

THE VORONOI GAME IN ROBOT COORDINATION

James Law, Jeffrey Johnson

Department of Design and Innovation
The Open University
Walton Hall
Milton Keynes, MK7 6AA, UK

ABSTRACT

Cooperation and communication are two of the fundamental problems associated with research into multi-robot systems. In our work we examine the interactions of a team of robot footballers. Like computer chess and Go, robot football 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. As a research experiment we are investigating a ‘space-time possession game’ in which the only issue is the space controlled by the agents both individually and in combination. In this paper, we abstract a simple geometric representation of player interaction which facilitates both cooperation and communication through visual structures in the environment. We introduce the Voronoi game as a representation for the game of football and demonstrate the effectiveness of some competitive strategies.

1. INTRODUCTION

In a multi-agent system there invariably exist a number of observable spatial relationships between agents which can be linked to the objectives of that system. Often these relationships are secondary by-products of many complex interacting rules which define a task, however, they can also be governing rules themselves. For example, in a traffic system, drivers of vehicles maintain spaces between each other which are loosely based on the concepts of speed, thinking and braking distances. These change whether the vehicles are following each other along a road, or emerging from a junction. On the other hand, figure skaters must coordinate to perform set moves and holds, which are the focus of their routines. We are interested in these types of spatial relationship, and suggest that they can be used to create powerful multi-agent control strategies for use in complex and dynamic environments. Control structures for these types of environment are typically inflexible, suffering from over simplified definitions of the task.

Our interest stems from our involvement with robot football [1], particularly our participation in the competition of Mirobot [2]. Over recent years much progress has been made in the areas of robot design, machine vision, and control, and we are increasingly witnessing diminishing returns on the application of more advanced technologies. One area which is now becoming increasingly important, particularly with the recent introduction of 11-a-side competitions, is that of team strategy. Controllers have developed in stages as the rules of Mirobot have developed from 3, to 5, to 11 robots per team, and are typically role based [3-5]. By this we mean that a strategy is composed of a number of behaviours (goal keeper, defender, striker, etc.) which are grouped into ‘plays’. A simple role based strategy is shown in figure 1.

This type of architecture is easy to adopt in teams of five or less, when explicit roles can be defined for each player, but becomes awkward with the addition of more players. Roles are invariably repeated, leading to much

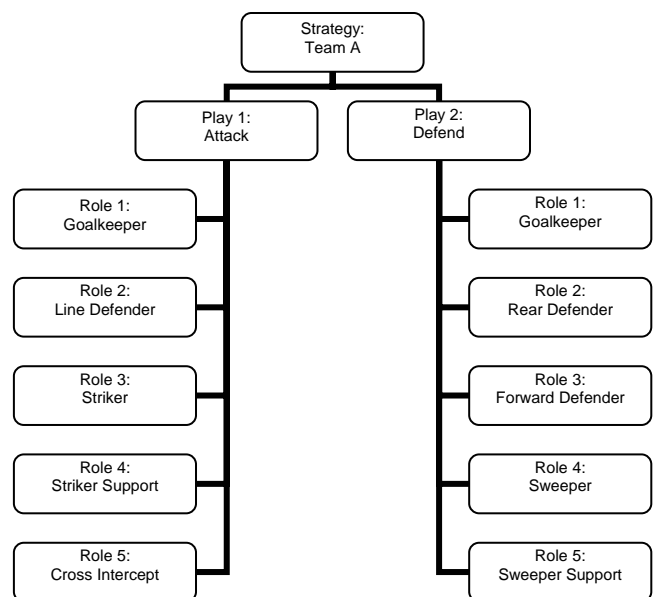


Figure 1: A simple role base architecture



redundancy, and their strict definitions make them repetitive and inflexible. So, how do we define strategies which are scalable across multiple team sizes, and that are also responsive to the environment? 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 these players, particularly how they are able to construe the soccer pitch, and find useful structure in what they see to guide their play [6].

This work focuses on competitive games, though the principles are relevant to any dynamic and competitive environment. Consider a team of fire-fighters tackling a forest fire. The fire grows and spreads as the wind changes. Given that each fire-fighter can only cover a limited area, and if there are only n fire-fighters, how do they coordinate their movements to restrict the growth of the fire and eventually extinguish it? This is a resource problem which requires dynamic manipulation of spatial structures. The solution is to distribute fire-fighters along the expanding edge of the fire, at the maximum separation as to stop the fire passing. As the wind changes, this line must continually reconfigure to trap localised pockets of fire, and restrict its growth in new directions. Such a situation can be expressed as a competitive game with similar characteristics to football.

2. SPATIAL STRUCTURES IN GAMES

We begin our research by considering the well known AI benchmark of computer chess, which is quintessentially concerned with structuring space [7]. This is illustrated in figure 2(a) where we give *names* to configurations of squares on the chess board. In figure 2(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.

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'



(a) Structured space in chess (b) The knight fork

Figure 2: Spatial relationships in chess

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.

On a higher level of complexity, and with more obvious reliance on spatial structures, is the game of 'Go'. In Go, players take turns to place coloured stones on a 19×19 position grid until both players pass. The objective is to surround the opponent's stones, or to surround contiguous sets of the opponent's stones, and to end owning the majority of territory once captured stones are accounted for. If chess has an estimated game tree complexity of 10^{123} then Go has a complexity of 10^{360} [8]. Traditional computer players for both games use game tree search algorithms. For chess these 'brute force' algorithms are comparable to human players, but they fail in Go, due to the larger tree sizes. Both computer Go and computer chess can be highly tactical, using properties of the particular pieces, and thus there are many recognised set-piece openings and gambits, which can reduce the overall complexity. We see these sequences and structures as discrete-time analogies to tactical plays in football.

Similarly, robot football can be considered to be a discrete game played on a rectangular grid. If the pixels of the vision system are seen as squares on the playing field, and turns are measured as frames, then the complexity can be calculated, similarly to chess and Go, as a game tree considering every possible move. At each turn, a robot can move anywhere within a circle, with radius proportional to its velocity. For a camera with a resolution of 640×480 imaging a 180×220 cm pitch at 30 frames per second, a robot moving at 1ms^{-1} can move to any of 558.5 squares. Therefore, with 10 robots on the pitch, 5585 possible moves can be made each turn. If a game lasts for two 5-minute halves, the number of turns, is 18,000 meaning the total complexity can theoretically reach $5585^{18,000}$. A number of factors reduce the effective complexity of the game, such as acceleration limits, obstructions and periods of inactivity, but brute force search algorithms are clearly inappropriate for problems of this magnitude.

Human footballers are experts at mastering space. They demonstrate remarkable skills in movement and perception, well beyond the current state of the art in robotics. Although they base their game on the skills and set pieces they practice before a match, the successful implementation of these tactics depends on the players' abilities to control space, to identify predefined plays from the positions of players around them, and create formations on the pitch to enable these plays. Players do not even need to touch the ball to be able to make a great contribution to their team.

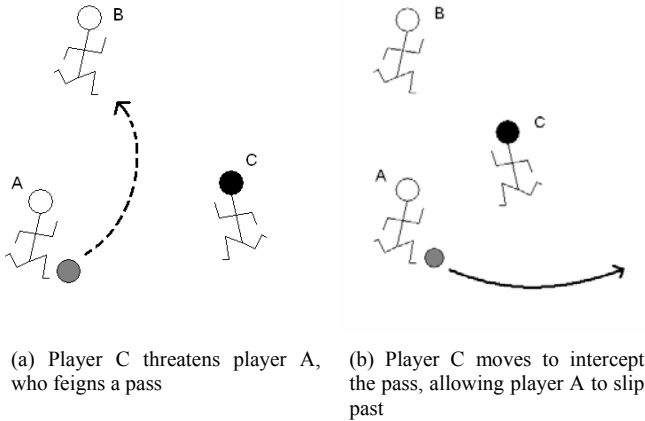


Figure 3: A set-piece in football

Consider the well-known set piece described in figure 3. Players A and B are attackers from the same team. Player C is an opponent defender, who threatens to tackle player A for the ball. If player A feigns a pass to player B, player C must move to intercept that pass. In doing so, player C moves out of position, and player A can slip past. We say that player B has drawn player C out of position. Human players find it relatively easy to spot these spatial structures, which enable players to cooperate in useful ways. In contrast, these spatial configurations are difficult to spot, and for the most part overlooked in robot football.

By representing these ideas in a form comprehensible to our robots, we aim to create a form of perception which will simplify the problem of controlling a team of cooperative agents, to one that is almost intuitive.

Our interest in these spatial configurations led us to develop the space-time possession game.

3. THE SPACE-TIME AND VORONOI GAMES

In [9], we separated the concept of spatial representation from the game of football. The result was the space-time possession game, a cellular automata in which two teams of agents competed to control space on a 2-dimensional pitch. In the game, the pitch is divided up into cells, each of which is owned by the closest agent (player space), and, by extension, that agent's team (team space). By outmanoeuvring the opposition, it is possible for one team to control a larger area of the pitch than that of the opponent. Results from this work showed that a team in which agents cooperated outperformed a team composed of non-cooperating individuals.

An extension of this work was given in [10], where a bounded Voronoi diagram was used to analyse the change in team space during a simulated robot football game (figure 4). The results showed a clear correlation between the state of play and the overall area each team controlled on the pitch. From these findings, we concluded that robot football can be represented as a game of spatial competition. We propose to generate an abstracted

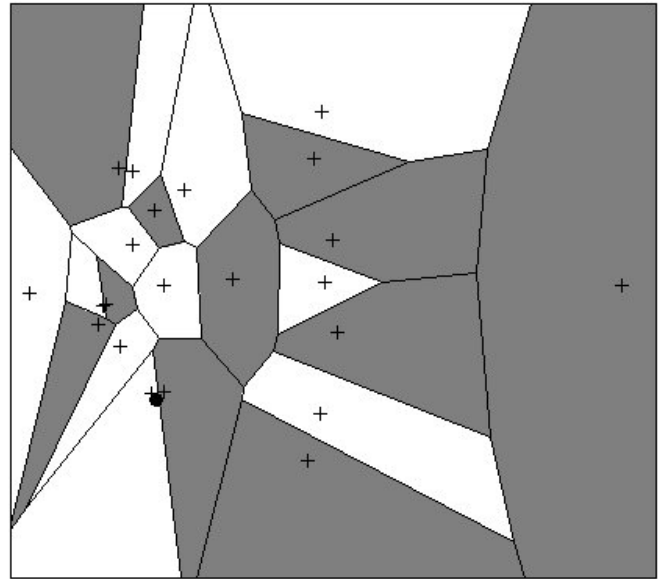


Figure 4: Voronoi analysis of a frame in simulated football

strategy composed of movements which cause players to occupy key spaces, open up areas for moving the ball, block opponent players and passes, and set up shots at goal. The actual passes and ball movements are considered to be separate events made possible by these underlying spatial relationships.

In both experiments we assumed that every player was omni-directional and could move with the same velocity and acceleration. In real systems this is not the case, and so a weighted Voronoi diagram is required. However, the principles under investigation here relate to both types of diagram, and so we examine the more general case.

Having identified the relevance of spatial possession in robot football, we continue by investigating suitable structures, that effectively compete for space, to form the basis of our abstracted team strategies. We take inspiration from a similar set of problems called Voronoi games.

The one-dimensional Voronoi game was introduced by Ahn et al. [11]. In this work, used as a model for competitively placing facilities along a road, players take turns to place n facilities on a line or circle (figure 5). The game is composed of n rounds, each player placing one site in each turn. At the end of the game, the arena is subdivided into sections according to the nearest neighbour rule, and the player with the largest area wins. The analysis provides a set of rules for placing sites, which enables the second player to force a win in every game. A modified one-round circle game, is presented where each player places their n sites in one turn. It is shown that he first player can force a win by placing sites on the odd integer points $\{1, 3, \dots, 2n - 1\}$.

Cheong et al. [12] extend this one-round Voronoi game to two or more dimensions. In this game, which is similar to our space-time possession game, a piece controls the area of pitch P closer to it than any other piece. Player one, white, places a set of pieces W , which is followed by player two, black, placing a set of pieces B . When all

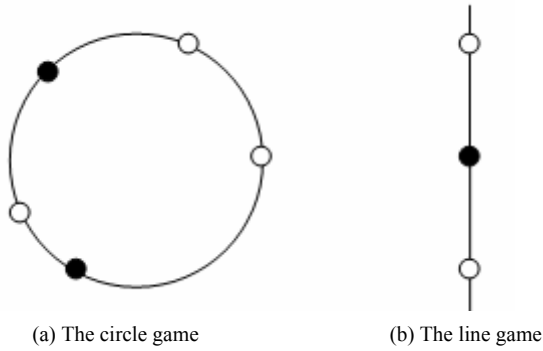


Figure 5: The one-dimensional Voronoi games

pieces are placed, the Voronoi diagram of $A \cup B$ is constructed, and the player which owns the largest area of P is declared the winner. Cheong et al. show that given certain criteria, the second player can always steal at least half of the pitch. This proof is extended by Fekete and Meijer [13] who show that for a rectangular pitch of aspect ratio ρ , black has a winning strategy for $n \geq 3$ and $\rho > \sqrt{2/n}$, and for $n = 2$ and $\rho > \sqrt{3/2}$. White wins in all remaining cases. It should be noted that, these strategies all require white to place its pieces on a rectangular grid.

It follows that the space-time game as applied to robot football can be considered as a multi-round Voronoi game, with some constraints: Players place pieces simultaneously, with only knowledge of their opponents' previous positions. Pieces have a limited movement between turns, and must remain within moving distance of their last position. There are no specific winning conditions, the aim being to continually out manoeuvre the opponent, and some areas are more important than others (such as the ball and goals).

In the following sections we propose some useful spatial formations, and examine their application to the one-round Voronoi game. Movement of pieces across multiple rounds will be examined in our further work.

4. SPATIAL STRATEGIES

Figure 6 shows the Voronoi cell of an opponent piece, O . h , is a home piece, and p is the centroid of the Voronoi polygon with area A . Our goal is to steal the maximum possible area from O . If p and O are coincident, then there is no position for which the area of A closer to h is greater than that closer to O . However, if p and O are not coincident, then placing h on the line between p and O will cause h to capture to a larger proportion of A , as shown. This can be considered as a *strong marking* strategy. Provided opponents do not lie on the centre of their Voronoi polygons, then it is always possible to steal a slightly greater area of the pitch from the opponent team using this technique.

This type of strategy is called a *takeover* and a variation is proposed by Cheong whereby two home pieces are allocated to the $n/2$ opponent pieces holding the largest

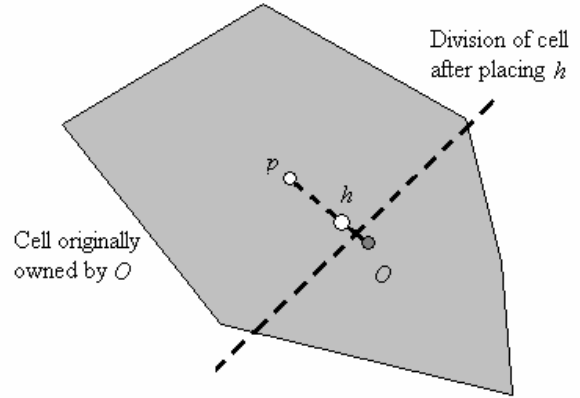


Figure 6: Voronoi cell capture

areas. By placing two pieces close to, and on opposite sides of O , the home player captures almost the entire area of A . Provided the areas of all the opponent pieces are not similar, this takeover strategy will capture at least half of the playing field.

These strategies are suitable for strongly marking n opponents with m pieces if $m \geq n$. However, if $m < n$ (say we have already allocated two home pieces to the largest opponent cell), the best strategy for the remaining m players may be to spread them between the n opponents, and *weakly mark* multiple opponents. Consider figure 7. Three opponent pieces O_i describe a Delaunay circle D , centred on a Voronoi vertex (not shown). A piece h placed within that circle will form a new Voronoi cell which neighbours the cells of O_i (as shown). If $i > 3$ and O_i are not co-circular, then multiple Delaunay circles can be constructed. If a piece is placed within the overlapping segment of two or more Delaunay circles, such as that between O_2 and h in figure 7, then its Voronoi cell will neighbour those of all the pieces lying on those circles. In general, the more Delaunay circles enclosing a point, the more neighbouring pieces that point will have. Also, the larger the radius of those circles, the further away those neighbours will be, and the larger the Voronoi cell associated with that point. As a stand-alone strategy, we position pieces on the centre of the most frequently overlapped segments, or the segments overlapped by circles with the largest cumulative radii.




Figure 7: White neighbours multiple opponents

5. EXPERIMENTAL RESULTS

Each of the above strategies was played second in 100 one-round Voronoi games, each using teams of 5 players on a 64×48 arena, against randomly positioned opponents. The seeds for the random position generator were set for each set of games to ensure all strategies were played against the same set of random opponents. As a control, games were also played using a random strategy, and a brute force best-position search algorithm. A brief description of each strategy and its outcome is given below, and statistical results are shown in figure 8. The total playing area is 3072, with the dashed line at $x = 1536$ indicating the half pitch area. Scores above this are counted as wins, and scores below are losses. Tabulated results for the mean, median and standard deviation of scores are given in table 1.

Random: Pieces are placed at random. As would be expected, there is a normal distribution of area captured over 100 games, with a mean score of 1593.6, which is within 2% of the half pitch area.

Optimal: A brute force search of all integer coordinates for positions which give the greatest returns. This strategy always wins in our tests, with a confident margin over the opponent. However, the lengthy computation makes it impractical for real-time applications or large pitch sizes. It is included as a benchmark for our other strategies.

One on one: A strong marking strategy with each piece paired with a single opponent. This is a very competitive strategy, giving results with a mean within 0.9% of our optimal benchmark strategy, but using a much simpler search algorithm. The spatial structures employed here are very different from those observed in the optimal strategy, but produce very similar outcomes. These particular structures perform most competitively in situations where each opponent piece controls a similar sized area.

Two on one: A strong marking strategy, using two pieces to mark each of the strongest opponents. The remaining piece is allocated to the 3rd strongest opponent. Again, this is a strong strategy, consistently winning all 100 games, and with a mean falling within 5.4% of that of our optimal benchmark strategy. These structures perform best against opponents where some pieces occupy more space than others.

Overlaps: A weak marking strategy. Pieces are placed at the centre of the most overlapped Delaunay segments. Effectively these configurations place pieces as to neighbour the maximum possible number of opponents. Although not as competitive as the strong marking strategies, this approach still wins in 72% of the games. A main benefit of this structure is its flexibility. The two strong marking strategies require pieces to be very close to the opponents at all times. To change between the one and two marker strategies requires single pieces to make relatively large movements, which will take time to perform. An advantage of both of the weak marking configurations is that pieces are well distributed amongst

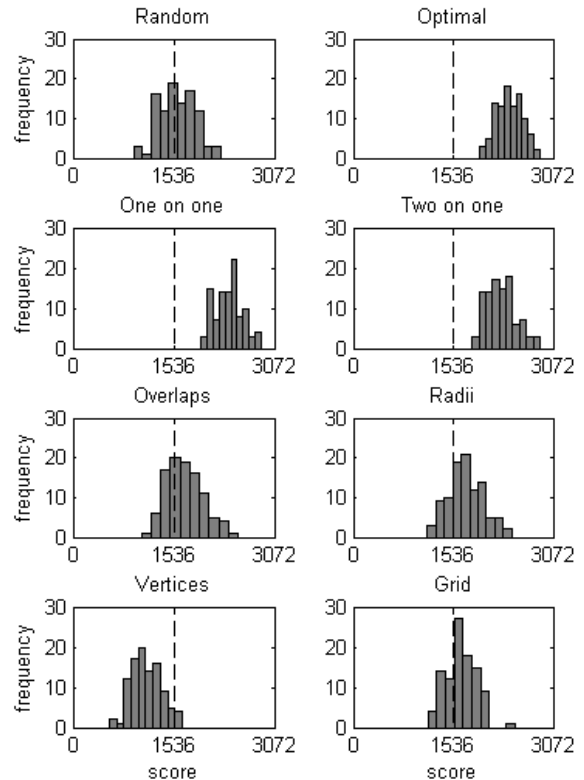


Figure 8: Strategy scores over 100 games

the opponents, allowing an easy switch between strategies.

Radii: A weak marking strategy. Pieces are placed at the centre of overlapping segments with the highest cumulative radii of enclosing Delaunay circles. Effectively pieces are placed in large but highly neighboured spaces. Again, this performs less well than the strong marking strategies, but out performs the ‘Overlaps’ strategy with a 2% higher mean score, and 6 more wins..

Vertices: Players are placed at the furthest points from all opponent pieces, i.e. on the most remote Voronoi vertices. This is a control experiment to demonstrate a poorly abstracted spatial structure. Intuition may suggest that by simply moving pieces far from their neighbours, they will move into and occupy large empty spaces. Losing 95 of the games indicates that this is not the case.

Grid: The first player strategy proposed by Fekete and Meijer for $n \geq 3$. In [13] they indicate positioning pieces on regular grids minimises the gains of an opponent. Here we implement the $1 \times n$ grid, and demonstrate its performance on a pitch of aspect ratio $\rho > \sqrt{2/n}$ (Fekete and Meijer propose it as a winning strategy only if $\rho \leq \sqrt{2/n}$). The outcome is much worse than our strong marking strategies, with only 73 wins and an average possession of 53.9%, making it more comparable to our weak marking strategies. A drawback of this style of play as applied to an N -round game is its inability to adapt to the changing configurations on the pitch.

Strategy	Mean Score	Median Score	Standard Deviation
Random	1593.6	1569.3	270.1978
Optimal	2399.1	2392.7	204.0187
One on one	2378.5	2385.9	207.3706
Two on one	2270.9	2261.1	228.1404
Overlaps	1696.2	1666.0	274.0008
Radii	1730.9	1686.9	271.2585
Vertices	1123.6	1101.8	227.3849
Grid	1654.4	1633.4	229.7723

Table 1: Statistics for each strategy

6. CONCLUSIONS

Based on the established research on competitive games, we have introduced the concept of spatial representation as a basis for multi-robot control in dynamic and unpredictable environments. Using robot football as an example, we have highlighted the lack of flexibility of traditional control architectures, and shown that the complexity of the task is too great for standard AI search techniques. Through analysis of simulated football matches, we have identified the significance of controlling areas of pitch, and have created an abstracted generalisation of football in the form of an N -round Voronoi game.

From our knowledge of human football, strategies for the one-round Voronoi game, and analysis of Delaunay and Voronoi structures, we have identified a set of spatial structures which correspond to ideas we consider to be useful in spatial competition. These structures can be easily identified from the positions of local players, and provide an adaptive strategy for player positioning.

Using the one-round Voronoi game as an experiment, we show how our spatial configurations respond to a set of opponent positions. The results indicate that the structures we have identified are at best near-optimal, and at worst, above average, and all more competitive than some arbitrarily chosen configurations. We hypothesise that while our strong marking strategies perform best in these games, a combination of the strong and weak marking configurations will be more appropriate for the N -round game and, by extension, robot football. This is being examined in our ongoing work.

As well as using general structures for coordinating players at a strategic level, we can predict that similar relationships will be useful in tactical plays, such as passing a ball or scoring a goal. By identifying a hierarchical set of spatial structures, and applying them to a dynamic environment, such as robot football, we aim to create an adaptive control architecture which is based on the dynamics of an environment, and not a set of pre-defined, restrictive roles.

7. REFERENCES

- [1] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, and H. Matsubara, "RoboCup: A Challenge Problem for AI.," *AI Magazine*, vol. 18, pp. 73-85, 1997.
- [2] P. Robinson, J. C. Wolf, J. A. Law, J. D. Oliver, K. W. Young, and D. A. Harewood-Gill, "Recent Mirobot Developments in the UK," presented at FIRA Robot World Congress, Busan, Korea, 2004.
- [3] J.-H. Kim, K.-C. Kim, D.-C. Kim, Y.-J. Kim, and P. Vadakkepat, "Path planning and role selection mechanism for soccer robots," presented at IEEE Int. conference on Robotics and Automation - ICRA'98, Leuven, Belgium, 1998.
- [4] H. L. Sng, G. Sen Gupta, and C. H. Messom, "Strategy for Collaboration in Robot Soccer," presented at IEEE DELTA, New Zealand, 2002.
- [5] M. Bowling, B. Browning, and M. Veloso, "Plays as effective multiagent plans enabling opponent-adaptive play selection," presented at International Conference on Automated Planning and Scheduling (ICAPS'04), 2004.
- [6] J. H. Johnson and B. A. Price, "Representing Patterns of Autonomous Agent Dynamics in Multi-Robot Systems," presented at AROB-8, Japan, 2003.
- [7] R. H. Atkin, W. R. Hartston, and I. H. Witten, "Fred CHAMP, positional-chess analyst," *International Journal of Man-Machine Studies*, vol. 8, pp. 517-529, 1976.
- [8] B. Bouzy and G. Chaslot, "Monte-Carlo Go Reinforcement Learning Experiments," presented at IEEE 2006 Symposium on Computational Intelligence in Games, Reno, USA, 2006.
- [9] J. Law and J. Johnson, "Discrete Dynamics and the Space Time Possession Game," presented at FIRA Robot World Congress, Busan, Korea, 2004.
- [10] J. Law, "Analysis of Multi-Robot Cooperation using Voronoi Diagrams," presented at The 3rd International RCL / VNIITransmash Workshop on Planetary Rovers, Space Robotics and Earth-Based Robots, St. Petersburg, Russia, 2005.
- [11] H.-K. Ahn, S.-W. Cheng, O. Cheong, M. Golin, and R. Oostrum, "Competitive Facility Location along a Highway," presented at Computing and Combinatorics: 7th Annual International Conference, Guilin, China, 2001.
- [12] O. Cheong, S. Har-Peled, N. Linial, and J. Matousek, "The One-Round Voronoi Game," presented at 18th Annual ACM Symposium on Computational Geometry, Barcelona, Spain, 2002.
- [13] S. P. Fekete and H. Meijer, "The one-round Voronoi game replayed," presented at 19th European Workshop on Computational Geometry, Bonn, Germany, 2003.