Aberystwyth University

Discrete dynamics and the space-time possession game
Law, James; Johnson, Jeffrey

Published in:
FIRA Robot World Congress

Publication date:
2004

Citation for published version (APA):
Law, J., & Johnson, J. (2004). Discrete dynamics and the space-time possession game. In *FIRA Robot World Congress*

DISCRETE DYNAMICS AND THE SPACE-TIME POSSESSION GAME

James Law and Jeffrey Johnson

Department of Design and Innovation
The Open University
Milton Keynes, MK7 6AA, UK
j.a.law@open.ac.uk j.h.johnson@open.ac.uk

ABSTRACT

Like computer chess, robot soccer can be thought of as a game in which the agents on one team cooperate to control more space more comprehensively than the opposing team. Tactical issues such as winning and holding the ball add a complication to the issue of controlling space. As a research experiment we are investigating a 'space-time possession game' in which the only issue is the space controlled by the agents individually and in combination. The resultant game is similar to 'Go', but the objectives and rules are different. We have begun to study this using cellular automata, and we report on a series of experiments in which the 'robots' move at random, move to maximise their individual space, or move to maximise the whole team's space. Our results suggest the last is the most successful approach.

1. INTRODUCTION

Many problems involve humans structuring real and abstract spaces, and we are remarkably good at it. These include everyday problems of getting around in two or three physical dimensions, both living for the moment and planning for the future. They also include everyday problems of structuring multidimensional information spaces of great complexity. One goal in AI and cognitive science is to understand how we achieve amazing feats of understanding by effortlessly finding useful structure in an astonishing range of real and abstract environments.

We come to these problems through our interest in robot football [1], in which teams of robots collaborate to score goals against other teams of robots. The problem is elegant and powerful because it is easy to express and understand, while taking researchers to the frontiers of AI, cognitive science, and engineering. As we reflect on human soccer players we realise that they have wonderful abilities in perception, cognition, and movement, compared to the most advanced robots of today. In this paper we are concerned with the cognitive ability of soccer players. In particular, we are concerned with the problem of how human soccer players are able to construe the soccer pitch, and find useful structure in what they see to guide their play.

The research starts with a set definition. We have two teams of n robots, and a ball. We will suppose $n = 11$, even though for most robot soccer games the number has predominantly been three, five or seven. We also have a 'pitch', and this is where our problems begin. What is a 'pitch'? In robot soccer the pitch is, ideally, an array of pixels in an image. This provides the lowest level of representation. At a higher level of representation, some sets of these pixels are special, being designated the 'centre spot', the 'goal area', the 'penalty spot', the 'corners', the 'goals', and so on. But the really interesting higher level spatial structures are those highly dynamic configurations that emerge during the game. As the players move and position themselves they create spatial structures that enable and constrain the action [2]. It can be argued that part of the pleasure of watching football is to observe the formation of these dynamic structures, whether or not goals are actually scored.

This then is the subject of this paper. How do human beings structure space to promote the occurrence of desirable events, and inhibit undesirable ones?

The benchmark of computer chess reached a scientifically inconclusive crescendo in 1997 when IBM's Deep Blue computer beat the world champion, Garry Kasparov. Nonetheless, there is much to learn from chess as played by humans and machines, since it is quintessentially concerned with *structuring space*.
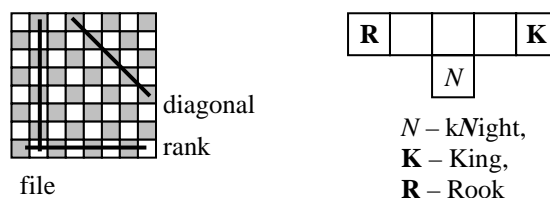
(a) structured space in chess   (b) the knight-fork

Figure 1. Structured space in chess

This is illustrated in Figure 1(a), where we give *names* to configurations of squares on the chess board. In Figure 1(b) the spatial structure of the three pieces forms a structure called the *knight fork* in which the knight checks the opponent's king, and threatens the more valuable rook. These structures were known long before the invention of electronic computers, and the way that humans understand and manipulate them has long been held as an indicator of human intelligence. Small wonder then that Alan Turing and others at the forefront of computer technology and Artificial Intelligence should seize on chess as a test of machine intelligence.

From the perspective of today, it can be seen that one of the very attractive features of chess for testing machine intelligence is the simplicity of its form and its rules. A grid of sixty four squares and thirty two pieces is a 'small' system. The rules of the system are also relatively straightforward, determining how the pieces can move, and what constitutes a win or draw. Crucially the dynamics of chess are very simple seen from a modern viewpoint: chess has a very simple time structure, and it is non-chaotic. In other words, (i) time in chess is governed by simple alternate move events (although human players are constrained to another time governed by the clock, bringing in an element of psychology), and (ii) when a chess game is started from the same position, and the same moves are played, the same outcome will be observed as on previous occasions.

In contrast, the new benchmark problem of robot soccer is an order of magnitude more complex. The robot soccer pitch on a computer screen has millions of squares, and the granularity of real robot soccer is of the order of molecular distances, if not actually continuous. Similarly, the time of robot soccer is measured in fractions of a second according to digital computer clocks, while the embodied behaviour of the robots exists in the micro-time of physics. If chess has combinatorial explosion in its space-time structure, robot soccer has hyper-combinatorial explosion. Furthermore, real robots are undoubtedly chaotic in their dynamics, and simulated robots are also very sensitive to initial conditions.

Robot soccer, as a physical system, is subject to the space-time laws of physics. So are humans. Humans exhibit great intelligence in the way that they handle space and time, although how we do it remains mysterious, and it is the spur for much research in AI. Those who have played soccer will know that spatial configurations are very important, and that a player can make a great contribution to his or her team without touching the ball at all. To be in the right place at the right time is everything.

Figure 2(a) shows how individual players command parts of the pitch, while Figure 2(b) shows the composite picture for the teams. In this example the winning team consistently controlled most of the pitch, and this is clearly a great advantage. However, as we can learn from chess,

sometimes the important thing is to be able to control special parts of the system at just the right time.

The example shown in Figure 2 is taken from the 2002 finals of the RoboCup simulation league. As we studied this example, we have been struck that the underlying problem in robot soccer, and much of 'intelligent human activity', is the ability to structure space. For that reason we have tried to strip away particular contexts such as chess or soccer, and create a universe that focuses purely on structuring space and time.
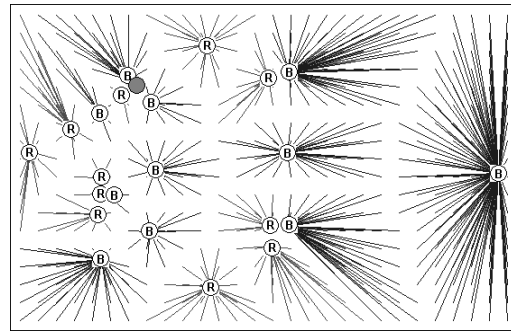
(a) relationship between players and grid squares
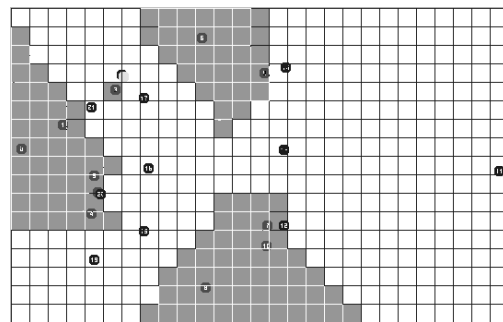
(b) the winning team 'owns' most of the pitch

Figure 2 Structuring space in robot soccer.

We will use a cellular automaton in our preliminary experiments. Such systems have been used reliably in path planning for robot football [3], and multi-agent cooperation [4].

2. THE SPACE-TIME POSSESSION GAME

Our interest is in how humans structure space in useful ways. To investigate this we wanted to get away from the details of robot soccer, and tactics regarding moving ball around in order to score goals. This led us to define the *space-time possession* game described below.

Let G be a grid of cells. For simplicity we'll assume G is composed of squares, although other planar tessellations are possible. Each cell has eight *neighbours* in the usual way. Let A and B be two sets of *players*. Each player

occupies a cell at any given time, which will be represented by the notation (x, y, z), where z is a member of A or B, or a symbol 'a', 'b', or 'c' denoting unoccupied areas 'claimed' by team A or B, or both 'c'. We assume that the system has a discrete *clock*. When the clock ticks any player can move to an adjacent unoccupied cell.

A player's claimed area is a function of distance. Each player possesses all the squares which are closer to it than any other player, such that the whole grid is owned by one or both teams (Figure 3). Squares equidistant from either team are considered shared, with distances measured using chessboard distances i.e.

$$D_{Chess} = \max\left(\left|x_2 - x_1\right|, \left|y_2 - y_1\right|\right)$$

a	A	a	c	c
c	c	c	c	a
b	B	b	a	a
b	b	b	a	A
b	b	c	a	a

Figure 3. Grid possession by teams A and B

Figure 3 shows the convention for pitch ownership. The two squares marked 'A' are players from one team and posses areas of the grid marked 'a'. 'B' is an opposition member, and its claimed area is notated 'b'. Furthermore, 'c' are the squares equidistant from both teams, and so jointly claimed by both sides.

The objective of the game is to control strategic areas of pitch, by outmanoeuvring the opposition. Initially, each team starts in an opposing half of the pitch, with players able to take up any position in their team's half. At every time step, the players are free to move a distance of one square in any direction, unless it is already occupied by another player. Player movements are controlled by a team strategy, in the same way as in robot soccer, with two strategy programmes pitted against one another. The winning team is that which meets the following criteria, only one of which will be applicable to each game:

i. The team that holds the largest distributed area after N_d clock ticks.
ii. The team that holds the largest contiguous area after N_c clock ticks.
iii. The first team to hold M_d distributed grid squares.
iv. The first team to hold M_c contiguous grid squares.
v. The first team to hold M_d distributed grid squares for N_d clock ticks.
vi. The first team to hold M_c contiguous grid squares for N_c clock ticks
vii. The first team to link either end of the pitch with one contiguous set of claimed grid squares.

These simplified objectives and the constraints of the environment, combined with the competitive nature of the game, provide ideal conditions for developing advanced multi-agent controllers, and investigating complex interacting systems.

The goals also represent generalisations of wider problems. Possible applications for game strategies include network science, mapping, search and rescue, surveillance, planetary robotics and mine sweeping. For our research we will begin by focusing on developing distributed control strategies for the first winning condition.

3. EXPERIMENTAL RESULTS

We have programmed the possession game in MATLAB, and implemented three simple strategies. In our game, players take it in turn to move, with Team A moving first. Each team consists of 5 players. The strategies run once at every time frame, and are based on the state of the pitch in the moving player's neighbourhood. They are:

1. Random - the moving player is assigned a random direction to move in. This can be into any of its 8-neighbours, or it can remain stationary.
2. Local player maxima (LPM) - the player's pitch possession is calculated for its current position, and for moves to every available 8-neighbour. The move which results in the player owning the largest area is taken. If there will be no change, the player remains stationary.
3. Local team maxima (LTM) - for each of the nine moves a player can make, its effect on the global team area is calculated. The player moves to the cell which gives maximum benefit to the team. If there will be no change, the player remains stationary.

We have experimented with playing each strategy against one another, on either team. There are nine combinations with each team adopting random, local, and global maximising strategies. Each run began with the players in the same symmetrical opening positions.

The following comments describe our observations. The accompanying graphs show the change in area occupied by both teams, during the course of each match. Each graph shows a single game, conducted with a different set of competing strategies:

Team A random strategy – Team B random strategy
Results are radically different for each match. Either team may end up with a larger area, irrespective of who moves first. No graphical data is given for this pair of strategies, since the areas owned by either team fluctuate randomly.

Team A local player maxima – Team B random strategy
Initially Team A collects most ground (Figure 4). However, team mates begin to compete against each other, rather than cooperate. This causes the rate of accumulation of pitch to slow down between 100-200 clock cycles. Pairs of same-team players are then repelled from the randomly moving opponents, which then retake ground from Team A during clock cycles 200-300.
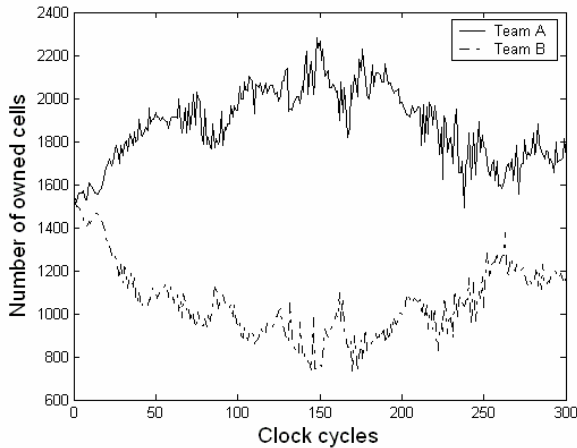

Figure 4. A-Local player max, B-Random

Team A local team maxima – Team B random strategy
Team A rapidly collects most of the pitch and holds it (Figure 5). As time progresses the space occupied by Team A fluctuates, since it is not building 'robust' structures, and the random motion of Team B may fortuitously gain ground for one clock tick, only to have Team A recapture it after the next clock tick. In the first 75 cycles only two Team A players are in a position to take pitch from team B, causing the initial gains. At cycle 75, two more players become actively involved, causing the renewed growth in pitch possession. The fifth player finally joins the attack around cycle 275.
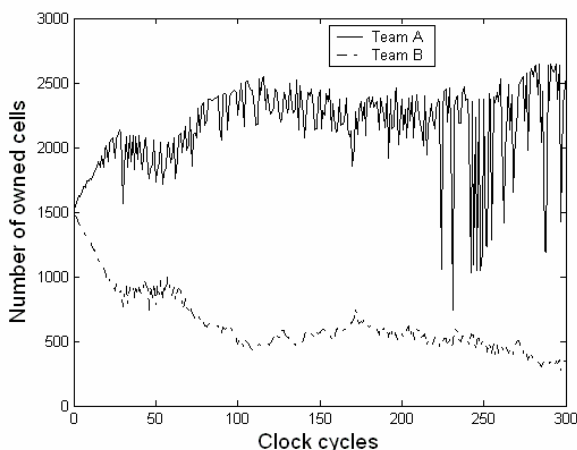

Figure 5. A-Local team max, B-Random

Team A random strategy – Team B local player maxima
Observations show similar effects of play to those described for the strategies reversed (shown opposite in figure 4). However, in this game the competition between Team B players allowed opportunistic advances by the randomly moving Team A players around cycle 150. The resulting pitch possession is shown in figure 6. At cycle 160, Team B manages to re-take the lost ground, but only because of unfortunate movements by Team A.
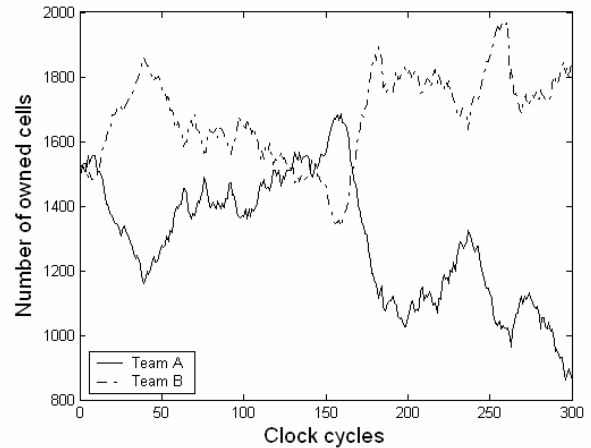

Figure 6. A-Random, B-Local player maximum

Team A local player maxima-Team B local player maxima
Initially team A takes most ground. Over time, possession evens out, as pairs of team mates compete against each other. The remaining two players, one each from Team A and Team B account for most of the changes. Play and possession of space oscillate after 130 ticks, with each team gaining the advantage at alternating ticks due to the action of the mixed pair. There is no obvious winner (Figure 7). The magnitude of oscillations in pitch possession occurring after cycle 80 are different for each team since Team B always moves last, and Team A only protects the neutral ground on every other move.
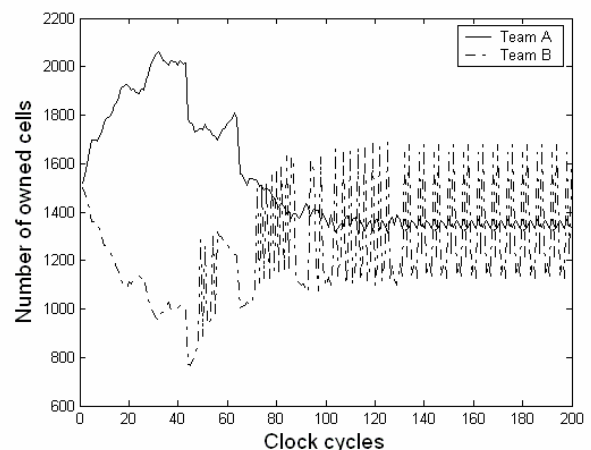

Figure 7. A-Local player max, B-Local player max

Team A local team maxima – Team B local player maxima
During the first 40 clock cycles, Team A rapidly takes ground whilst Team B competes against itself (Figure 8). The aggressive movements of Team A manage to split up the competing pairs in Team B, allowing them to gain ground between cycles 40-60. The lone Team B players then begin to compete with Team A players and mimic their moves, causing the constant difference in possessed areas from cycle 60. Although Team A wins overall, it does not maintain its promising initial success due to its beneficial effect on Team B's strategy (Figure 8).
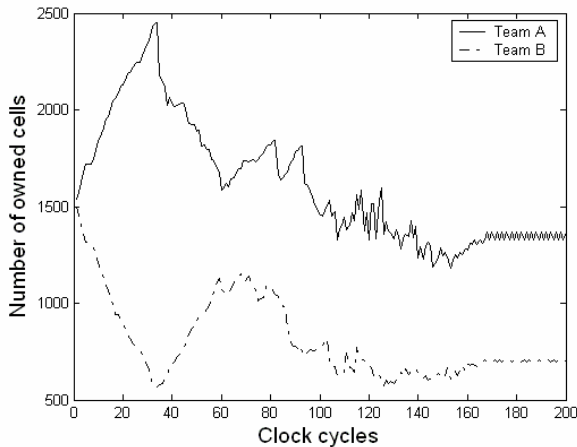

Figure 8. A-Local team max, B-Local player max

Team A Random Strategy – Team B local team maxima
Team B rapidly collects most of the pitch (Figure 9). The fluctuations shown in Team A's ownership in figure 5 are not repeated here. This is due to Team B moving second, so any fortuitous gains made by Team A are recaptured by Team B before the pitch possession is calculated. For the first 100 cycles, only two Team B players actively contribute to the cells owned by B. As they slow down, two more players begin to get drawn out, and fully join the attack at cycle100, causing the surge in pitch ownership.
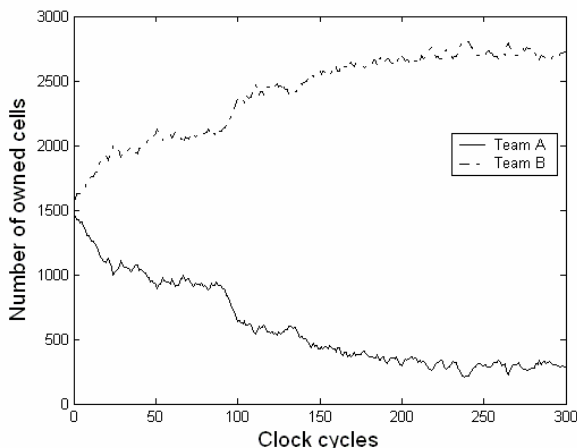

Figure 9. A-Random, B-Local team maximum

Team A local player maxima - Team B local team maxima
After an even start, team B eventually takes more ground (Figure 10). Mirroring the events shown in figure 8, Team A competes against itself allowing Team B to advance between cycles 30-40. Between 40-60 Team B breaks apart the Team A pairs, which then cause it to loose ground rapidly during cycles 60-80. From cycle 80 onward, Team A mimics the movements of Team B. Play reaches a stable state after 130 ticks, and Team B wins. Play is comparable to that given by swapping the strategies between A and B.
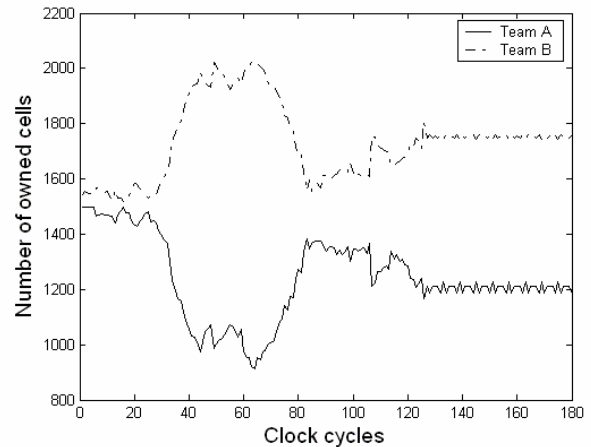

Figure 10. A-Local player max, B-Local team max

Team A local team maxima – Team B local team maxima
Team B slowly looses out to Team A. In this pairing, and with symmetrical starting positions, it is advantageous to move first. The two front Team A players move toward the edges of the pitch, forcing their opposite numbers into the centre. Eventually the movements of Team B bring a third Team A member into play, causing the sudden increase in A's superiority at cycle 20 (Figure 11). With an extra active player on Team A, it has no problems taking most ground.
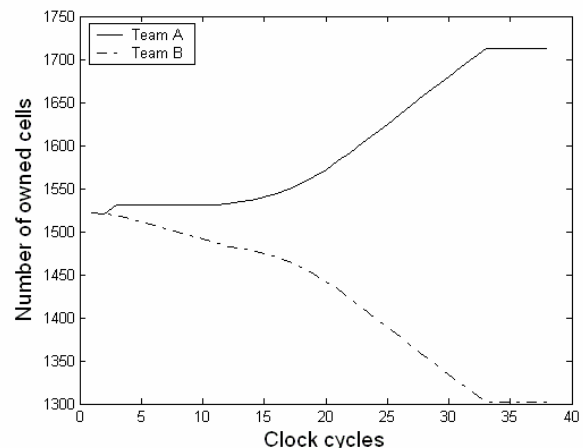

Figure 11. A-Local team max, B-Local team max

LTM is clearly the strongest strategy, winning five of the competitions. LPM comes in second, with three wins, and random comes in last. LTM teams are also most efficient, initially only using the foremost players to take ground.

When played against random strategies, the LTM exhibits herding behaviour. Players move to cells adjacent to the opposition, blocking them from moving forward. The majority of random opposition moves therefore have to be away from the player, resulting in them being herded into small pockets.

LPM teams are disadvantaged by their tendency to compete against home players. Against random strategies this can be unfavourable, resulting in random players herding LPM pairs. Surprisingly, it fairs better against LTM teams, due to the aggressive nature of the LTM splitting up the LPM pairs.

Future generations of strategies will focus on global, rather than local conditions. Following this, strategies involving planning will be considered. We will also update the simulation to enable both teams to move at the same time.

4. COMPARISON WITH GO

The space-time possession game described here can be compared with computer implementations of the game of Go [6]. The main difference between our game and Go is that the latter has more rules and more game-specific structure. In Go the objective is to surround the opponent's 'stones', or to surround contiguous sets of the opponents stones. We have suggested the space-time possession game to allow exploration of special features not tied to specific objectives such as surrounding pieces in Go, checkmating the king in chess, or scoring goals in robot soccer. Thus, although research into Go is highly relevant, the space-time possession game we have proposed is intended to allow the investigation of spatial structures in general.

Both Computer Go and Computer chess can be highly tactical, using properties of the particular pieces. Thus there are many set-piece openings and gambits which, currently, are unknown in the space-time possession game. It is our hope that the greater simplicity of the space-time possession game will give insights into spatial structures in general, rather than into spatial structures that are opportunistically useful in particular applications.

5. DISCUSSION AND CONCLUSIONS

Structuring the pitch in a soccer game is essential in that it provides a method for identifying good passes. A good pass might be to set up a shot on goal, to move the ball out of a dangerous situation, or simply to open up more passing possibilities. Strategic pitch possession is essential in implementing these tactics. In human football games, players try to structure the pitch by taking up positions to improve their team's chances of success. Although players take up these positions without explicitly communicating their intentions to their team mates, extensive training allows the players to recognise tactical opportunities based on these positions alone [5]. Furthermore, players can use their positions to weaken the opponent's by using feints, just as in the game of chess. To achieve the goal of beating the human world football champions, robot soccer will undoubtedly require this type of cooperation.

By forming a contiguous area of possessed pitch in the possession game, we identify a sequence of passes along which a ball could travel, so that it is always closer to home players than it is to opponents. Hence, we aim to structure a coordinated set of passes, similar to the way real footballers do. These ideas build on previous work to create a strategy as a string of tactics [5]. Continuation of the work will focus on linking any ball position to a predefined goal area, which can then be implemented as a passing strategy in robot football. Moreover, by simultaneously aiming to block opponent paths, we will also be forming defensive strategies.

The possession game provides a basic environment for research into many aspects of multi-agent systems, including centralised control, distributed control, planning, pattern recognition, machine learning, genetic algorithms, adaptive behaviour, and cooperation. Our aim is to develop a control architecture which combines both tactical and strategic operations in a single multi-agent controller.

3. REFERENCES

[1]   FIRA, (2004, 22 July), Mirosot, [Web page], http://fira.net/soccer/mirosot/overview.html

[2]   J. H. Johnson and B. A. Price, "Representing Patterns of Autonomous Agent Dynamics in Multi-Robot Systems", AROB-8, Japan, 2003.

[3]   C. Behring, M. Bracho, M. Castro, and J. A. Moreno, "An Algorithm for Robot Path Planning with Cellular Automata", Fourth International Conference on Cellular Autómata for Research and Industry ACRI'2000 in Theoretical and Practical Issues on Cellular Automata, 2000.

[4]   J. Thangavelauthma, T. D. Barfoot, and G. M. T. D'Eleuterio, "Coevolving Communication and Cooperation for Lattice Formation Tasks", 7th European Conference on Artificial Life (ECAL), Dortmund, Germany, 2003.

[5]   J. H. Johnson, "Visual communication in swarms of intelligent robot agents", *Artificial Life and Robotics*, **5**, 1-9, 2001.

[6]   M. Müller, "Computer Go", [Web page], http://citeseer.ist.psu.edu/cache/papers/cs/18669/http:zSzzSzwww.cs.ualberta.cazSz~mmuellerzSpszSzGo2000.pdf/computer-go.pdfbut