Multi-site development of a FIRA large league robot football system

J.C. Wolf[1], J.D. Oliver[2], P. Robinson[1] and C. Diot[3]
[1]School of Computing, Communications and Electronics, University of Plymouth,
Plymouth, PL4 8AA, England
{ joerg.wolf, paul.robinson }@plymouth.ac.uk
[2]School of Engineering, University of Warwick, Coventry, CV4 7AL, England
[3] Automatique et robotique, Centrale Nantes, Nantes, 44000, France

Abstract

The difficulties associated with developing a competitive FIRA team using undergraduate students, spread across multiple sites are discussed. System modularisation and common interface standards are shown to be key elements for success. Details are provided of the approach used for each module with particular emphasis given to the vision system and time delays. A new prototype robot body design resulting from recent international match experience is discussed.

1. Introduction

Robot football has been a source of very many undergraduate individual and group projects in the UK [1]. The main FIRA centres are the Universities of Plymouth and Warwick plus the Open University. More recently the University of Nottingham fielded a team at the 2005 FIRA UK championships and the University of Oxford competed in the SimuroSot league. It is hoped that several more UK universities will become involved in the near future but funding remains a perennial problem and the search for appropriate sponsors seems never ending.

At Plymouth FIRA development is centred on final year BEng individual projects whereas at Warwick it is a predominantly multidisciplinary final year M.Eng group project activity. In both cases undergraduate students concentrate on part of a robot football system as a basis for their project. Often, due to the course constraints and time pressure the projects divert from the ultimate goal of producing a winning team. At both universities these projects are limited to a single academic year. During this time students from the participating universities will meet once or twice to take part in a national competition and share ideas and technology. Consequently experience is gained and then rapidly lost when the students graduate. Continuity is provided by academic staff and the odd PhD student whose main focus is usually elsewhere. This 'ignorance of history' among potential students wishing to work on robot football projects helps explain the unusual content of many of the team UK publications. A major aim of these papers is to provide basic source material for final undergraduate students. The resulting dilemma is how can a robot football system be designed to provide a basis for undergraduate student projects and at the same time be of sufficient standard to compete internationally?

This problem is typical for product cycles in industry, and the same principals for solving these can be applied here:

a) Modular design. This allows members of a team to work in parallel.
b) Common standards. Teams can work together confident that each part will fit.
c) Documentation. Ensures continuity and prevents wheels being reinvented.

It is therefore proposed that a modular robot football system, consisting of interlinked software modules with agreed interfacing and efficient documentation, is the best way to overcome the short termism of the development personnel, Figure 1. This structure may be compared to conventional robot football systems [2, 3] where many of the same problems pertain.

In practice software modules are interconnected using TCP/IP. This not only allows each student to work entirely independently on individual modules but also allows complete freedom in terms of choice of language and user interface. Additionally modules may be distributed easily over a number of PC's, thereby allowing the system to be moved from 5 to 11-a-side which generally requires multiple cameras and thus potentially multiple processors. Multiple cameras are usually preferred due to the limitations in terms of height (3m) and field of view of a typical lens. Another advantage of using multiple cameras is the balance between the cost of a single high resolution camera and two standard resolution cameras. Two standard cameras are usually less costly. In fact, a camera system using two 640x480 IEEE1349 FireWire web-cams, available for about $300US, is a low cost, practical option.

All groups have access to common robots, cooperatively designed and manufactured using the strengths of each group. This allows accurate comparison between systems while allowing new machines to be tested against a standard.

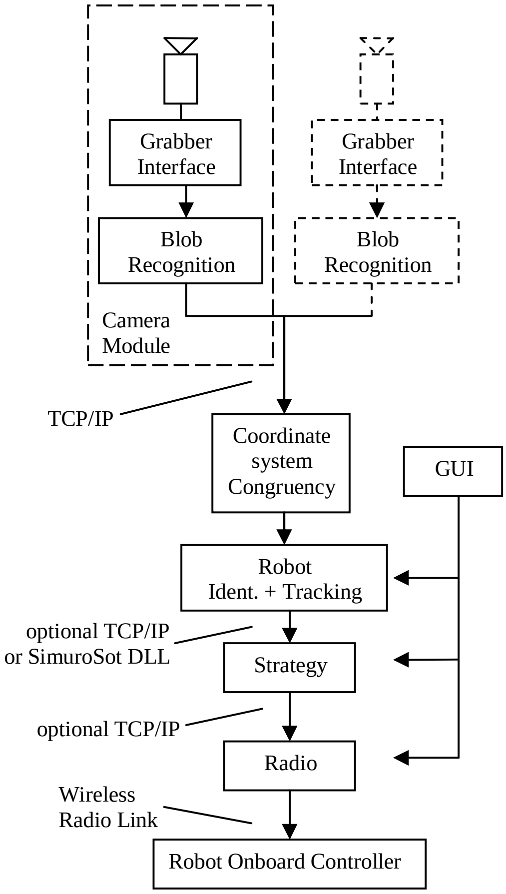

Figure 1. The modular system

2. Camera Module

This module directly interfaces with each camera, segments the images produced and then provides a list of unsorted blobs (segments with similar colour) for further processing. In order to simplify module interfacing, the amount of information forwarded from one module to the next is minimised, i.e. each module processes, as far as is possible, its own information. Therefore before information is forwarded spherical aberration corrections and parallax error corrections are applied to the segmented blobs.

2.1 Segmentation Process

The output of Firewire cameras is typically YUV, therefore in order to eliminate the overhead of colour space conversion the segmentation process is performed in this space. The radial method of dividing the colour space is used [4], i.e. colour is classified depending upon the angle of a given pixel from the centre of the UV space, Figure 2. Minimum thresholds were also placed on each colour Y component. In addition to providing reasonable lighting independence this method reduces the number of parameters that need to be calibrated from 4{Umax, Umin, Vmax, Vmin} to 2{βmax, βmin} , a significant benefit given the limited setup time available in robot football competitions. A fixed threshold D_{min} is also placed in order to eliminate the centre of the colour where colours are somewhat spurious.

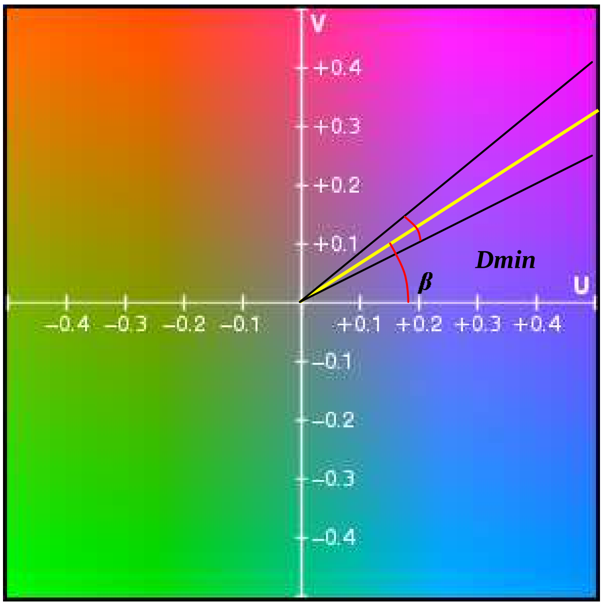

Figure 2. The UV-Space segmented.

2.2 Multiple cameras

Each camera interface module acts as a TCP/IP client and sends the information extracted from the pictures to a server module called Coordinate System Congruency (CSC) module. The following information is provided to the CSC module:

1. The age of the data, i.e. when the picture was taken.
2. The location, orientation and zoom of each camera.
3. A list of colour segments identified by each camera module.

Since the CSC module is connected to all cameras, it is able to monitor and synchronize the triggering of the cameras. The module's main task is to transform the coordinates of each identified colour blob (regions with similar colour) from the local camera coordinate system to a global reference coordinate system. Firstly the incoming coordinates are scaled from pixels to millimetres using a calibrated constant k in pixels per millimetre. Secondly the coordinates of the blobs are transformed by multiplication with the transformation matrix ${}^{R}T_{C}$.

$$ {}^{R}T_{C} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 & \Delta X \\ \sin(\theta) & \cos(\theta) & 0 & \Delta Y \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} $$

where: $\Delta X, \Delta Y$ represent the offset p_C between the global frame R and the camera frame C and θ is the orientation of the camera relative to the orientation of the global frame. Applying ${}^{R}T_{C}$ transforms a blob from the camera frame to p_R the reference frame:

$$ p_R = {}^{R}T_{C} * p_C $$

The origin of the global reference frame R is usually chosen to be the bottom left corner of the football pitch. Figure 3 shows that two cameras, when used with an 11-a-side pitch, will usually have an overlapping field of view. In this case he effect of imperfect orientation of the cameras is exaggerated.

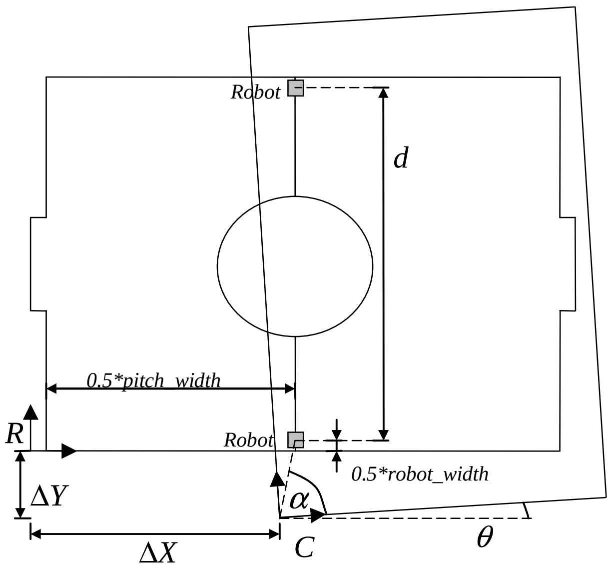

Figure 3 – Calibration process

2.3 The calibration process

Calibrating $\Delta X, \Delta Y, \theta$ and k of a camera can be carried out automatically. In the case of two cameras, Figure 3, the calibration process is carried out once per camera. Assume two known points on the pitch, indicated by the two robots on the centre line. The coordinates of these robots are known both in the camera frame from the vision system and in the global reference frame, since they are always put on the same spot for calibration purposes. (The robot positions can be placed in the overlapping fields of view, to speed up the calibration process.) The process starts with the determination of k.

$$ k = \frac{d_{pixel}}{d_{mm}} $$

The angle α is the angle between the x-axis of the camera coordinate system and a vector pointing from the camera coordinate system to the robot.

$$ \alpha = a\tan 2(p_{cy}, p_{cx}) $$

The line between the two robots is parallel to the y-axis of the global reference frame. Therefore θ is given by $\theta = 90° - \alpha$

In order to determine ΔX and ΔY, a vector diagram can be constructed:

$$ \vec{RC} = \vec{RP} - \vec{CP} $$

$$ \begin{bmatrix} \Delta X \\ \Delta Y \end{bmatrix} = \begin{bmatrix} 0.5 * pitch\_width \\ 0.5 * robot\_width \end{bmatrix} - l \begin{bmatrix} \cos(\alpha + \theta) \\ \sin(\alpha + \theta) \end{bmatrix} $$

where ℓ is the distance from camera frame origin to robot position P.

To confirm the automatic calibration system for multiple cameras a simple test was carried out. Two 640x480 web-cams were located 2.4m above a 7-a-side pitch, i.e. 2.2 x 2.8m, and a robot placed at different positions along the vertical centre line of the pitch, Figure 3. Table 1 shows the real robot coordinates, and the coordinates where the two cameras measured the colour blob (X,Y Blob). The root mean square error between the two positions is given in the last column. All dimensions are in mm.

X real	Y real	Camera Number	X Blob	Y Blob	RMS Error
1400,	25	1	1400.49,	25.05	0.86
		0	1399.63,	24.99	
1400,	700	1	1387.48,	701.00	16.33
		0	1399.22,	689.65	
1400,	1100	0	1402.38,	1107.94	17.51
		1	1392.88,	1122.66	
1400,	1470	1	1393.02,	1477.78	17.59
		0	1404.66,	1464.60	
1400,	1780	0	1406.12,	1782.33	12.21
		1	1400.29,	1793.06	
1400,	2175	1	1400.29,	2174.57	0.78
		0	1399.63,	2174.98	

Table 1. Measured positional errors between two calibrated cameras

The RMS error (distance error) approaches zero near the top and bottom end of the pitch since these coordinates have been used during the calibration process. Overall an accuracy of better than 18 mm was achieved. This remaining error is possibly due to the simplified model for correcting spherical aberration used, where tangential and radial distortion is combined into one constant.

3. System time delays

Sophisticated robot football systems use predictive control methods [5, 6] in order to compensate for the delays. Ultimately in order to control the robots at high speed the delays must be measured and accounted for. The measured position of a moving robot is relatively noisy and the error is magnified by the prediction component in the control system. For example, a velocity based prediction model including errors is given by:

$$P_{t+1} = P_t + \Delta t * [P_t + e_{Pt} - (P_{t-1} + e_{Pt-1})]$$

where P_{t+1}, P_t and P_{t-1} are the predicted position, the current measured position and past position respectively. e_{Pt} and e_{Pt-1} are the corresponding measurement errors in position, which are usually variable and unknown to the controller. The error in position e_{Pt} at time t and the error at time t-1 are multiplied by the delay time Δt. It is therefore desirable to reduce Δt. The error e can be reduced by higher resolution cameras, high quality cameras and a good calibration of iris and shutter for given lighting conditions. Since there will always be a disturbance e, the predictor is not perfect, resulting in a sub-optimal controller gain setting for motion control to keep the system stable. Ultimately slow motion control has a huge effect on the team's competitiveness.

The following sections describe how latencies in robot football systems can be measured.

3.1 Vision Delay

The vision delay comprises of image acquisition and image analysis (colour blob recognition). The delay can be measured by switching on a light and waiting until the PC has reported the light as switched on. The light is coloured red (a large LED is suitable) and recognized as a coloured region in the picture. The PC reports the light as switched on through the serial port. A storage oscilloscope records the time difference between the time when the light was switched on and the PC report over the serial port, Figure 4.

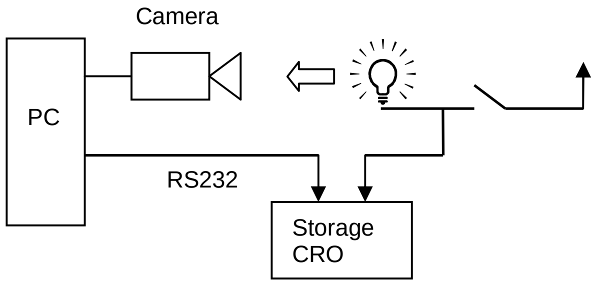

Figure 4. Delay measurements for vision systems

After subtracting the delays for the transmission and image analysis the results are shown below:

Camera system	fps	Delay
jAi CV-M91 PAL camera with a Matrox MeteorII card	50	12 - 31 ms
Logitech USB web-Cam	30	36 - 96 ms
Apple iSight IEEE1349 web-cam	30	46 ms

Table 2 – Time delays for three camera systems.

3.2 Processing delay

Given the complexity of the back end processing, the time taken to perform the image analysis, robot tracking and strategy calculations is one of the hardest to predict. Depending upon the situation the processing can take differing lengths of time, e.g. simple defending calculations to block a ball requires less time than striking calculations. However as a controllable delay it was decided that enforcing a fixed processing period was preferable in order to reduce uncertainty in the system parameters. The Pentium instruction read time stamp counter (RDTSC) is used to provide a real time clock and each module processing time was padded (extended) to produce a constant processing time. A longer, but defined, time delay is considered preferable to an unknown time delay.

3.3 Transmission Delays

Movement commands are sent in packages from the host computer to the robots via a wireless link. Each package usually contains the robot's target position and/or target velocity. Delays in this process may be classified as RS232 and radio module delays.

The RS232 delay is the delay of the operating system when accessing the serial port. Using a null-modem cable between two RS232 ports on a single PC allows for testing serial transmission delays of the PC only. After subtracting the time taken for the packet length a dead-time of 2 ms was measured. The test was performed using a dual Xeon 3.0 GHz PC, with MS Visual C++ 6.0 operating under Microsoft Windows XP Pro SP2. Radio modules were tested with a 19-byte packet at 19200 baud. The delay was measured from the start of the data entering the transmitter, to the start of the data coming out of the receiver.

Radio Module	**Delay**	**Std. Deviation**
Radiometrix RX2 / TX2	< 5ms	0 ms
EasyRadio ER900TS/RS	44 ms	0 ms
Bluetooth BlueCOM1	39 ms	16 ms

Table 3. Comparison of radio delays

In practice the Radiometrix module has virtually no delay; however a preamble is required to synchronize the receiver and transmitter and a maximum synchronization time of about 5 ms seems reasonable. By comparison, the EasyRadio module has a microcontroller, with an onboard buffer that waits for a full packet to arrive in the transmitter before transmitting [7] resulting in a 44 mS delay. The Bluetooth latency of 39 ms was measured with only 2 modules in the piconet. The latency would rapidly increase with more modules in the network. The Bluetooth delay is variable. Unpredictable latencies from 14 ms to up to 70 ms have been measured; the standard deviation of 20 measurements was 16ms.

Knowing the delay caused by the central program and the estimated delay of all other components, the data provided to the strategy module may be significantly improved using prediction. This enables the strategy module to make decisions based upon estimated states at the time of actuation as opposed to the time of sampling. This can be done via a relatively simple model of the differential drive robot and ball.

4. The tracking system

In order to reduce setup times the number of colours used on team shirts is kept to a minimum. Consequently robots are not uniquely identified. This causes the classic 'data association problem' of maintaining the robot's identity. Traditional approaches of greedy and global minima searches were applied, using the distance of each new sample from the robots last known coordinates as a heuristic [6]. These were found to be reasonably reliable under the majority of play conditions. The global search was however viewed to be inefficient since the complexity of the search is of the order o(n!). Efficiency can be increased by the simple expedient of matching all robots which have been designated with slow or purely rotational movements greedily, and then performing a global search on the remainder. This soon reduced the complexity of the algorithm.

Performance was further improved by the addition of an automated ID recovery algorithm. Implementation of this algorithm is triggered when, as time elapses, the confidence in each robot's identity decays. As soon as confidence falls below a given threshold and the robots current influence on the game is viewed as negligible, a set movement is performed (generally a pure rotation) to confirm that robots identity.

5. A new robot design

A major benefit of robot football competitions is that the robots are tested under harsh competitive environments and this allows comparisons to be made between robot footballer designs. It became clear at the recent European FIRA 2005 Championships that some parts of the standard UK design needed to be improved, e.g. the COG (centre of gravity) should be lower and the ball scoop redesigned to improve ball control. Subsequently a model of the existing robot has been created using CAD software [8], Figure 5. Analysis of this model compared very favourably to measurement made on the real robot. Of particular interest is the position of the COG which is above the wheel axle and off-centre, i.e. it is closer to the front of the robot. This has serious implications for robot control.

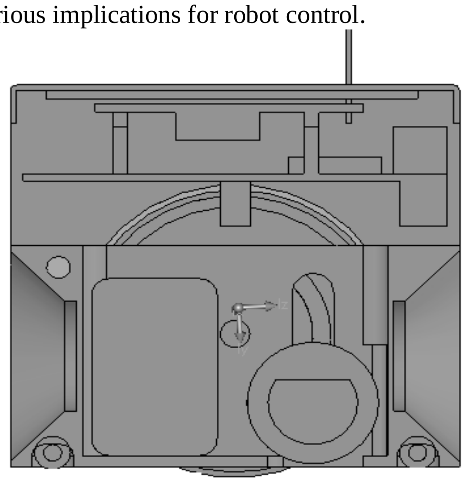

Figure5. The current robot design with the centre of gravity indicated

Comparisons between the model and real robot yielded;

- Weight:
  1. real: 523g
  2. model: 504g (without screws, bearings, pillars…etc.)
- COG : (millimetres)
  1. real: $\{X = 37.82 \quad Y = 39.05 \quad Z = 29.70\}$
  2. model: $\{X = 38.05 \quad Y = 37.35 \quad Z = 28.77\}$
- Principal axes of inertia and principal moments of inertia, gm.cm^2 taken at the centre of mass.

Ix = (-0.05, -0.99, 0.09) Px = 4804
Iy = (0.06, 0.09, -0.99) Py = 4909
Iz = (1.00, -0.05, 0.06) Pz = 5317

- Moments of inertia in gm.cm^2, taken at the COG and aligned with the output coordinate system and rounded to the nearest integer.

$$II(G, R_G) = \begin{bmatrix} L_{xx} = 5314 & L_{xy} = 27 & L_{xz} = 22 \\ L_{yx} = 27 & L_{yy} = 4806 & L_{yz} = 11 \\ L_{zx} = -22 & L_{zy} = 11 & L_{zz} = 4909 \end{bmatrix}$$

The result of this exercise is that a new prototype robot has been designed incorporating a range of desired improvements. Firstly the ball scoops at the front and back faces of the robot have been changed so that the ball is held steady when in contact with robot. This should allow higher velocity ball control. A much lower profile will lower the COG to below the axles and a careful choice of battery and construction materials will ensure that it is central in the robot XY plane thereby improving dynamic control. Wider wheels and improved tyre material will increase ground traction. New tyre designs are currently under test.

6. Conclusion

Most robot football development in the UK is carried out as part of final year undergraduate student projects. Due to the short duration of such projects there is no realistic chance of any one group producing an internationally competitive team. This has led to a core group of interested universities forming what is effectively a FIRA UK robot football league where competition and cooperation go hand-in-hand. Nationally there is competition and internationally cooperation.

FIRA UK members have agreed on a modular system design with common interfacing protocols. Such an agreement allows individuals or groups to work independently on particular parts of the system confident that their work will easily fit in to the overall scheme. An overview of the main modules and the progress made to date has been described. Colour segmentation when used with Firewire cameras is shown to provide robust illumination independence and at the same time reduce calibration parameters from 4 to 2. Automatic calibration suitable for low cost, two camera systems covering an 11-a-side pitch, demonstrated accuracy over the whole pitch of better than 17.6mm while experiments with a selection of commonly available camera and radio systems measured time delays of between 12 and 46 mS and 5 and 44 mS respectively. Bluetooth is shown to be unsuited for use in large league competition. A new robot body has been designed with a lower COG and improved ball control. The prototype will be built and tested shortly.

A major advantage of the described FIRA UK cooperative arrangement is that modules are language and machine independent. It is hoped that by pooling resources and expertise an 11-a-side UK FIRA team will be competing internationally in the near future.

References

[1] P. Robinson, G. Bugmann, T. Kyriacou, P. Culverhouse, M. Norman, *MIROSOT as a Teaching & Learning Tool*, in the Proc. of FIRA World Congress, Seoul ,2002

[2] Mi-Hoi Koo, Yun-Ki Lee, Kang-Hee Lee, Tae-Hun Kim, Jae-Kyung Lee, and Jong-Hwan Kim, *Development of a Robot Soccer System for MiroSot* in the Proc. of FIRA World Congress, Busan, 2004

[3] DongHun Lee, KyungHun Hwang, ByeongKyu Ahn, SeungMin Baek and TaeYong Kuc, *A Dual Camera Based Vision Processing System for Large League MIROSOT* in the Proc. of FIRA World Congress, Busan, 2004

[4] P.P. Jonker, J. Caarls, and W. Bokhove, *Fast and Accurate Robot Vision for Vision based Motio*n. RoboCup 2000: Robot Soccer World Cup IV, Proc. 4th Int. Workshop on Robocup, Lecture Notes in Artificial Intelligence, vol. 2019, Springer Verlag, Berlin, 2001, 149-158.

[5] J.C. Wolf, P. Robinson, J.M. Davies, *Vector Field path Planning and Control of an Autonomous Robot in a Dynamic Environment* in the Proceedings of the FIRA Robot World Congress, Paper 151, Busan 2004

[6] Kwun Hart and Manuela Veloso. *Reactive visual control of multiple non-holonomic robotic agent*, in Proceedings of the IEEE International Conference on Robotics and Automation, Belgium, May 1998.

[7] Easy-Radio ER900 Datasheet, Low Power Radio Solutions Ltd, 2004, (http://www.lprs.co.uk/)

[8] P. Chedmail, E. Dombre, P. Wenger, "*La CAO en robotique, outils et méthodologies*", Hermes, Nantes, 1998