

Aberystwyth University

Analysis of multi-robot cooperation using Voronoi diagrams

Law, James

Published in:

3rd International RCL/VNIItransmash Workshop on Planetary Rovers, Space Robotics and Earth-Based Robots

Publication date:

2005

Citation for published version (APA):

Law, J. (2005). Analysis of multi-robot cooperation using Voronoi diagrams. In *3rd International RCL/VNIItransmash Workshop on Planetary Rovers, Space Robotics and Earth-Based Robots*

General rights

Copyright and moral rights for the publications made accessible in the Aberystwyth Research Portal (the Institutional Repository) are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Aberystwyth Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Aberystwyth Research Portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

tel: +44 1970 62 2400
email: is@aber.ac.uk

Analysis of Multi-Robot Cooperation using Voronoi Diagrams

James Law

J.A.Law@open.ac.uk

Department of Design and Innovation

The Open University

Milton Keynes, MK7 6AA, UK

Abstract - Achieving cooperation in a team of autonomous agents is a complex task. To achieve success requires each agent to have a well grounded understanding of the task and how to achieve the objectives. As a basis for designing new multi-agent control architectures, we introduce the notion of spatial representation in a group of autonomous football playing robots. Using Voronoi diagrams to examine data from simulated football matches, we identify a correlation between our spatial structures, and the events during match.

I. INTRODUCTION

A common question in robotics and AI is “how do we get a group of agents to self organise?” Humans have an acute ability in this area, whether as part of small autonomous teams or large hierarchical organisations. We are capable of forming both practical workgroups to undertake tasks on the spur of the moment, and coordinating by committee to plan for future events. Our ability to cooperate is one that governs our whole society. In robotics we endeavour to reproduce these complex relationships to create robotic teams that are capable of a wider variety of tasks, and more effective operation. Given this objective, how can we get robots to plan, divide workloads, coordinate, and interact to make teams more capable than the sum of their parts?

We approach these problems through our interest in robot soccer [1], in which robots compete in teams against one another to score goals. Traditionally the aim of these competitions has been to drive forward robotic technology, with particular focus on machine vision, hardware, communications and low level control. Now, teams across the globe are reaching a stage where there are diminishing returns in having the fastest robots, highest resolution vision system, or most accurate control algorithms. We have now reached the point when robots need to use teamwork to win.

Robot football provides us with an excellent foundation on which to investigate the challenges of implementing cooperation. It encompasses the traditional problems associated with multi-robot research, such as machine vision, communication, task allocation and planning, but does so in a highly dynamic and competitive environment, which is unusual in most other team working applications. The aim, to score goals against an opponent team is easy to express and understand, and has the added benefit of allowing direct measurement of success against alternative control methods.

In contrast, we can also reflect on aspects of human football and compare their performance with that of our robots. Specifically we are interested in the complex relationships human players form on the pitch.

A. Mirobot Robot Football

We are involved in the competition of Mirobot robot football, as regulated under the association of FIRA (Federation of International Robot-soccer Association), and with our partners at the University of Plymouth and the University of Warwick, we represent the UK at international competitions [2]. The Mirobot leagues are played with teams of 3, 5, 7, or 11 robots. In this paper we will focus on the most demanding of these, 11-a-side.

The rules of 11-a-side Mirobot state that robots should be less than 7.5cm*7.5cm*7.5cm in size, and play with an orange golf ball on a black pitch, 440cm*280cm. Each team of robots is controlled from a central PC through a camera mounted directly above the pitch (Figure 1). Robots are identified by a colour patch on their top sides, and controlled remotely from the PC. The vision system identifies the position and orientation of each robot, and the position of the ball, and passes this information to the strategy software, which calculates the actions of each robot and broadcasts them via a radio link. Although essentially a centralised control system, we are developing a distributed architecture, simulating the multi-agent dynamics within the PC.

Strategies in Mirobot robot football are currently based on hierarchical roles, plays and strategies [3-5], with the high level strategy layer governing the overall direction of play. Based on factors such as ball state, player state and play




Figure 1: Mirobot robot football system

feedback, the strategy level selects plays from a precompiled playbook. Each play contains information concerning the role of every robot in that play, additional specific role information, and any sequencing or rules for switching roles. At the lowest level, the role layer holds pre-programmed positions and actions for each robot. Roles are usually defined for the duration of a play, such as goal keeper, defender, and sweeper etc. but they can also be switched or assigned temporarily to extend functionality [6]. An example would be for a kick off, where robots might be given a sequence of movements to perform before defaulting into their main role. Simple role, play and strategy structures are given in Figure 2.

Strategies may consist of as few as only one or two plays. For example, a one-play strategy would simply use the same number of attackers and defenders throughout a match, perhaps with some role switching. A two-play strategy would consist of a defensive play, for when the ball is in the home half, and an attacking play, for when the ball is in the opponent's half, with different numbers of attacking and defending players. Although sufficient for 5-a-side matches, the recent introduction of 7 and 11-a-side Mirosoft leagues has highlighted weaknesses in this approach. A team with 11 robots requires a much wider variety of roles and more complex plays. For human coaches this is a difficult mental task, and the mathematical complexity causes problems for automatic role selection algorithms.

Furthermore, this type of architecture is limited in its

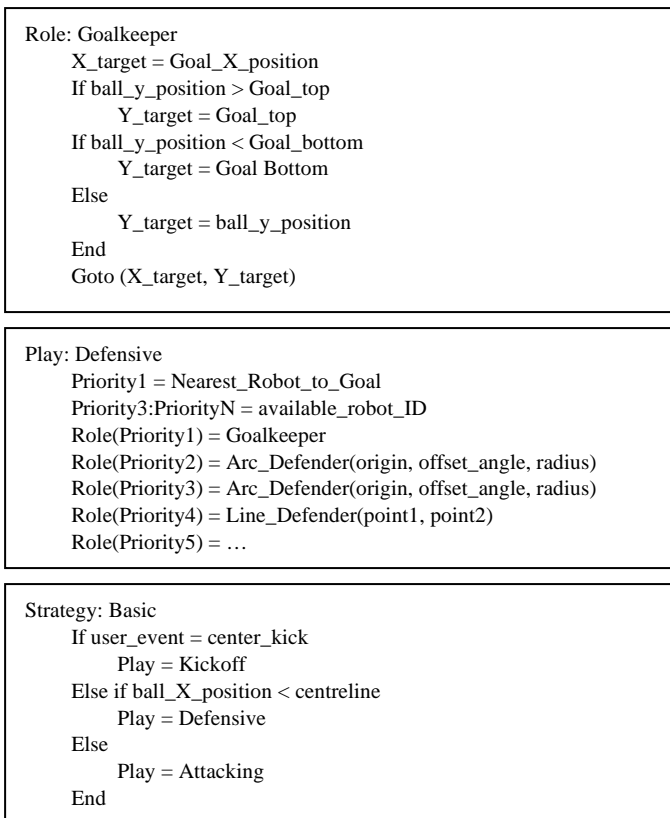


Figure 2: An example of simplified role, play and strategy functions

ability to adapt, and tends not to incorporate true cooperation. Any apparent cooperation is usually short lived and the effect of a pre-programmed set piece, such as a kick off, where robots might be issued with a set sequence of passes and moves. At other times, players tend to work collaboratively, working toward the same goal, and supporting one another, but there is no explicit cooperation between them. Each player, in this case, is usually only a back up in case a play goes foul.

An exception to this rule is the work of Bowling et al. [7] Their architecture holds a number of alternative plays, and switches between them during a game by using an on-line selection algorithm, which adapts to the playing style of the opponents and promotes plays which produce better outcomes. However, the plays themselves are still pre-programmed, with any cooperation confined to set pieces as described above.

B. Multi-Agent Architectures

Cooperative multi-robot architectures are mainly focused on two areas: Task allocation and task execution. Task allocation is the autonomous division of a task inputted by a user, between available robots. Schemes for task allocation are based on several approaches, the most popular being:

- Auctions and markets – based on the contract net protocol. An auctioneer agent broadcasts a task for execution. Each agent makes a bid for that task based on its estimated costs. By holding a number of rounds of auction, a near optimal solution can be found. In some systems it is possible to sub-auction tasks [8-10].
- Voting – each agent votes on which task it should perform based on its perception of the situation and its abilities [11, 12].
- Motivation – agents are motivated by concepts such as impatience and acquiescence. If an agent spends too long doing a task (which may indicate robot failure) another agent will become bored of waiting, and seize the task. Similarly, if a robot repeatedly fails to perform a task, and can sense its failure, it will retire from that job and search for another [13].

Task execution is the problem of carrying out a task shared among robots. For example, a popular problem is two or more robots moving an object too large or heavy for one robot to handle. Schemes for task execution include:

- Behaviour exchange – a behaviour-based approach where high level functionality is composed by coordination of more basic behaviours distributed across a set of robots. Behaviours can interact across the network, allowing sensors on one robot to drive motors on another [14].
- Leader-follower – One agent becomes a team leader and takes responsibility for the entire task, delegating commands to other agents for the duration of its leadership [15].
- Markets – similar to the auction method of assigning tasks, but using shorter duration tasks and holding auctions more frequently. This is adequate for tasks

where jobs are not closely coupled, i.e. where jobs are not closely orchestrated simultaneously [16].

All the above schemes require well defined tasks, such as move object A to point B, or assign destinations D_1 , D_2 , and D_3 to robots R_1 , R_2 , and R_3 . Although the end result of robot football, to score goals, is well defined, how to achieve them is not so obvious. At a simple level, we could say it is simply to move object BALL to point GOAL, but of course there are many more factors involved. How do we move BALL to GOAL whilst avoiding OPPOSITION? How do we know when to pass, and when to shoot? Before we can tackle the problem of coordination in a team of robots, we need to be able to clearly identify the tasks required to reach their objective. The role selection protocols used in current robot football architectures are based on an over simplified view of the game, and so our aim is to find new and more comprehensive ways of representing the robot football task.

This, then, is our problem. How can we represent the intricate requirements imposed upon a team in a game of robot football? Can we identify structures and sub-goals which simplify the overall objective of robot football into comprehensible tasks?

In this paper, we describe our experiments in representing the task of robot soccer at a team level by analysing the spatial distribution of agents using Voronoi diagrams.

C. Spatial Awareness and Perception

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. 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




Figure 3: The two-on-one configuration

pass to player B, player C must move 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 under utilised in robot football.

By representing these ideas in a form comprehensible by our robot footballers, we aim to create a form of robot 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.

II. THE SPACE-TIME POSSESSION GAME

In our earlier work [17], 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 outmaneuvering the opposition, it is possible for one team to control a larger area of the pitch. Results from this work showed that a team in which agents cooperated outperformed a team composed of non-cooperating individuals.

In these initial experiments we used image processing techniques to evaluate the ownership map. The processing time for this analysis increased exponentially with pitch size, leading to a slow turn rate on pitches representing real world camera resolutions. Although not critical in the requirements of the space-time possession game, a faster method is required to analyse the images captured in a robot football system. Predicting these problems, we improved our methods, implementing an analysis based on Voronoi diagrams.

A. The Voronoi Diagram

A standard Voronoi diagram is composed of n tessellating convex hulls, or Voronoi polygons, which are defined as follows. Consider a set of points in a plane, $P = \{p_1, p_2, \dots, p_n\}$. For any point p_i , there exists a locus of points (x, y) in the plane that are closer to p_i than any other point in the set P . These loci form the Voronoi polygons, which we have referred to previously as a player's space. In ordinary Voronoi Diagrams, boundaries between external points in P stretch of to infinity, and have infinite area. Since we are only interested in a player's space within the perimeter of the pitch, we constructed a bounded Voronoi diagram, in which the pitch boundaries were added to all open Voronoi cells (Figure 4).

Since the Voronoi Diagram is calculated directly from the positions of the players, it is much more efficient in segmenting large areas than the cellular automata, which analyses the empty spaces. This means a Voronoi Diagram can be used on much larger pitches and create higher resolution divisions of space. The complexity of the Voronoi diagram varies with the number of agents, whereas the cellular automata varies with pitch size. Since we will be using a




Figure 4: Bounded Voronoi diagram showing players and their associated areas

relatively small number of agents in comparison to the size of the playing area, the Voronoi diagram is the more appropriate option for our task.

Figure 5 shows the time taken to run 1000 game cycles using both the cellular automata and the Voronoi versions of the possession game, with players making random movements. For large numbers of players on a small pitch, the cellular automata gives the fastest response, but its speed rapidly decreases as pitch size is increased. Conversely, the Voronoi method is much faster, slowing only as large numbers of players are introduced.

Mirosot robot football is typically played using images of 640*480 pixels, at 30 frames per second (fps). For 11-a-side games, the Voronoi method can analyse one frame, or game cycle in $22.9 \cdot 10^{-3}$ s, or a rate of 43fps, within the requirements of our robot football system.




Figure 5: Speed comparison of the cellular automata and Voronoi possession games

III. ROBOCUP SIMULATION LEAGUE ANALYSIS

Using our new model of spatial perception, we turned back to the game of football to investigate the dynamics of team space during a match. Our aim was to find whether a relationship exists between the distribution of team space and the states of play during a football match. Similar work by Kim [18], examined player space with relation to victory conditions in a simulation of real football. He concluded that to win, a team did not necessarily have to control the largest area on average during a match, but that in order to score a goal, a team did need to be in control of a larger area of pitch at that moment. Our experiments differ from Kim's in that we are not only examining victory conditions, but searching for relationships which exist throughout a match.

We based our own tests on data from the RoboCup simulation league. In this league, teams of 11 agents compete in a simulated environment, based on the RoboCup small sized league, which is itself similar to Mirosoft. The simulator logs information about the agents and the ball, for every match played, and records are published on the internet. Using our possession game, we examined seven different matches, representing a variety of winning conditions, and observed the changes in team space, player space, the goals scored, and the position of the ball during play.

A. Team Space Analysis

We began by measuring the amount of pitch owned by either team in each of the seven matches, and compared their average ownership to the number of goals scored. The results are shown in Table 1. In Match 5 we observed the second largest goal difference of any match, and the largest average margin in pitch possession by the winning team. However, the largest goal difference is in match 2, which has one of the smallest average pitch possession margins. Examining the relations between goal difference and pitch margin for the remaining matches, it is difficult to suggest that controlling the majority of the pitch is sufficient to win a match.

We furthered this research by analysing the change in possession throughout game 5. By monitoring and recording key events, such as an intercepted pass, we formed relationships between our definition of team space and the changing state of the game. Figure 6 shows how often team A controlled specific quantities of pitch.

The total area of the pitch is 7140 units, and we can clearly see that team A mainly controls only a fraction of this,

TABLE I
AVERAGE POSSESSION SCORES IN SIMULATED ROBOT FOOTBALL MATCHES

Match	Score (A-B)	Average team possession as % of pitch		Average possession difference	Maximum team possession as % of pitch	
		A	B		A	B
		1	0 - 0	44.89	55.11	10.22
2	10 - 0	51.73	48.27	3.46	71.54	78.59
3	1 - 2	48.39	51.61	3.22	75.41	75.86
4	0 - 0	54.48	45.52	8.96	74.95	79.13
5	0 - 6	42.14	57.86	15.72	76.01	79.07
6	4 - 3	54.23	45.77	8.46	80.36	74.68
7	3 - 0	51.28	48.72	2.56	76.76	74.81




Figure 6: Frequency with which team A controlled areas of pitch

around 2500 units. This low ownership is reflected by the team B being in possession of the ball for 78% of the match, and team A playing defensively. It should be noted that the significant feature relating to a possession score of 3600 is an effect of the time spent in the kickoff position after each goal is scored, and as such is not a proportional representation of team A's influence.

From the simulations, we observe that larger team spaces are usually linked to attacking plays, and smaller ones to defensive plays. In terms of spatial configurations, a large team space facilitates easier passing and movement to intercept stray balls, which is desirable in an attacking formation. In contrast, small team and player spaces indicate tight configurations of players, which are better for protecting a small area and intercepting passes and shots in that region.

However, as concluded earlier, it is not sufficient to state that by controlling more space a team is more likely to score goals. Neither is it appropriate to state that a team in control of a larger area will be on the attack. For either of these to hold any merit, the team in question must be in possession of the ball.

B. Movement on the Ball

We examined the relationships between team space and ball position. Figure 7 shows ball position data, and space distribution from a portion of match 5. The area plot is taken from the perspective of team A, with equal control between teams indicated by a magnitude of zero. Positive values indicate more control exerted by team A, and negative values indicate increased control by team B. The plot of ball x-position has been rescaled in amplitude, but has not been altered frame wise. The x-axis for ball position is defined as the line passing through the centre of both goal mouths.

A relationship can be seen between the two plots, with both having similar major features. These features appear in close phase to one another, with a difference varying between 20 and 30 frames. This relationship between ball position and team space is evident throughout each of the seven simulated matches. The relation between these two signals is more




Figure 7: Comparison of ball movements and controlled space

clearly shown in Figure 8, which depicts the spectra of the two signals over the entirety of the match. From these results, it is clear that there is a relationship between team space and ball position.

Our experiments indicate that once a team has possession of the ball, it adopts a broad spatial configuration, which facilitates passing and safe movement about the pitch. At the same time, the opposition forms a much tighter spatial configuration to protect specific areas of pitch, or block opponent players. As the ball is moved further toward the goal, the attacking players increase their control over the pitch, whilst their opponents form tighter, more defensive structures around their home goal. The phase lag between the signals in Figure 7 is due to the reaction times of the robot football system.

An example of this spatial structuring is shown in Figure 9. Here, team A (controlling the area in white) have the ball and are attempting to shoot at team B's goal, on the right hand edge of the image. Team B (grey area) has responded by forming a tight defensive structure around the player in possession. To keep the ball, team A must make a pass to a




Figure 8: Spatial control and ball variance spectra




Figure 9: A tight defensive configuration restricts an opponent's passing opportunities

safe, team controlled area (indicated by the arrows), which is made difficult by the intermediate opposition space. By focusing on the key area surrounding the ball, team B control more useful areas, enabling them to intercept any passes, or stray balls in the region. In contrast, the spatial configuration of team A, although covering more, gives it little influence in the region of interest, making it difficult to safely pass the ball. In this instance, Team A fails to make a pass, and team B gains possession of the ball.

IV. CONCLUSIONS

In this paper we have introduced spatial representation as a technique for control and analysis of teams of dynamically interacting agents. We have identified robot football as a task with interesting spatial structures, and highlighted the shortcomings of typical control strategies in this field. After considering popular multi-agent architectures, we conclude that a better representation of the game is firstly required.

Taking inspiration from the spatial perception of human footballers, we introduce the Voronoi diagram as a technique for representing a teams control over the pitch, and investigate the ideas of player space and team space in simulated robot football matches. Our results show a strong correlation between the movements of the ball, and the distribution of team space, whilst also proving that simply controlling more of the pitch is not sufficient to win a match. We show that the control a team exerts over the pitch is proportional to its attacking or defending stance, and demonstrate the importance of controlling specific areas of useful space.

The findings of this work provide us with a basis for forming more flexible team strategies. Using the ideas presented here, we will implement a high level strategy on our MiroSot system.

REFERENCES

- [1] J. Johnson and P. Irvani, "Robotics in the Emergence of Complexity Science," presented at The Ninth International Symposium on Artificial Life and Robotics, Beppu, Oita, Japan, 2004.
- [2] P. Robinson, J. C. Wolf, J. A. Law, J. D. Oliver, K. W. Young, and D. A. Harewood-Gill, "Recent MiroSot Developments in the UK," presented at FIRA Robot World Congress, Busan, Korea, 2004.
- [3] K.-H. Han, K.-H. Lee, C.-K. Moon, H.-B. Lee, and J.-H. Kim, "Robot Soccer System of SOTY 5 for Middle League MiroSot," presented at FIRA Robot World Congress, Seoul, Korea, 2002.
- [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] 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.
- [6] R. Emery, K. Sikorski, and T. Balch, "Protocols for Collaboration, Coordination and Dynamic Role Assignment in a Robot Team," presented at IEEE International Conference on Robotics and Automation, Washington DC, 2002.
- [7] 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.
- [8] B. P. Gerkey and M. J. Mataric, "Sold!: auction methods for multirobot coordination," *Robotics and Automation, IEEE Transactions on*, vol. 18, pp. 758-768, 2002.
- [9] A. Stentz and M. B. Dias, "A Free Market Architecture for Coordinating Multiple Robots," Robotics Institute, Carnegie Mellon University CMU-RI-TR-99-42, December 1999.
- [10] T. Lemaire, R. Alami, and S. Lacroix, "A distributed tasks allocation scheme in multi-UAV context," presented at IEEE International Conference on Robotics and Automation ICRA '04, 2004.
- [11] J. Hoff and G. Bekey, "An architecture for behaviour coordination learning," presented at IEEE International Conference on Neural Networks, 1995.
- [12] R. Sorbello, A. Chella, and R. C. Arkin, "Metaphor of Politics: A Mechanism of Coalition Formation," presented at Forming and Maintaining Coalitions and Teams in Adaptive Multiagent Systems, AAAI 2004 Workshop, San Jose, California, 2004.
- [13] L. E. Parker, "ALLIANCE: an architecture for fault tolerant multirobot cooperation," *Robotics and Automation, IEEE Transactions on*, vol. 14, pp. 220-240, 1998.
- [14] T. Huntsberger, P. Pirjanian, A. Trebi-Ollennu, H. D. Nayar, H. Aghazarian, A. J. Ganino, M. Garrett, S. S. Joshi, and P. S. Schenker, "CAMPOUT: A Control Architecture for Tightly Coupled Coordination of Multirobot Systems for Planetary Surface Exploration.," in *IEEE Transactions on Systems, Man & Cybernetics: Part A*, vol. 33, 2003, pp. 550-559.
- [15] L. Chaimowicz, T. Sugar, V. Kumar, and M. F. M. Campos, "An architecture for tightly coupled multi-robot cooperation," presented at IEEE International Conference on Robotics and Automation ICRA '01, 2001.
- [16] A. Stentz, M. B. Dias, R. M. Zlot, and N. Kalra, "Market-Based Approaches for Coordination of Multi-Robot Teams at Different Granularities of Interaction," presented at ANS 10th International Conference on Robotics and Remote Systems for Hazardous Environments, 2004.
- [17] J. Law and J. Johnson, "Discrete Dynamics and the Space Time Possession Game," presented at FIRA Robot World Congress, Busan, Korea, 2004.
- [18] S. Kim, "Voronoi Analysis of a Soccer Game," *Nonlinear Analysis: Modelling and Control*, vol. 9, pp. 233-240, 2004.