Charlie 5: A Robot Ping Pong Player using a novel Vision System

Robin Bradbeer, City Polytechnic of Hong Kong,
and
John Knight, Highbury College, Hampshire, UK

Abstract

This paper presents the results of a novel attempt at machine vision using a series of rotating mirrors. Its use in robot ping pong player, Charlie V, is described. This robot came second in the 1992 World Robot Ping Pong Contest.

Further enhancements to the successful system are discussed in detail.

1 Introduction

Robot Ping Pong was first suggested by John Billingsley in 1983. The rules were published in 1984 [1], and the first contest was held in 1984, with the first international contest in 1985.

The first attempts were not very successful, and the contest was seen as providing more entertainment than serious research.

However, with the entry of large international corporations, such as AT&T Bell Labs in the USA, along with the reducing cost of machine vision systems, the research potentials were soon recognised.

In 1988 the seminal work on robot ping pong was published by Russell Andersson at AT&T [2]. Since then, the idea of ping pong playing robots has fired the imagination of most of those working with machine vision.

Robot ping pong differs in many ways from human ping pong. First, the table is smaller, being 2 m long and 0.5 m wide. The 'net' in the middle of the table is a wire frame 0.25 high with another frame 0.5m high on top. (Fig. 1). The ball must be served from the centre by a special mechanism as robots have no hands! The ball must pass through a 0.5 x 0.5 m frame at the end of the table before it can be hit. The robot must not occupy any space beyond the egde of the table.

Scoring is similar to normal table tennis.

The time taken for the ball to traverse the whole table is in the order of 1 sec, although ball speeds can be as high as 8 m/sec. This means that there is a small and finite time for any camera system to pick up the trajectory of the ball, estimate its flight characteristics and position the bat for a return hit. Unlike a human most robots cannot retreat from the table to give themselves longer processing time!

Many different techiques have been tried to overcome these problems. Most have used ccd camera systems with much display processing. This requires lots of computer power. Another approach is to use rotating mirrors, as described in this paper.

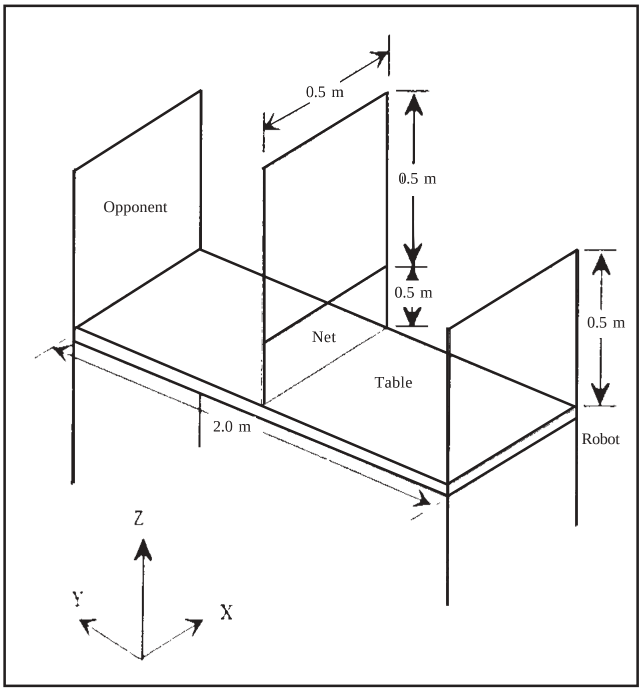

Fig 1. Configuration of robot ping pong table

2 Charlie 5's ancestors

Charlie 5 is the latest in a line of ping pong playing robots built at Highbury College that go back to 1985.

Charlie 1 had what can now be called a traditional configuration; an XY plotter up-ended and placed squarely up against the end hoop. The bat was spring loaded, cocked by a Beldon cable and released using a solenoid. The main carriage moved in the X axis (width-wise) and this in turn carried the bat carriage which moved in the Y axis (vertical). In addition to the bat carriage the main carriage also carried the camera, which was perched on the top, some 70 cm above the table [3]. The camera was mobile so that the field of view could be reduced and consequently the background interference also reduced. It also simplified the X axis tracking control system. When play started the camera was centred and from there the whole width of the table could just be covered. Once the ball had been sighted the control system moved the camera so that the ball remained centre camera and in doing so ensured the alignment of the bat and ball in the x axis.

The camera used photosensors and revolving optics. It provided 3 fixes per revolution, which, by the use of mirrors, were effectively from different view points. These were used to get a sequence of fixes on the ball in the Y and Z axes (Z axis being down the table). These were then used to predict the intercept point in the Y axis. Finally a light curtain was used to release the bat.

The overall performance was reasonable and it won a few contests. It's best feature was its x axis tracking and its resistance to background interference, but it's capacity to predict the height of the ball on intercept could have been out performed by a random number generator!

Charlie 2 retained the basic configuration of Charlie 1 but adopted a different approach to the Y axis intercept problem. The tracking technique had been successful for the X axis on Charlie 1 so Charlie 2 was to retain the technique for the X axis and also apply it to the Y axis. Y axis tracking is difficult because of the high speeds in the Y axis, pre and post bounce. The solution adopted was to use a light bat carriage freely running on a vertical track. Instead of being driven by a motor the carriage rested at table height up against a cocked spring. When the ball was "heard" to bounce by a microphone the spring was released and the bat carriage sent skywards in free fall. As the ball is also in free fall the two maintained a similar height as long as the following conditions were met.

1 The ball and carriage have a similar resistance to vertical movement

2 The ball and carriage have similar initial heights and they both start to decelerate at the same time.

3 The ball and carriage both have the same initial velocity (post bounce).

Condition 3 was the one most difficult to meet. The camera was used to detect the maximum height that the ball attained on its incoming flight. This is proportional to the post bounce vertical velocity and could be used to adjust the tension on the spring used to launch the carriage. The results were quite good and a notable improvement on Charlie 1.

Charlie 3 used exactly the same principles as Charlie 2 but aimed to execute them better. The microphone was replaced by a more sophisticated vision system, designed to predict when the ball would bounce. In that way condition 2 above could be better met. By knowing in advance when the ball would bounce the bat carriage could be held below table height and released early. If all went to plan the carriage would be at the right velocity and at table height by the time the ball actually bounced.

The result of all this added complexity was only a marginal improvement on Charlie 2. Charlie 4 did not get beyond the design stage!

3. Ping Pong physics

The flight of a ping pong ball is essentially affected by two forces - gravity and air drag. The equations of the ball's motion have been well documented [2], [4] and [5].

Basically, given the velocity vector *v*, the ball accelerates at *a*, where

$$\underset{\rightarrow}{a} = -C_d \,|\, \underset{\rightarrow}{v} \,|\, \underset{\rightarrow}{v} - \underset{\rightarrow}{g} \qquad (1)$$

The vector *g* is the acceleration due to gravity. The drag coefficient, C_d, is such that the terminal velocity comes out to be 9.5 m/sec [4].

The position of the ball at any time can be found when the initial velocity vector is given using classical techniques.

The most important problem from a robot's point of view is what happens when the ball hits the table. Theoretical work [4] estimates that, without spin, around 40% of the initial horizontal velocity should be lost. Spin will cause the velocity to increase.

In the vertical direction, deformation at bounce will absorb some of the energy. At the same time some of the original kinetic energy will heat up the ball and induces vibrations in the ball, table and atmosphere.

These forces have been derived [6] and can be defined as:

$$v_{fy} = \epsilon\, v_{iy} \qquad (2)$$

where v_{fy} is the vertical velocity before bounce, v_{iy} is the vertical velocity after bounce and ϵ is the coefficient of restitution.

In the horizontal plane

$$v_{fx} = v_{ix} - \mu\, v_{iy}\, (1 + \epsilon) \qquad (3)$$

where v_{fx} is the horizontal velocity after bounce, v_{ix} the horizontal velocity before bounce and μ is the coefficient of friction.

Initial experiments carried out by Knight showed that, for a single table surface and ball combination the reduction in horizontal velocity was around 50%. Further work by Bradbeer, using different surfaces and balls showed similar results. Experiments by both authors also confirmed the coefficient of restitution.

The results shown only refer to those carried out with different surfaces and different balls. The initial experimental results from Knight are within the same ranges.

From Fig 2, it can be seen that the vertical velocity before bounce was directly proportional to the vertical velocity after bounce. The slope of the best line was found to be -0.875. The outer limits were found to be -1.143 and -0.668. Therefore v_{fy} = -(0.875 +/- 0.268). This result agrees with the coeffiicient of restitution from previous theoretical work carried out on

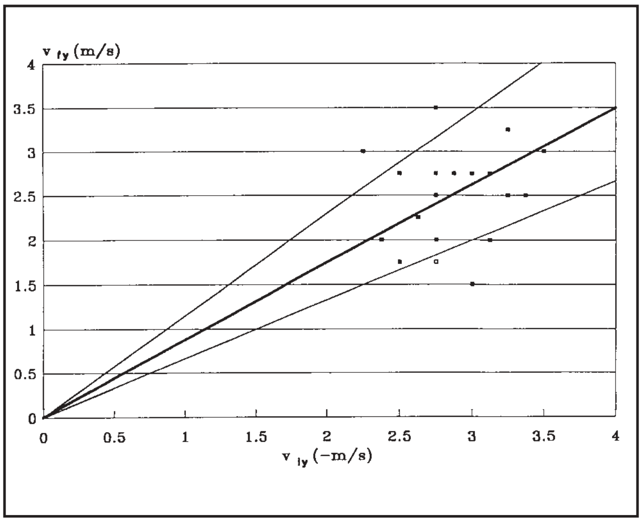

Fig 2. The graph of vertical velocity before bounce against vertical velocity after bounce.

tournament table tennis tables.[7].

Fig 3 shows relationship between v_{fx}/v_{iy} and v_{ix}/v_{iy}, . The slope of the best line was found to be 0.124, the x-intercept -2.6. The outer limits of the slope were found to be 0.140 and 0.122. The outer limits of the x-intercept were found to be -2.9 and -2.2. Therefore the slope was 0.124 +/- 0.016, and the x-intercept -2.6 +/- 0.4. From equation (3) the coefficient of friction was calculated as 1.387. Again, this was within the range of previous theoretical results.

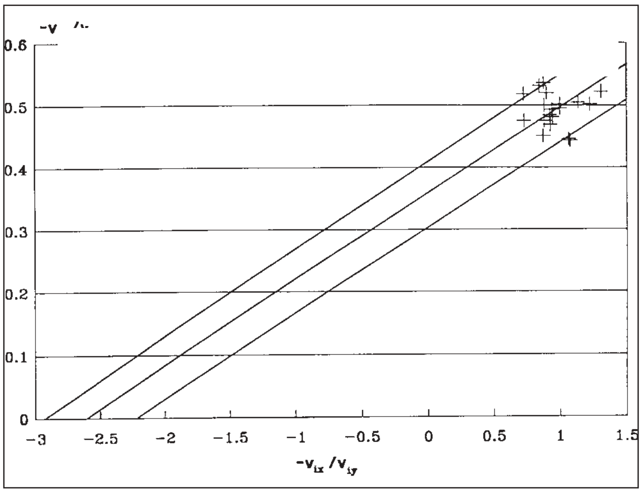

Fig 3. The graph of -v_{fx}/v_{iy} against - v_{ix}/v_{iy}

Fig 4 shows the velocity before and after bounce. This result shows that the velocity after bounce decreased by a factor of 2.03 +/- 0.22, ie about 50%, which was in agreement with theoretical results.

Similar experiments have been carried out to determine the position of the bat with reference to the point of bounce.

Up to Charlie 3, the main problem was to intercept the ball; now it was to play the stroke. Charlie 3 used Y axis tracking which could result in the bat carriage having a high vertical velocity at intercept. A higher velocity in fact than the Z axis velocity of the bat when the stroke was played. The combined movements gave the bat a direction of stroke that was often inappropriate. Charlie 5 was to address this problem by giving the bat carriage an extra degree of freedom, so that it could move back from the table end and play the ball when it attained post bounce maximum height (PBMH). This is when the vertical velocity is at a minimum. It would not be possible to hit all balls in this way as some would reach PBMH before they had passed through the end hoop while others would be through the end hoop and a metre or more beyond before reaching PBMH.

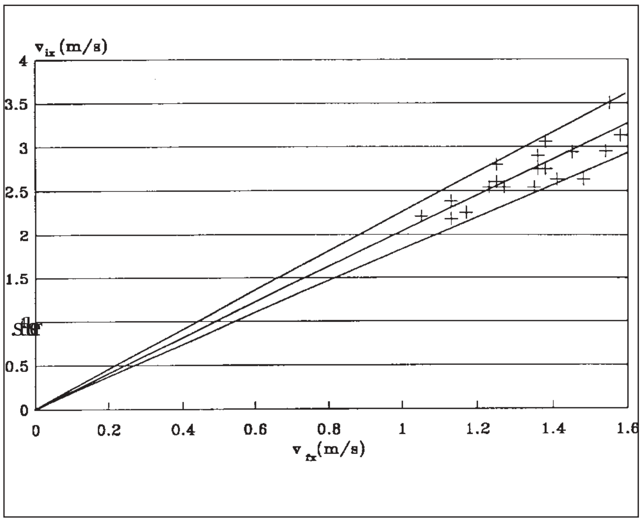

Fig 4. The graph of horizontal velocity before bounce against horizontal velocity after bounce

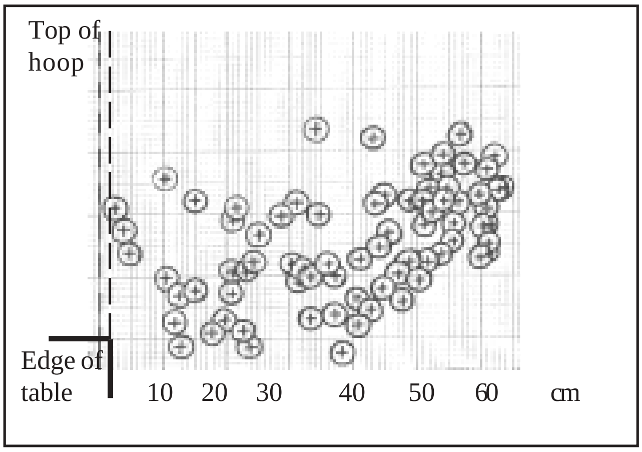

Fig 5. Bat/ball intercept point

A compromise was looked for and research by Knight established an envelope within which the ball could be struck at or near PBMH - Fig 5. In the course of the research it was found that ideally the ball should be stuck after PBMH at a point where the ball's line of flight coincided with the line the bat should take to play the stroke, normally between 20 and 30 degrees to the horizontal.The envelope turned out to be quite flat, such that with a little more compromise the area could be reduced to a straight line. This opened up the attractive possibility that the bat carriage would only have to move in one plane, an inclined plane approximately 40 degrees to the horizontal. The familiar XY plotter mechanism would therefore do the job, but now laid back at an angle of 40 degrees to the horizontal.

Further work by Bradbeer has shown that, for different surfaces and balls, the envelope could be better described as an arc at 40 degrees to the horizontal with a subtended angle of 10 degrees.

4. The camera

At the heart of the system is an array of 16 lens tubes, at the focal point of each is positioned a TIL 78 phototransistor - Fig 6.

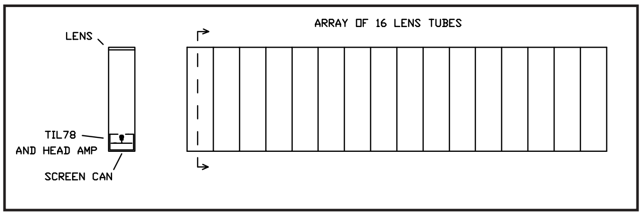

Fig 6. The lens array

The lenses are plain single element fresnel lenses with a focal length of 110mm. These thin plastic lenses have a good transparency at the 800nm wavelength and can be cut simply to fit the lens tubes. The lens tubes are square as shown in Fig 1 permitting the maximum aperture for the lens and giving a good packing density. Each tube has a narrow angle of view, sensing the ball only when two-thirds of the ball is within 2 degrees of the tube's axis.

Each lens tube is therefore a highly directional ball sensor with a range of 1.5m to 2m. A closely packed array of such tubes forms a curtain which the ball cannot get through undetected. When this array is positioned horizontally at the end of the table the curtain fits the playing space in width and length. To cover the whole playing space, that is width, length and depth, the array must scan the space and that is done by reflecting the curtain through a revolving mirror, as shown in Fig. 7.

The mirror is double sided and covers the full width of the array, measuring 7cm by 51cm. The maximum scan angle is 105 deg. and the mirror speed of 1300 rpm. gives a scan every 23ms.

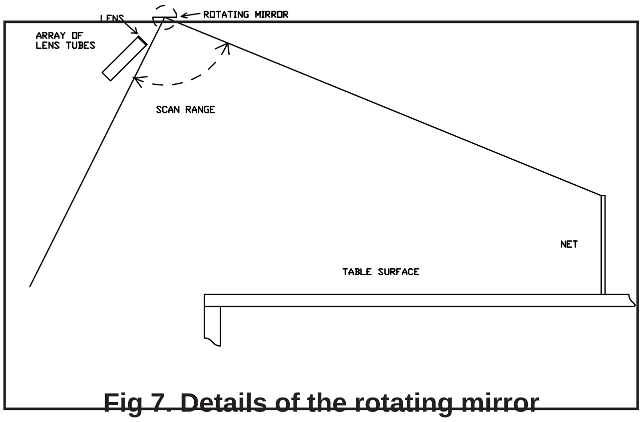

Fig 7. Details of the rotating mirror

5. The control system

The X and Y tracking system had proved effective for Charlie 3 and it was decided to retain the X tracking system which had survived from Charlie 1 but to try another new system for the Y axis. The camera's low resolution made the prediction of an intercept point by analysis and modelling a poor candidate while the camera's ability to cover the ball through a large part of it's flight meant that a system which could use that information could give Charlie an edge. An empirical system could do that and to my knowledge the technique had not been applied to a ping pong playing robot before, that was the technique selected.

The method chosen uses a database containing Track Records, these are sets of data describing known flight trajectories for incoming balls . Each track record is a set of bearings on the ball's position for one incoming flight. The bearings are not Cartesian data, they are simply the angle that the ball's apparent position makes to a reference point.

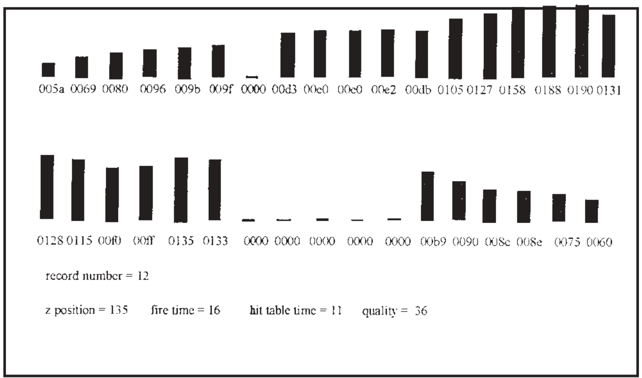

Fig 7. Typical track record

First the machine must go into Learn mode. 20 of these bearings are recorded for each flight, the operator enters the best position for the bat and the time that the stroke should be played to return that or a similar incoming ball. These two parameters are then stored as part of the same record. Once a range of trajectories have been stored in this way the machine can enter Play mode. In Play mode the database is searched for a similar set of bearings to the set it is collecting from the camera. The retrieved record should then contain a suitable position for the bat and a time for the stroke to be played.

Other information to be recorded in the track record must be entered manually. This 'open loop' system is one reason why the bat can only track and hit about 50% of the balls. This information includes the estimated PBMH, the estimated delay in firing the bat mechanism, the time to hit the table and the 'quality' of the information. All these parameters have to be estimated by the operator.

The angle of bearing and track record system is inherently a very powerful one. It lends itself to both fuzzy logic and neural network approaches in determining the firing and positional information for the bat. However it does have some disadvantages, which current work aims to overcome.

6. Current work

The first problem is that the track records are not unique enough for closely related but critically different trajectories.

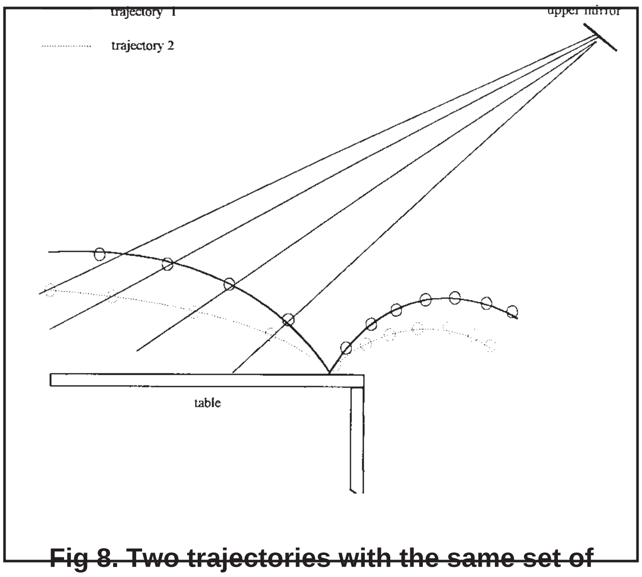

Fig 8. Two trajectories with the same set of bearings and hit table time

A number of attempts have been made to overcome this problem. A more defined look-up algorithm has been tried, but does not seem to yield the uniqeness required, even when faster computers are used, so that more information can be processed before the bat is committed.

Using more than one mirror is clearly the way to go, and a number of different ideas have recently been tested.

First, a second, stationery, mirror was placed low down on the table, Fig 9.

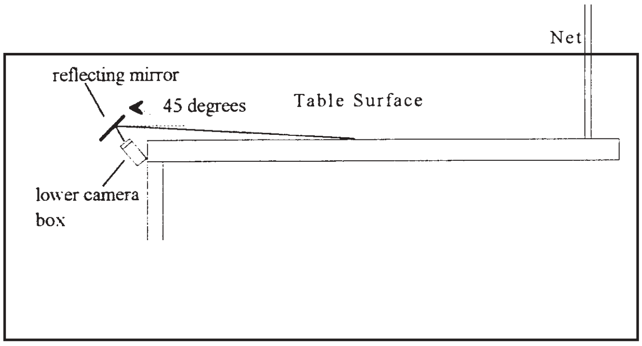

Fig 9. Position of lower mirror

This was not very successful! A stationery mirror at the top, however, has proved far more useful, and provides one extra bit of key information. Even if two trajectories have very similar bearings, the bearing, and thus the time, at which the ball is momentarily in view of the upper sensors is different. This is shown in Fig 10.

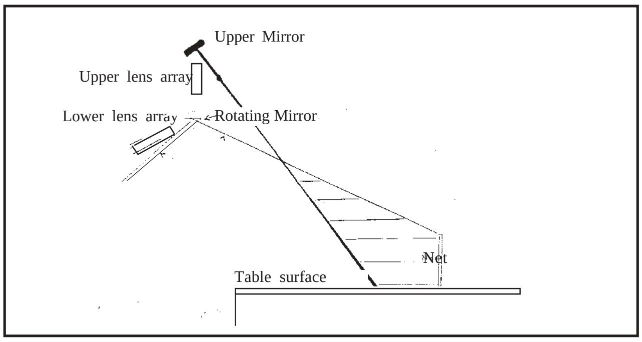

Fig 10. Position of upper lens array

The upper lens array has only six phototransistor detectors as its only function is to detect when the ball passes through its very narrow angle of view. Initial results from this system indicate that it has the potential to overcome the problems identified.

One other area of research is to make the whole operation of learning more closed loop. It is planned to have an led light curtain around the front frame of the robot so that an accurate x-y coordinate can be obtained to reduce any errors in the detection system.

Another area of interest is the control of the bat platform. At the moment this moves along the 40 degree plane identified above. The current wooden frame is not adequate to support the curved motion necessray to refine the bat hitting position, so a light metal frame with better feedback mechanisms on the bat and carriage postions is being constructed.

7. Conclusions

The original concept behind Charlie 5 has proven itself in competition, so the basic concept is correct. However there are some fundamental problems that have to be overcome for it to hit the ball more than 50% of the time.

The rotating mirror approach is clearly well suited to this type of application and has potential in others. The addition of extra mirrors to provide a cross correlation and more information to the computer means that less 'guess work' is needed in the software. With the addition of fuzzy logic software, maybe allied with a neural network approach, the slowness of the look-up table algorithm should be overcome.

The rotating mirror approach to following fast moving, small ojects is a less costly, and more robust method that that traditionally using video cameras and video signal processing.

8. References

[1] Billingsley, J. :Machinroe joins new title fight: Practical Robotics, May/June 1984, pp14-16 (UK)

[2] Andersson R. L. : A robot ping-pong player: Exeriment in real-time control: The MIT Press, 1988 (USA)

[3] Knight, J. and Lowery, D : Ping pong playing robot controlled by a microcomputer: Microprocessors and Microsystems, vol. 10, no 6, pp332-5, 1986 (UK)

[4] Brandan, M. E., Gutierrez, M. et al : Measurement of the terminal velocity in air of a ping-pong ball using a time-amplitude converter in the millisecond range: American Journal of Physics, vol 52, no 10, pp890-3, October 1984

[5] Frohlich, C. : Aerodynamic drag crisis and its possible effect on the flight of baseballs: American Journal of Physics, vol 52, no 4, pp325-334, April 1984

[6] Brody, H. : That's how the ball bounces: The Physics Teacher, pp494-497, November 1984

[7] International Table Tennis Federation: technical File: ITTF, East Sussex, UK