming available and new technologies such as Bluetooth offering the possibility of robust transmission at low cost. Mirosot competitions have likewise evolved to take advantage of these changes. Matches have increased from 3-a-side to 5-a-side. This year 7-a-side matches will become commonplace and in the near future there will be 11-a-side competitions. Each improvement in technology has enabled a corresponding increase in the complexity of the matches played. This is in line with the long-term aims of FIRA (International Federation of Robot Soccer associations) organized robot football, i.e. the creation of a bi-pedal robot footballer team capable of beating a human team by 2050 [2].

Notwithstanding recent improvements in technology many of the original challenges facing Mirosot robot football development remain. In order to examine the Plymouth experience it is useful to deconstruct a robot football system into six constituent parts and examine each part separately. The six parts are,

1. Robot body
2. Drive train
3. On board electronics
4. Communication link
5. Vision recognition
6. Strategy and control

Part 1 and 2 are exclusively hardware whereas parts 3, 4, and 5 are a mixture of hardware and software. Part 6 is exclusively software and for the sake of simplicity includes the simulation competition. This deconstruction is dictated by a combination of convenience and the fact that team members tend to focus their work on one or more of these parts. Clearly the overall system performance is only as good as the weakest link. For success in competition the performance of each part needs to be optimized. However, under match conditions experience has shown that it is often inherent low-tech problems that lead to defeat [3].

2. Robot Body

In practice many Mirosot robots can move at velocities of up to 2 m/s. Worse case conditions can therefore lead to possible head on collisions with absolute velocities of around 4 m/s. The forces involved are equivalent to dropping the robot onto a solid floor from a height of about 0.8 m, hence the need for robust body design.

During matches circumstances will occur where robots from opposing teams find themselves in a one-on-one 'pushing' competition. The ability to hold ones own position, or better still, push the opposing robot backwards is a further requirement for success. This suggests that the robot should be reasonably heavy, i.e. at least comparable to the opposition, and have good pitch surface traction. Traction is a function of several variables, e.g. robot weight, motor power, gear ratio, tyre material, surface coefficient of friction and surface contact area between the tyres and pitch. In this section discussion is confined to Mirosot robot body design.

Robot body design at Plymouth has evolved through four main stages. The first designs, circa 1997 and inherited from the Open University, were built using Lego bricks. During matches there was a tendency for these to fall to pieces. A further problem was that they did not meet the required MiroSot size limitation of a maximum 7.5 cm cube although at the time the FIRA authorities, and opposition teams, kindly overlooked this infringement. If the robots had been more successful then this lenient attitude may not have prevailed. By 1998 the Lego bricks had been replaced by bent sheet aluminum bodies. Although a big improvement over Lego bricks, the Mark 2 design had batteries access and on-board electronics problems. The third incarnation was the modular body concept. Each robot was constructed from a five 7cm square aluminum sheet panels held together by extruded solid aluminum pillars. The top face was the electronic printed circuit board. This Mark 3 design formed the basis of the Miabot robot footballers produced commercially by Merlin Systems Corporation Ltd. [4]. Due to its modular design it is also popular with undergraduate engineering students. Students interested in investigating mobile robots are able to construct and programme a Miabot type robot using only skills acquired during their degree studies.

The Mark 4 design, introduced in 2003, has followed the trend of internationally successful teams by machining each body from a solid block of aluminium. These robots are very robust but much more expensive to produce compared to earlier designs. Specialist milling machines and CAD design skills are required. Manufacture is not possible in laboratories equipped with only simple machines and hand tools. Although the basic body design is considered to be acceptable the robot 'lid' remains problematic and is a feature of on-going discussion and modification.

3. Drive Train

Each robot body design has required an investigation of possible drive trains. The Lego body was build around Lego D.C. motors and their associated plastic gears and small wheels. Size problems meant that the wheels were positioned at diagonally opposite corners of the body resulting in some interesting velocity control problems. Feedback was provided by hand made optical encoders, i.e. small circular sheets of white paper marked with radiating black lines. Ambient light, especially during matches, could sometimes result in intermittent signal loss. A review of possible motor drive chain combinations resulted in the choice of the Swallow Systems [5] A062 drive chain assembly for the Mark 2 Plymouth robot footballers. Several factors contributed to the final choice, not least of which was cost. At about £30 per pair the A062 motor/gearbox/shaft encoder combination was much less expensive than anything equivalent. They were also the correct size; two motor/wheel units side by side are approximately 80mm, i.e. just about fitting into the MiroSot size category. In addition they were supplied ready built thereby reducing building times. Weaknesses in the Mark 2 body design (wheels at one corner leading to control problems, inaccessibility etc.) rapidly led to its abandonment in favour of the modular Mark 3 design. However, the Swallow Systems motor gearbox combination proved to be a great success and was adopted for the Mark 3 robot.

More than 50 Mark 3 robots have been built using Swallow drive systems. Reliability and ruggedness have proved to be excellent. Their main weaknesses when compared to international premier league Mirosot robots is their relatively small wheels, narrow tyres, low power and worm/wheel gearing which has very high reverse friction thereby preventing 'free wheeling'. Notwithstanding these difficulties it was not until 2003 that this system was abandoned, at least for the first team, in favour of using bespoke designed large wheels and in line gearing linked to the relatively expensive Faulhaber 006 SR DC micro motor . This new design compliments the new solid aluminium body and is expected to remain the main competitive Plymouth Mirosot robot for the foreseeable future.

It is useful to review some of the simple calculations that informed the design of the new robot body and drive train.

Mass of vehicle = 0.6 kg
Wheel Radius = 50 mm
Gear Ratio = 8:1
Assuming a desired acceleration of 4m/s^2

Force = Ma
= 0.6 x 4 = 2.4 N
= 1.2N per wheel

Wheel Torque required = 1.2N x 0.025m = 0.03Nm
= 30 mNm

Motor Torque required = 30/8
= 3.75 mNm

Current required = 3.75 mNm x 0.144 A/mNm
= 0.54 A

To achieve this, the resistive voltage drop across the armature is IR = 0.54 x 1.94 Ω
= 1V

The voltage drop across the H-bridge, assuming about 1Ω per on-transistor ≈ 1V also. At constant no-load speed, the motor current is quoted as 29mA, giving a resistive drop of 0.029 x 1.94 = 56 mV. In practice, the vehicle will offer some load at constant speed. But the armature voltage drop is quite small. The theoretical maximum attainable robot speed can now be calculated. The nominal voltage for the motor is quoted as 6V. At this voltage, ignoring resistive voltage drop, the speed will be

6V / .000725 V/rpm = 8276 rpm.

(The maximum motor speed is specified as 8200 rpm). At 8276 rpm, the forward velocity is rω and

$$r\omega = 0.025 \times 1000\text{rpm} \times 2\pi/60 = 2.62 \text{ m/s}.$$

In theory it should be possible to accelerate from rest to 2.62 m/s at a constant 4m/s^2. The drive voltage at 2.62 m/s and 4m/s^2 would be 6V + 3V = 9V, assuming continued use of the H-bridge.

These simple calculations demonstrate that an acceleration of 4 m/s^2 seems a good figure to aim at. Given a coefficient of friction of around 0.7, then a theoretical acceleration of 0.7g (7 m/s^2) is attainable if the full weight of vehicle is on the drive wheels. It has been noted that the standard Mirosot robot, i.e. the robot manufactured and supplied by the Yujin Robot company [6] is uni-directional, i.e. the wheel axels are not in the centre of the robot but off-set towards the rear. This means that when accelerating the backward reaction force puts the weight over the rear wheels thereby reducing the chance of wheel slip. The downside to this is that wheel slip will be much more likely on deceleration. It was considerations such as these which resulted in the Mark 4 robot body, designed in house by Alan Martin [7], having its wheels placed centrally in the robot body.

4. On Board Elecronics

On board control may be divided into two parts; hardware which includes both the microprocessor and power electronics and the embedded software control algorithms. Much debate was devoted to choosing the microprocessor family most appropriate for this application. In 1999 the choice was between PIC or Atmel chips. Atmel won this debate. The deciding factors include its ability to be ‘flash’ programmed in situ through the serial port of the PC, lower cost, faster speeds and free programming software. Since 1999 there have been major improvements in both PIC and Atmel processors. Upgraded Atmels (e.g. speeds have increased from 8 MHz to 16MHz) with surface mount technology have been introduced and used on the new robots. However there has been no major debate leading to a review of the original decision

The drive electronics of the mobile robots can prove to be one of the most difficult areas to deal with. The first problem is the selection of components appropriate to the requirements of the system. Early versions of the control electronics used a common 5 volt voltage regulator linked to a H-bridge motor control chip. The original multiple control board solution quickly developed into a single surface mount board; this was again modified and produced as a through-hole version. It may seem strange that a redesign would progress from surface mount to through-hole technology. However, this was primarily done to encourage undergraduates to participate in building and programming the robots. The construction techniques for through-hole are well established in the undergraduate community and are easily repaired. Having said that, the design has proved to be very robust with few failures in operation. It is interesting to note that the design is now moving back to a mix of through-hole and surface mount technology. The fact that manufacturer’s only produce the latest Atmel microprocessors in surface mount form is making this change necessary.

As with any battery equipment making the most of the energy available is important. The first requirement is the need to provide regulated power to the system. A poor choice can waste power. Another potential problem is that at full turn-on, the combined motor currents can reach between 8 and 10 Amps. If this happen the battery terminal voltage falls below that required by the microprocessor resulting in loss of control. Some form of current limiting is therefore required.

The chosen board design uses a low dropout regulator (LP2954) to maintain a viable supply to the microcontroller and the H bridge control circuitry for as long as possible from the nominal 9.6 V (i.e. eight AAA 650 mAh, NiMH cells) battery supply. The dual full bridge driver (L298) was selected to provide the drive circuitry (rather than a discrete solution) because of the integrated features it provided in terms of drive capability, ease of interfacing to the microcontroller and board space. However it does not altogether support the imperative of minimum waste of power. The saturation voltage for the device is quite high thereby reducing the voltage ceiling available for the motors and, of course, dissipating power in the device. The latest control board design uses a Motorola MC33887 MOSFET dual full bridge driver, which exhibits a low drain to source on resistance and addresses the power efficiency issues as well as providing some additional features. Again it is only available as surface mount part. This was originally designed to fit on a sub board connected to the main control board as a through-hole assembly, but in the latest design is now part of the main board. By using this device it has been possible to very easily design in a hardware current limit using the provided ground referenced current feedback facility. This supplies an output current of 0.00266 of the H-bridge high side current.

The MOSFET driver integrated circuits are surface mount with the robot requiring one per motor. These H-bridges are controlled using a combination of a hardware current limiting circuit (described above) and an input from the microprocessor. Each H-bridge has its frequency of operation set up by the P.W.M. output from the microprocessor. With the ATmega8, a frequency of around 7KHz is setup up by a 256 multiplier. The H-bridges are capable of operating at frequencies up to about 20KHz. However the next multiplier step on the micro is 1024. This means the next frequency is around 28KHz, i.e. too far over the specified maximum for the H-bridges. The lower frequency does not pose any real technical problem; it just means that the H-bridges are operating within the audible hearing range, i.e. when the robots are moving an audible squeal can be heard.

5. Communication Link

The basic function of the communications system is to provide the players with reliable movement commands. Therefore a one-way communications link with the host computer is required. (In the not-too-distant future it is hoped that the robot footballers will be fully autonomous and therefore require two-way communications as they ‘talk’ to each other.) Due to the nature of the game and playing environment a UHF wireless system is used by all competing teams. The Plymouth system, in common with most others, comprises of an FM serial radio link between the host and each player. Currently the system may operate on one of two frequencies, chosen so as not to interfere with the transmission frequency of the opposition. The frequencies in use are 418MHz and 433MHz. At first sight it would appear that this a sufficient choice to compete with since one team can use one frequency, the other the alternative. In practice this is often not the case because there may be more than one match being played in the competition hall. Therefore the possibility of interference from nearby matches is always present.

System communications is controlled by the host computer which broadcasts packets of information at up to 50 times a second, i.e. after each field of the video signal is processed. The robots, fitted with a matching receiver, detect the availability of new data and parse the packet, each for its own commands. The packet structure currently in use is as follows:

[control 1 control 2 L R L R L R L R L R chksum]

Each part of the packet is one byte long, therefore in the above packet, used for both three and five a side games, the packet length is 15 bytes. The radio link operates at 9600 baud (bits per second) and uses the standard serial data format of 10 bits to a byte, i.e. 8 data bits plus one start and one stop bit. It is therefore possible to transmit (9600/10) bytes per second. Since 15 bytes per packet are required, the above system allows up to (960/15) packets per second to be sent, i.e. or 64 packets/sec. This is more than adequate to keep up with the video-processing rate.

As previously mentioned this system is fine for games up to five-a-side but FIRA have recently announced the introduction of seven and eleven a side competitions. It is clear that the current communications system will not be able to send out commands to all robots fast enough. An upgraded system has been designed which not only uses higher frequencies (869MHz), but also a faster baud rates (19200) thereby allowing data transfers of 1920 bytes per second. With the possibility of there being up to eleven players per side, this would mean a packet length of 27 bytes. 1920 divided by 27 gives 71 packets per second, i.e. plenty enough to cope with the video processing rate of 50 fields per second.

As the information carrying content of the communications channel increases then the tendency is for the demands upon the channel to increase in order to fill the vacuum. In this case a suggested extension to the above protocol includes sending upgraded PID (Proportional, Integral and Derivative – see section 7) controller values to each robot to enable improved control. Team performance could then be changed in order to effectively cope with dynamic environments.

The information revolution is continually developing new radio communications systems, many of which may prove useful for robot football. After a somewhat hesitant start Bluetooth systems operating in the 2.45 GHz ISM band are now becoming common. They employ frequency hopping (typically 79 different frequencies hopping at 1600 hops/sec, may be used) and spread spectrum techniques. Data rates are about 1 Mbits/sec are combined with an operation range from 10 to 100 metres for power outputs of between 1mW to 100mW. In theory Bluetooth should be an ideal system for robot football. It has good bandwidth and several devices can be connected in a communications web that is immune from interference. A Bluetooth controlled robot, built in Plymouth by Merlin Systems Corp., scored the first ‘Bluetooth’ goal in an international FIRA robot football competition during the 2003 World Championships in Vienna. Sadly Bluetooth, as presently marketed, has a fatal weakness as far as robot football is concerned; it has an undefined stack protocol initiation delay. Until this is rectified it seems unlikely that Bluetooth will be used in competitive robot football matches. Bluetooth competitors include the IEEE 802.11b (up to 11 Mbits/sec) standard, wireless USB and application specific RF as used for car remote locking systems. All these await investigation by the authors.

6. Vision Recognition

Plymouth’s robot soccer vision recognition systems have typically utilised an analogue camera connected to a frame grabber card in a PC. When considering the path taken by the image signal in such a system, it will come as no surprise that the image arriving at the interpretation algorithm is of poor quality. The reasons for this may be explained as followed.

1. The image is captured via a zoom lens onto a charge coupled device (CCD) image sensor. This has an array of light sensitive pixels coupled to an on chip analogue to digital converter (ADC) circuit.
2. The output from the image sensor is read by a processor within the camera and converted back into an analogue signal by a PAL encoder. This allows onward passage of the signal, in a well-known format, via a coaxial cable. The PAL encoding imposes a limit of 25 frames per second, each of which comprises of an odd and even pair of fields, i.e. the usual interlacing producing 50 fields per second.
3. The signal travels the length of the cable suffering losses (albeit small) as it goes.
4. The signal arrives at a frame grabber card in a PC. The card uses a further ADC to convert the signal back into a digital format. In some implementations the resulting signal has a different x & y resolution to that created by the original image sensor.

Items 1. to 4. combine to reduce the quality of the image arriving at the image interpretation algorithm. Using existing vision systems the following image degrading effects can be directly observed:

5. The pitch (which should appear as a mass of black) is infested with coloured dots (usually of just a few pixels) and broad horizontal bands of colour. Although it is easy for a human to ignore these bands it is difficult to perform the same task within an vision recognition algorithm.
6. Coloured blobs in close proximity bleed into one another producing one region with a slowly

changing hue.

7. Edges separating black from a colour exhibit the same effect with the black encroaching into the coloured blob on the left hand side. This causes the blob to appear smaller to the algorithm than a human would naturally interpret and causes a right shift in the blob of two or three pixels.
8. White lines on the pitch often exhibit ghost lines that appear as a blue line to the right of the original line.

These effects force the use of very robust algorithms along with wide pixel value thresholds. Unfortunately these robust algorithms impose a high computational cost, needing a powerful PC to run them or yielding a low frame interpretation rate on a low power PC.

The poor image quality, limitation on frame rate and the fact that a PC is a general purpose machine (so therefore not ideally suited to the task of image interpretation) lead the team to consider the use of a lower power processor coupled directly to an image sensor. This could be expected to exhibit the following benefits:

9. The image appearing at the processor would be of a higher quality since it has only undergone a single analogue to digital conversion and will not have incurred the losses caused by the PAL encoder and analogue circuitry. Because the image will be clearer it should also be possible to use more straightforward, computationally cheaper interpretation algorithms.
10. The frame rate will only be limited by hardware availability rather than by an artificial limit imposed by the PAL encoder.
11. A low power processor can be used since it will be specialised for the task of processing the image. In the case of a robot soccer system this will consist mainly of blob location, but may also include the location of match objects from the blobs.
12. Since the camera will perform the image interpretation it will not be necessary to transmit entire images to the PC. This should enable the use of a general-purpose port on the PC (USB or LAN) for connection to the camera.

Current robot football image analysis systems operate in a well-constrained environment. Robot size and surface markings are specified by FIRA, as are the pitch colour and markings. The only uncontrolled variables are illumination and off-pitch contributions to clutter and reflected colour interference. An early mobile robot demonstrated much that is required of a modern robot footballer [8]. Illumination is, of course, difficult to tightly constrain where colour temperature and lighting intensity variation can be influenced by external sources such as sun light, TV and camera lights or reflectance from adjacent bodies.

Recently FIRA organized robot football competitions have moved from five-a-side to seven-a-side matches and shortly eleven-a-side competitions will be commonplace. This taxes the control computer and demands higher processing speeds and improved robot control as well as more sophisticated strategy control. But it does not push the image analysis to become more intelligent. Cameras on players is a recent innovation that has the potential to extend the complexity of image analysis, but the pace of development is constrained by the physical size of the Mirosot league players and battery capacity.

Image processing techniques make much use of shape template and colour matching techniques. Unfortunately these are weak Artificial Intelligence (AI) methods. They do not model human or animal perception, nor do they exhibit any flexibility to cope with colour variation under natural lighting. To stay in the field of research, image analysis for robot football must move on from existing techniques. It must begin to tackle strong AI methods as proposed by Searle [9] where human intelligence is embodied in machines that grow up and learn complex interactions with their environment. This said, it is perhaps a more reasonable short-term goal to explore the capabilities of lower animals and insects as examined in Srinivasan [10]. In this case machines could be given more modest intelligence that enables them to track objects, avoid obstacles etc. in arbitrary scenes. These are goals of cognitive vision. Can a robot play football if its vision system provides a hierarchal visual attention process? Can it be guided using task, scene, function and object contextual knowledge gained through visual perception alone? The Cavier project is exploring this at present [11]

So, what impact will this AI have on robot football? Ideally it would allow a robot player;
(i) to operate under any lighting conditions, dynamic or static,
(ii) to recognise objects by learning rather than explicit programming,
(iii) to communicate visual information with both its colleagues and controller to schedule group behaviours,
(iv) to use robot-based cameras for real-time navigation and planning, only using scene camera for strategic control and
(v) to react to new unknown events in a controlled fashion rather than just ignoring them or halting.

Using robot football as a test-bed for AI is exciting and challenging The absolute constraints of real-time operation, where slow responses lose the game, give the research an edge. Presently hardware speed of processing and new statistical methods of image clustering and analysis offer response times in the 100ms-1second ranges. For example DiCANN [12] can now process and distinguish closely related marine plankton in monochrome images every 250ms, i.e. as accurately as experts but faster. As yet these response times are too slow for robot football where recognition delays of less than 20 ms are required. However, it seems reasonable to suppose that in the near future these, and other, AI based vision recognition techniques will find a place in robot football systems.

7. Strategy & Control

The University of Plymouth strategies for the MiroSot league are based on the SimuroSot strategy template. This template makes programming and simulating of the strategy easier. In addition strategies are interchangeable and platform independent. The strategy can be broken up into 4 levels

Mission planning and Role selection
Trajectory generation
Trajectory following controller
Robot internal controller

Improving the system from the lowest level up is best. This strategy model seems to be the obvious choice since other teams have taken the same approach independently [13].

Matching the robot trajectory to the vehicle dynamics is an important requirement of the strategy. Robot football robots are required to move as fast as possible. Therefore the vehicle speed and controller gains are pushed to the limits. If the demanded trajectory is beyond the system's capabilities the following problems may occur:

1. non-linearities, such as slip and saturation, can make the system unstable and the robot spins out of control,
2. the robot can not follow the required path and may hit an obstacle
3. the robot can not follow the required path close enough and misses the target

Therefore there is a need to generate a trajectory that matches the vehicle dynamics.

One possible approach to trajectory generation is the potential vector field method [14]. Vector fields have been implemented for;

- line of sight guidance to a target (positive field)
- approaching a target position from a certain angle with modified potential fields [14]
- obstacle avoidance (negative field)

If all these implementations are applied at the same time, a weighting function is required to fuse the vector fields together. Experiments have shown that the Gaussian normal distribution function is an acceptable method of combining these fields. (Alternative functions such as cylinders or cones could create a sudden change in heading angle, possibly exciting instability.)

As an example a typical combination of vector-field is analyzed where a Robot R avoids an obstacle (Robot O) on the way to a target point T, Figure 1.

The angle α is the difference between the instant heading angle θ of the robot and the vector ro. that points from robot to obstacle.

$$\alpha = a\tan 2(\vec{ro}) - \theta \tag{1}$$

Thus α is an indication of how close the robot is to a collision course with the obstacle. The smaller the angle, the more likely is a collision and consequently the importance of avoiding the obstacle is higher.

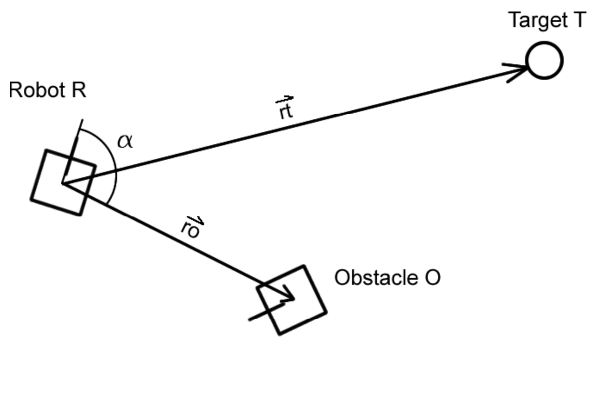

Figure 1 – Obstacle avoidance technique

The mission of the robot is to go to its defined target T. (In robot football this is likely to be the ball.) In order to take into account the obstacle on its way towards the target it must consider how close the obstacle is. The distance to the obstacle is defined as the magnitude of ro. A smaller distance to an obstacle means that is more important to avoid it.

An avoidance vector field V_O is defined which is normal to the mission vector field rt. The normalized target vector is V_T.

$$V_T = \text{normalize}\,(rt) \tag{2}$$

Suppose two vectors V_T and V_O are added together – 'fused' - with a Gaussian weighting function m*G(d).

$$\vec{V}_{MT} = \vec{V}_T + mG(|ro|, \mu, \sigma)\vec{V}_O \tag{3}$$

Where:
V_{MT} is the resultant modified target vector
m is a additional constant weighting factor
$G(\ )$ is the Gaussian distribution function
μ is the offset of the Gaussian hat
σ is the distribution of the Gaussian hat

As stated above there are essentially two factors that define how important it is to avoid the obstacle, namely α and $|ro|$. The the principle of vector field fusion is applied by relating the length of each vector to its importance towards the success of the mission at a particular point in the field. Thus α and $|ro|$ can be modelled as follows to influence the length of V_O.

$$\mu = -r_1 2(1 - \frac{1}{1 + e^{-\frac{\alpha}{\tau}}}) \tag{4}$$

Where;

r_1 is the maximum offset that α can cause.
τ is steepness of the slope (relationship of μ and α)

A larger τ will result in larger angles, already considered to be important. And the distance of the robot to the obstacle ro is modeled as the position parameter in the Gaussian function.

Finally, the resultant vector field V_{MT} indicates the new instant heading angle for the robot. A current attempt is to compare a path through a potential field with the robot's dynamic model in order to determine if the robot can follow it. This can be done in frequency domain, by comparing the bandwidth of the robot plus controller model to the bandwidth of the input signal when trying to follow the path. This approach can be taken further. A vector-field can be matched by design to the robot's bandwidth.

Robot performance may also be improved by tuning the controllers both in strategy and inside the robot. The Plymouth team has developed a simple way to measure the performance of the robot. It is done by measuring the time it takes the robot to run 10 circuits around the pitch using free-kick-corner dots as waypoints. Overshoot measurements provide extra data. This way of tuning takes into account communications and vision delays. It also allows a calibration of the guidance functions.

Control may be further improved by implementing some form of compensator into the low-level motor servo control loop. The classical PID (proportional, integral and differential) algorithm is a good stating point, i.e.

$$p(t) = \bar{p} + K_c\left[e(t) + \frac{1}{\tau_I}\int_0^t e(t)dt - \tau_D\frac{de(t)}{dt}\right] \quad (5)$$

Assuming that the sampling rate of the system is constant, then K is a straightforward constant multiplier. The integral term consists of a memory space, which takes in an error and adds the n incoming error to itself. The result is stored back to the original memory space for use in the p(t) computation. Since dt is constant (see above) it can be safely ignored. This can be thought of as being similar effects to a digital accumulator.

If the sampling rate is high, the integral could easily trigger an overflow problem, and therefore a range limiter may be required to limit the maximum value of the integral. The differential term itself implies computation by taking two values at two different instants of time. In an encoder system, commonly used for sensing the angular velocity of a motor, it is easy to run into a dynamic range problem when tuning a PID controlled system. It may seem that a PID controller will work best if it is set to the maximum allowable sampling rate, since the control loop is iterating at a faster rate. But, this could also result in a poorer resolution when computing the differential term. The higher the sampling rate, the shorter time it allows for the next error to develop before calculating the difference between the errors. When the motor is traveling at a lower speed with relatively high sampling rate, the resultant differential error could be zero. This usually causes serious computing problems. In summary, the instantaneous differential resolution of a system worsens with increased sampling rate due to digital quantization caused by the (usually optical) encoder. The real challenge is to achieve effective damping over a wide dynamic range of motor speed.

The problem can be effectively tackled by having a cyclic buffer, which holds a set of ***k*** errors. The instantaneous differential term can therefore be computed by subtracting the error at a time instant ***t*** with an error taken at ***t – k***, formerly stored in the buffer. The size of the buffer ***k*** is then adjusted systematically with τ_D to achieve the best damping effect. With better damping, the integral term can be tuned up to improve the PID respond.

Methods for tuning PID controllers cover a wide spectrum of possible applications and may be found in any good control engineering textbook [15]. Putting theoretical methods into practice can often prove problematical. Below is a summary of a successfully applied direct approach. Again it is assumed that the PID controller is modeled by equation 5. A trial and error (empirical) method is used.

1. Eliminate the integral and derivative action by setting the derivative and integral time constants, i.e. τ_D and τ_I ,to as near zero and infinity as is practicable.
2. Set K_c at a low value and put the controller on automatic
3. Increase the controller gain K_c by small increments until continuous cycling (limit cycle) occurs after a small set point or load change. The term "continuous cycling" refers to a sustained oscillation with constant amplitude.
4. Reduce K_c by a factor of 2.
5. Decrease τ_I in small increments (thus increasing integral control) until continuous cycling occurs again. Set τ_I to 3 times this value
6. Increase τ_D until continuous cycling occurs. Set τ_D equal to one third of this value.

The main disadvantage of the method is that it can often be very time consuming. Future investigations are planned using a Kalman filter approach to the control problem. Robot football by its very nature is stochastic, i.e. random fluctuations such as unexpected collisions, strange ball behavior, uncertain lighting conditions etc. make it impossible to predict with certainty the values of some signals at any given time. One view of Kalman filters is that they use statistical ideas and probability in an attempt to predict the future. Another way of looking at them is to say that they try to estimate something unknown from something whose value is uncertain. Therefore, on paper at least, they seem to hold much promise for robot football control.

8. Conclusion

Mirosot robot football technology can be divided into six separate, but interlocking, parts. Each part provides serious research challenges. From the author's perspective the limitations to good team performance are mainly situated in the game strategy, vision recognition and control sections. Robot football researchers around the world are working at the cutting edge of the relevant technologies in an attempt to meet these challenges. Dynamic obstacle avoidance is a topic is of interest to all teams. An obstacle avoidance strategy, based upon vector fields, has been successfully tested and seems to have potential for further development.

Biologically inspired AI vision and group dynamics will probably have a major role to play in future robot football teams. As these areas are improved greater strains will be placed upon the micro robots themselves. Players will need to be faster, more robust and intelligent. This suggests an improved awareness of the world around them. In order to achieve this, sensors will need to be incorporated into the robots. Inevitably, in the not too distant future, CCD camera miniaturisation and increased micro controller processing power will result in the abolition of the external PC and vision system. A knock on effect will be the added battery capacity required. Fitting battery power, at a realistic cost, sufficient for on-board cameras, intensive processing, high performance motor operation and two way communications, into an the 8 cm cube Mirosot robot requires technology not available at the moment.

Eleven a side Mirosot robot football brings with it another step jump in complexity. Perhaps the most obvious research challenge will be the enormous strain placed upon game strategies. It could be argued that in the highly nonlinear and stochastic robot football environment existing strategies are struggling to cope with seven side competitions. Robust methods will have to be developed using predictive strategies inspired perhaps from a combination of AI methods and biological behaviours. Cognitive based image analysis and embedded learning, especially in the area of group behaviours seem to hold much promise. Undoubtedly these technologies, if adopted, will have a major impact on team performance. Reliable, short range, two-way communications between 11 players in a noisy environment poses further problems. Bluetooth has been investigated but found wanting. What is certain is that existing techniques will struggle to cope and new solutions will have to be found.

The vision of 22 fully autonomous Mirosot robots playing a game of football independent of external computers or vision systems is the inspiration driving many researchers across the world. Such robots, which by definition will possess a form of distributed intelligence and learning capabilities, have application far beyond the confines of a football game. Clearly such machines will be able to evolve in some way by learning from their own experiences and adapting accordingly in much the same way as any animal. From here it is only a small step to the plot of many science fiction stories.

References

1. Robinson. P., Bugmann. G., Kyriacou. T., Culverhouse. P., & Norman. M., (2002): MIROSOT as a Teaching & Learning Tool, in Proc. of 2002 FIRA World Congress, edited by Ju-Jangn Lee (Seoul, Korea; Korea Robot Soccer Association, 2002), pp. 309-314
2. http://www.fira.net
3. Robinson, P. (1998), Soccer shock - England put seven past Brazil. IEE Computing and Control Engineering Journal, vol. 9, no.5, 231-232.
4. http://www.merlinsystemscorp.co.uk
5. http://www.swallow.co.uk
6. http://www.yujinrobort.com
7. Martin, A., (2003) Design of a new Mirosot robot. DCEE project report, University of Plymouth.
8. Nilsson, N., (ed.), (1984) Shakey the Robot, Technical Note 323, SRI Inernational, Menlo Park, CA.
9. Searle, John. R. (1980) Minds, brains, and programs. Behavioral and Brain Sciences 3 (3): pp. 417-457
10. Srinivasan and S. Venkatesh (eds) (1997), From Living Eyes to Seeing Machines, M.V. Oxford University Press, U.K.
11. Crowley J.L. and Reignier P, (2003) "An Architecture for Context Aware Observation of Human Activity", Workshop on Computer Vision System Control Architectures (VSCA 2003), Graz, Austria, April 2003.
12. Culverhouse PF, Williams R, Reguera B, Herry V, González-Gil S (2003) Do Experts Make Mistakes? Mar-Ecol-Prog-Ser. 247. 17-251]
13. Jong-Hwan Kim, Kwang-Choon Kim, Dong-Han Kim, Yong-Jae Kim and Prahlad Vadakkepat, (1998) Path Planning and Role Selection Mechanism for Soccer Robots,Proceedings of the 1998 IEEE International Conference on Robotics and Automation, KAIST, Korea.
14. Jakub Hrabec, Bohumil Honzík, (2002) Mobile Robots Playing Soccer, Brno University of Technology, Czech Republic, presented at the 7th International Workshop on Advanced Motion Control, 2002
15. Dutton, K., Thompson, S and Barraclough, B. (1997) The Art of Control Engineering, Pub. Addison-Wesley, pp 162-166.