Multi-level Network Analysis of Multi-agent Systems

Pejman Iravani

Department of Mechanical Engineering
University of Bath
Bath, BA2 7AY
p.iravani@bath.ac.uk

Abstract. This paper presents a multi-level network-based approach to study complex systems formed by multiple autonomous agents. The fundamental idea behind this approach is that elements of a system (represented by network vertices) and their interactions (represented by edges) can be assembled to form structures. Structures are considered to be at one hierarchical level above the elements and interactions that form them, leading to a multi-level organisation.

Analysing complex systems represented by multi-level networks make possible the study of the relationships between network topology and dynamics to the system's global outcome. The framework proposed in this paper is exemplified using data from the RoboCup Football Simulation League.

1 Introduction

Increasingly, engineers are faced with having to understand and design complex systems. For example, systems such as traffic in order to optimise time and reduce CO_2 emissions, urban transport systems to improve their efficiency, holistic manufacturing systems, health systems, electric power grid, large companies, etc. Traditional systems engineering approaches can not deal with some of the characteristics of complex systems such as emergent behaviour, evolution, co-evolution and adaptation [1]. In order to start understanding complex systems, it is of relevance to develop frameworks which can be used to analyse this phenomena in a systematic manner.

There is no real consensus about what are the essential characteristics that make systems complex. One of the most important characteristics and the focus of this paper is that the behaviour of complex systems is the result of the interactions of its parts. For example, traffic is the result of the interactions among the various vehicles that share the road network, traffic signals, roads, etc. The behaviour of the electric power grid depends on the number of generators connected, the loads, the transmission network, the protection elements, the weather conditions, etc. This is known as *upward causality*, from the interactions of the parts moving upwards to the system's behaviour as a whole.

L. Iocchi et al. (Eds.): RoboCup 2008, LNAI 5399, pp. 495–506, 2009.

Similarly, the behaviour of the whole will have effects on the behaviour of the parts, for example, traffic jams will force drivers to take alternative routes. This is known as *downward causality*. Thus, in complex systems, there is a constant upward-downward causality that drives the dynamics of the system.

Generally the engineer is interested in the system's *macro-dynamics* (dynamics of the system as a whole) rather than on the *micro-dynamics* (dynamics of the parts). For example, engineers will be interested in avoiding traffic jams rather than controlling the speeds and distances among individual vehicles. Unfortunately, the engineer can not directly influence the macro-dynamics of the system, and to make any changes, actuation must be at the micro-level. Therefore, a main research challenge is to understand the interplay between upward and downward causality.

Multi-robot football is an excellent test-bed to study this phenomena as the local interactions of players produce the global outcomes of the game (e.g. attacking, defending), and these global outcomes drive the players' behaviours (e.g. covering, breaking-away from a cover). This paper concentrates on the study of the *multi-level dynamics* of complex systems, and proposes to use a framework based on a multi-level networks to analyse it. The paper is organised in two large sections, in the first half it discusses how complex network can be used to represent and study multi-agent systems, in the second part it sketches the multi-level approach.

2 Network Representation of Robot Football Interactions

Networks can be used to represent a wide variety of systems in which vertices represent elements and edges represent their interactions. For example a network could use vertices to represent web-sites and edges to represent the existence of hyper-links from one to another. They could also be used to represent people as vertices and edges as their friendship relation. In this paper, vertices are used to represent players and edges interactions.

At all times during the simulated football game there are various types interactions occurring. Some interactions involve only players others also involve the environment (positions in the field, the ball and the goals), for example:

- *ClosestTeammate*: two players of the same team are closest to each other.
- *Supporting*: two players of the same team share controlled areas.
- *Cluster*: players from either team are at close distance.

These type of interactions can be defined using Boolean rules and they are considered to exist if the rule evaluates to *true*. For example, the *ClosestTeammate* rule will be true if two players belong to the same team and if the distance between them is the minimum of all the pairwise distances to the other teammate players.

Figure 1 illustrates some of the interactions occurring during the 208^{th} time-step of the 2007 RoboCup final game. It also illustrates their network representation. This paper focuses only on Boolean (existing or non-existing) interactions, thus two vertices will be connected through and edge if the defining rule is *true*. Interaction rules could be extended to have continuous values or directed links, but this is out of the scope of this paper.

Fig.1(a) displays the *ClosestTeammate* interaction as it happens during the game while Fig.1(b) illustrates its representation as a network. Fig.1(c) illustrates the voronoi diagram formed by the players' positions and the limits of the field. This diagram illustrates the areas that are closest to each player, in some sense these can be seen as the areas controlled by the players. Fig.1(d) illustrates the network of *Supporting* interactions which have been defined over the voronoi diagram as follows. If two players of the same team share two or more vertices of the voronoi diagram (their controlled areas are in contact) then the interaction is considered active. For example, the player w_{11} shares two edges of its voronoi area with w_1, w_2, w_3, w_4 and thus is defined to be supporting them (we will use b_i and w_i to refer to black and white player i respectively).

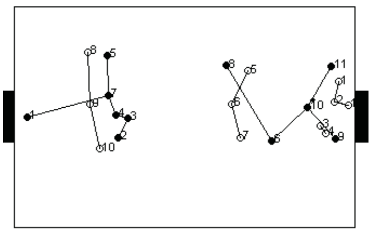

(a) *ClosestTeammate* interaction

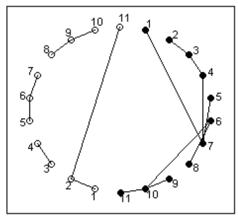

(b) *ClosestTeammate* network

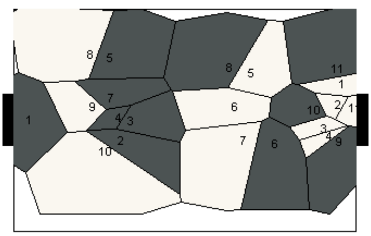

(c) Voronoi areas

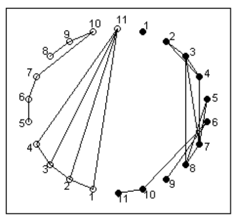

(d) *Supporting* network

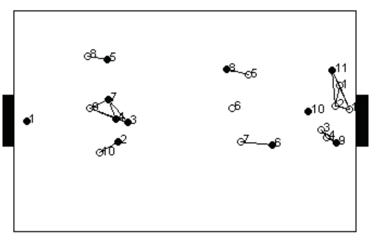

(e) Distance-based clusters

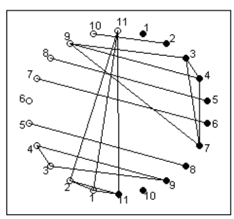

(f) *Cluster* network

Fig. 1. Some interactions in robot football

Figure 1(e) illustrates the *Cluster* interaction which is based on clustering players according to their distances, Fig.1(f) illustrates the network representation. In this particular network, edges between players from different teams can be seen as covering interactions, while edges within the same teams can be seen as support ones. These are just some of the interactions that take place during a football game and can be represented by networks.

The first questions that spring into mind after viewing multi-robot interactions represented by networks are the following:

– Do these networks have any structural patterns or are they random?
– If networks present structural patterns, which ones are beneficial for the team?

In order to address these questions, the next section reviews some of the theory of complex networks specially focusing on how topology affects network functionality.

3 Complex Networks

Mathematically a network can be represented as a graph that consists of two sets, a set of *vertices* $V = \{v_1, v_2, ..., v_n\}$ and a set of *edges* $E = \{e_1, e_2, ..., e_m\}$, such that $V \neq 0$ and E is a set of unordered pairs of elements of V. In other words, vertices are points and edges are links between pairs of points.

Since its inception in 1736, graph theory has been studying the properties of graphs and applied them to solve problems such as calculating flow in pipe networks, distinct colouring of neighbouring areas using minimum number of colours, etc. Recently, the study of networks has been extended to the analysis of complex networks, *i.e.* networks with complex topology which evolve over time [2,5]. Complex topology refers to characteristics such as high clustering coefficient, hierarchical structure and heavy-tail degree distribution.

Network topology refers to the way in which the vertices are connected to each other. There are several values that indicate differences in topology. For example, vertices have different *connectivity degree* that is, they have different number of incident edges, some will be lowly-connected by only a few edges while others may be highly-connected by a large number of edges. The *degree distribution* $P(k)$ of a network is defined as the probability that a randomly selected vertex has connectivity degree k, and gives an indication of how the connectivity is distributed among the vertices of a network. For example, Figure 2 illustrates four different network topologies and their degree distributions. Fig.2(a) illustrates a regular network which has $P(k_i) = 1$ meaning that all vertices have exactly k_i edges. Randomly connected networks Fig.2(c) have a gaussian degree distribution with the peak being the mean value of k. Scale free networks [3] Fig.2(d) have power law distributions meaning that the majority of vertices have few edges and only a few of them have large connectivity degree, this means that some vertices act as *hubs* connecting to many other vertices.

Another important measure that relates to topology is the *average shortest path length*, L, that is the average of the minimum number of edges that separate

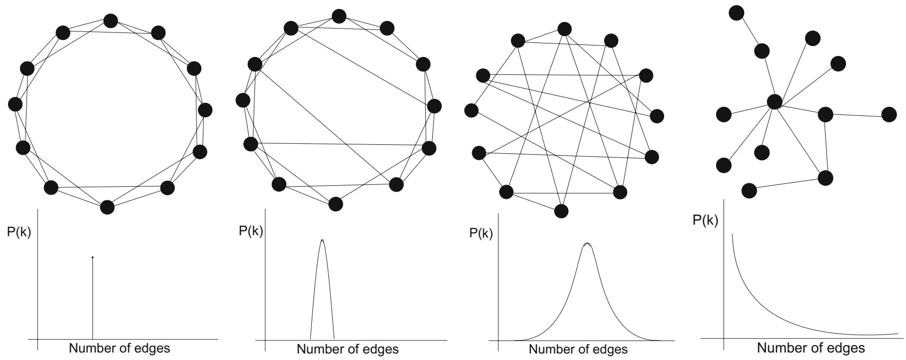

(a) Regular (b) Small-world (c) Random (d) Scale-free

Fig. 2. Different network topologies and their degree distribution

any two vertices, in other words, the average of the shortest paths connecting any two vertices. Equation 1 defines L, where n is the number of vertices in V and d_{ij} is the Euclidean distance between vertices i and j. Networks with small path lengths indicate that is some-how 'faster' to go from one vertex to any other vertex when compared to networks with larger path lengths. This is an important characteristic in networks as it relates to the efficiency of information transmission. A shortcoming with this equation is that L diverges with disconnected vertices, as their distance is infinite. There are two ways to solve this, one is to measure L of the largest connected group, and the other is to use the *efficiency* E measure as defined in Equation 1 where disconnected vertices have a zero effect over the E value.

$$L = \frac{1}{n(n-1)} \sum_{i,j \in V,\ i \neq j} d_{ij} \qquad E = \frac{1}{n(n-1)} \sum_{i,j \in V,\ i \neq j} \frac{1}{d_{ij}} \tag{1}$$

Watts and Strogatz [4] studied the effects of taking a regular network (Fig.2(a)) and randomly eliminating edges and re-wiring them into random vertices. This led into what is known as the *small-world* topology (Fig.2(b)) which is characterised by having a short L when compared with the same type of network (same number of vertices and edges) with a regular topology. This decrease in L is attributed to some edges that act as 'short-cuts' to otherwise distant parts of the network, allowing to travel long distances within the network in only a few steps. The authors presented various different networks with the small-world topology, including the US power grid and the neural network of the *Caenorhabditis* worm. Some of the implications of the small-world topology over the network's function is that there is an enhanced signal-propagation speed and computational power. From the dynamics viewpoint, small-world networks have enhanced synchronisation capabilities [4].

Scale-free networks [3] (Fig.2(d)) have even shorter average path lengths than small-world networks, this is due to the existence of hubs which connect many vertices at once. The implication of the scale-free topologies is that they are robust to random vertex failure but sensitive to targeted attack on hub vertices. In other words, losing random vertices has small effects on the network, but losing hub vertices changes topology and therefore function considerably.

4 Complex Networks of Robot Football Interactions

The previous section illustrated how the topology of networks can be characterised using some statistical measurements and how this affects their functions, such as information transmission, synchronisation and computational power. This section analyses the topology and dynamics of robot football interaction networks and discusses their possible effect over the teams' performance.

4.1 Statistical Analysis of Network Topology

These experiments are conduced over the 2001 to 2007 RoboCup log-files Simulation League finals. Networks are analysed for each time-step for the 7 finals, as each game contains approximately 6000 time-steps, this study is comparing about 42000 networks. For each final, and in order to study the implications of network topology in team performance, the teams have been divided into *winners* (champions of the simulation league) and *losers* (second best team). It should be noted that both finalist teams are high performing, thus only small network differences are expected.

This section studies the *ClosestTeammate*, *Supporting* and *Cluster* interactions (explained in Section 2) by calculating the previous topological measurements.

Figure 3(a) illustrates the degree distribution $P(k)$ of the *ClosestTeammate* network. This shows that although the majority of vertices (over 60%) have only $k = 1$ there are few cases in which players have up to $k = 4$ closest neighbours. This means that during the game there are players that act as hubs connecting to many others. This measurement does not show any significant differences between winner and loser teams, and may only be caused by the spatial distribution of players in the field.

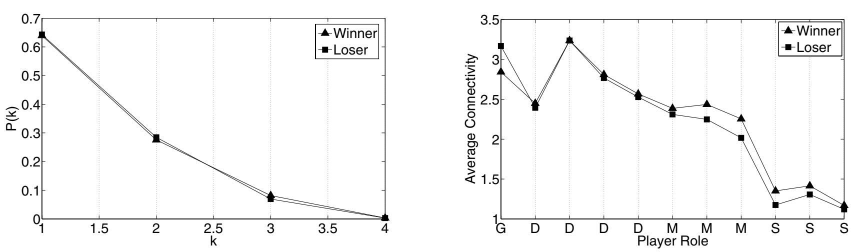

(a) $P(k)$ in *ClosestTeammate* (b) Average connectivity in *Supporting*

Fig. 3. Degree distributions of two different interaction networks

Table 1. Winner vs Loser statistics

Year	2001		2002		2003		2004		2005		2006		2007	
Result	W	L	W	L	W	L	W	L	W	L	W	L	W	L
E Supporting	0.21	0.19	0.19	0.22	0.19	0.18	0.23	0.20	0.23	0.18	0.19	0.21	0.22	0.19
%FC *Supporting*	36.2	21.4	18.3	35.9	23.1	21.2	40.9	24.6	45.9	14.6	22.7	39.7	41.3	17.6
#Clusters	12.49		11.92		11.82		10.88		9.66		9.85		9.72	

Figure 3(b) illustrates the average connectivity of the *Supporting* network (Fig.1(d)) with respect to each of the player's role, one goalie (G), four defenders (D), four mid-fielders (M) and three strikers (S). In other words, the graph illustrates the average number of supporting team mates in relation to the player's role. It shows that defenders are the group with highest connectivity while strikers are the ones with lowest, thus, teams tend to play with cohesive defenders and distributed strikers. The graph also shows connectivity differences between winner and loser teams, specially for their mid-fielders and strikers. The average of differences are 7.07% and 8.48% for mid-fielders and strikers respectively. A significance test (ttest) on the data showed over a 99% confidence on the difference of the mean values. This differences indicate that mid-field and striker players in winner teams are better supporting each other, in other words, they better position themselves to make and receive passes (in previous analysis it was shown that the *Supporting* interaction allows for higher probability of successful passes [8]).

Table 1 illustrates some more measurements related to network topology for winner (W) and loser (L) teams. The *E Supporting* row is the average *Efficiency* (Eq.1) of the *Supporting* network. Larger values indicate that the network has improved transmission capability, in the football scenario this implies teams which could pass the ball faster from one part of the field to another. It seems that winning teams have higher E value (5 out of 7 finals). This is not the case for the 2002 and 2006 finals, which indicates that having a good team-connectivity or high E is not sufficient to win or lose a game.

The %FC *Supporting* row indicates the percentage of time that the *Supporting* network is fully connected, or in other words, the times that there is a path in the network that connects all vertices. During the game this shows a path for fast passing the ball from defenders to strikers. The pattern for this measurement is the same as the discussed above and has the same explanation.

The #Clusters row indicates the average number of distance-based clusters that appear during each final, it is interesting to observe that this number has been decreasing during the evolution of the League from 2001 to 2007 indicating that teams are playing closer to each other, or in other words, they are involving more players into the active areas of the game.

4.2 Dynamics of Complex Networks

As games progress the interactions among players change and thus do the networks they define. Figure 4 illustrates a sequence from the 2007 final in which

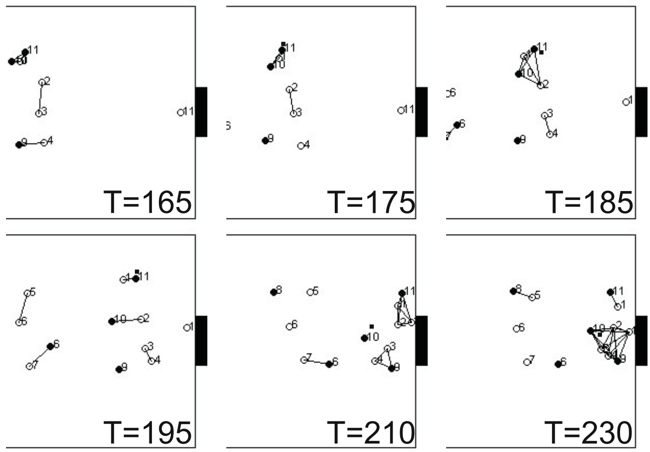

Fig. 4. Goal sequence in 2007 log-file

the winning team scores a goal. In summary, what occurs is that b_{11} dribbles from the mid-field up to the opponent's left-corner (from T=165 to T=195), at this point b_{10} has moved into a good position (being uncovered and in front of the goal). At T=210 player b_{11} passes to b_{10} which then dribbles and scores.

The edges between players in Figure 4 represent the *Cluster* interaction, an edge exists if players are close to each other (distance $<$ 10m). Between players of different teams, this interaction resembles a covering relation, whereas among players of the same team it resembles a support relation. It is important to observe what happens to this interaction between T=185 and T=195, in this period the b_{10} has moved away from its team mate b_{11} but also from the defender w_1 (the interaction with these players at T=195 is inactive). By T=210 b_{10} has broken the interaction with its cover w_2 and its now well positioned to receive a pass from b_{11} and score; this is exactly what happens.

Is it possible to abstract this type of information by studying how networks evolve in time? Figure 5 seems to illustrate some patterns in the dynamics of networks and the game's outcome. It displays the evolution of the *Cluster* interaction during a fraction of the game, between the kick-off (T=0) to the T=500 time-step. The continuous line indicates the number of clusters, the horizontal line indicates the mid-field position and the discontinuous line indicates the ball position. High number of clusters means that players are far from each other in the field, low number of clusters means that players are in close proximity. In this period of time, the winner team scores two goals, at T=247 and T=492.

The figure shows that when the winner team attacks (climbing ball position) the total number of clusters increases. Specially on the build-up towards the goal from T=120 to T=220 and T=330 to T=460. The opposite happens when the loser team is attacking (descending ball position), and the total number of clusters decreases. In other words, when the winner team attacks its players move in such a way that they separate from their covers (as seen in the sequence in Fig.4) increasing the probability of receiving a pass and scoring a goal. This

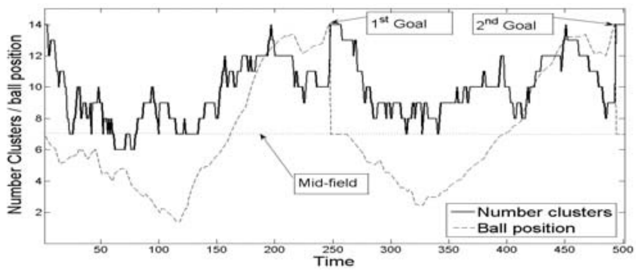

Fig. 5. Evolution of the *Cluster* network

exemplifies that different network dynamics may have different outcomes in the game, obviously one example is not enough to generalise to all games, but it is a good starting point to show the influence of network dynamics.

5 Multi-level Network Analysis Framework

A fundamental idea behind our work is that elements of sets can be assembled under network relations to form structures. Structures exist at a higher level than its elements in a multi-level hierarchy. The framework used to represent and study multi-level structures is defined below, these resembles the work presented in [6] but has some clear differences in the way structures are assembled. In their work, a special process is defined to assemble structures. In our work, structures are defined by the self-organising networks of interactions and thus only interaction rules need to be defined.

Let us assume that certain structures exist at the base level of the hierarchy (Level N), in our example the structures at this level correspond to players and they can be represented in the following manner:

$$\sigma_{i,t}^{l} = \langle s, r, p \rangle$$

where i is structure number, t is the time at which the structure is observed, l is the level, s is the state of the structure *e.g.* position, orientation, ball-possession; r is a set of interaction rules *e.g.* the *Cluster* interaction rule; and p is a set of properties, *e.g.* the probability of scoring a goal.

At the next level (Level N+1) structures $\sigma_{i,t}^{l+1}$ are formed by the assembly of Level N structures defined by their interaction network. Thus, if a number of Level N structures are interacting, then they assemble into a Level N+1 structure. Level N+1 structures have state, interaction rules and properties but these may be different from the ones defined at the lower level. Properties at Level N+1 can be considered to be *emergent* if they do not exist at Level N, *e.g.* if single players are structures at Level N they can not have the property of passing the ball as this property requires more than one player. The next sections present some of the multi-level analysis conducted over the 2007 RoboCup final game which was won by the black team.

5.1 Multi-level Properties

Figure 6 illustrates the previous notation over a sequence of the 2007 game.

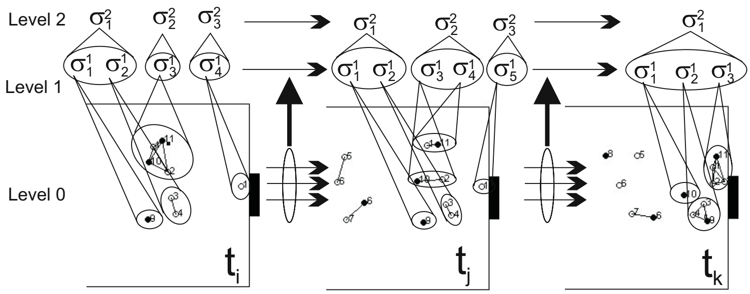

Fig. 6. Multi-level assembly and its dynamics

Let us refer to the base level as Level 0 (L0). Base level structures are defined as: i is the player number; t is the simulator time-step; s is the player's position in the field; r is the *Cluster* interaction rule (active when in close proximity of other players) and p is a property that indicates to what team the structure belongs to. For example, at time t_i player b_{10} generates the following structure:

$$\sigma^0_{b_{10},t_i} = \langle s = position(b_{10}), r = Cluster, p = black\rangle$$

The interaction of L0 structures results in Level 1 (L1) structures, where i is a number from $1, 2, ...m$ with m being the total number of networks[1] defined by the structures at L0. The state s is defined as the centroid of the positions of the L0 structures; r continues being a *Cluster* interaction, this time with threshold (distance < 20m). The property p indicates which team has more players in the structure, $p = W$ more white players, $p = B$ more black players, and $p = X$ same number of players. A $p = W$ structure is beneficial for white and a $p = B$ for black. For example, at time t_i the following is a L1 structure with $p = X$ as it contains as many white as black players:

$$\sigma^1_{3,t_i} = \langle s = centroid(\sigma^0_{b_{10},t_i}, \sigma^0_{b_{11},t_i}, \sigma^0_{w_1,t_i}, \sigma^0_{w_2,t_i}), r = Cluster, p = X\rangle$$

The interaction of L1 structures will generate new structures at Level 2 (L2). L2 structures have s defined as the centroid of L1 structures and p as before. This time, no further interaction rules are defined ($r = \phi$), thus L2 will be the highest on the hierarchy as seen in Figure 6. The following is a L2 structure:

$$\sigma^2_{1,t_i} = \langle s = centroid(\sigma^1_{1,t_i}, \sigma^1_{2,t_i}), r = \phi, p = X\rangle$$

[1] We also consider single vertices as networks.

Table 2. Multi-level p

	W	B	X
Level 0	50%	50%	0%
Level 1	27.7%	28.8%	43.4%
Level 2	34.7%	38.4%	26.7%

Table 2 shows the structures' properties at different levels in the hierarchy for the 2007 final. At L0 there is the same proportion of structures with the white (W) and black (B) properties, and none with the X. This is because both teams have the same number of players. At L1, structures start to show more interesting properties, with X being the largest followed by B and W. This indicates that most structures are equilibrated (43.4%) but with a slight advantage for the black team. L2 reveals a stronger presence of black structures over white ones (38.4% vs 34.7%).

This analysis shows how properties have different meanings at different levels, at L0 the property p indicates that both teams have the same number of players. At L1 most structures are in equilibrium ($p = X$) this is because at this level structures mainly represent local interactions between close-by players, usually by 1-vs-1 interactions. At L2 the proportion of structures with team advantage ($p = W$ or $p = B$) is larger than the equilibrated ones. This indicates that L2 is more strategic, and in this case, favourable to the black team.

6 Conclusion and Further Work

This paper has presented a framework to study complex multi-agent systems based on representing the system using multi-level networks. The main idea is that agents can be represented as nodes and interaction as edges of a network. Different interactions over the same agents define different networks, for example this paper showed how three different interactions (*ClosestTeammate*, *Supporting* and *Cluster*) result on different networks over the same set of agents.

The first part of the paper reviewed some of the complex network literature, the main purpose was to show that network topology affects network functionality. Some evidence of this was then provided in relation to the RoboCup Simulation League in which slightly better teams display different network topology during the game in comparison with the slightly worse teams. The main evidence was that better teams have slightly higher connectivity and efficiency in the *Supporting* network, specially in the mid-field and attacking positions.

The second part of the paper introduced the multi-level analysis framework. One of the main ideas in the paper was that (sub)networks formed by players and their interactions can be considered to be a structure that exists at a higher hierarchical level. It was argued that structures have different properties at different levels. For example the majority of Level 1 structures have the same number of players from both teams. Level 2 structures start to be more strategic and show how some teams position better their players to gain structural advantage by having one or more players free of cover. This analysis is interesting as many properties of complex system are not visible at low-levels but are the result of the interaction of their parts.

One of the main aims of the RoboCup Coach competition was to have an agent that observes the opposition team and is capable of modelling their behaviour. The basic principle exploited here was to define behaviours as functions know *a-priory* and using statistics to calculate the parameters of the function. For example, the positioning behaviour could be a function of the home-position and ball-position [9], therefore this type of pattern recognition was based on statistical function fitting. We believe that approach presented in this paper could be used to a similar extent in this league with the difference that patterns would not be defined as parametric functions but as relational functions and that multi-level information could be used. The results presented in this paper are based on past final games in which the teams have similar performance (being the best two of that year), the following experiments will be conducted on controlled teams with different capabilities.

References

1. Ottino, J.M.: Engineering Complex Systems. Nature 427, 398–399 (2004)
2. Boccaletti, S., et al.: Complex Networks: Structure and dynamics. Physics Reports 424, 175–308 (2006)
3. Barabasi, A., Bonabeau, E.: Scale-Free Networks. Scientific American, 50–59 (2003)
4. Watts, D.J., Strogatz, S.H.: Collective dyanamics of 'small-world' networks. Nature 393, 440–442 (1998)
5. Albert, R., Barabasi, A.L.: Statistical mechanics of complex networks. Reviews of Modern Physics 74, 47–97 (2002)
6. Rasmussen, S., et al.: Ansatz for Dynamical Hierarchies. Artificial Life 7, 329–353 (2001)
7. Johnson, J.H., Iravani, P.: The Multilevel Hypernetwork Dynamics of Complex Systems of Robot Soccer Agents. ACM Transactions on Autonomous Systems 2, 1–23 (2007)
8. Iravani, P.: Discovering relevant sensor data using Q-analysis. In: Bredenfeld, A., Jacoff, A., Noda, I., Takahashi, Y. (eds.) RoboCup 2005. LNCS, vol. 4020, pp. 81–92. Springer, Heidelberg (2006)
9. Kuhlmann, G., Knox, W.B., Stone, P.: Know Thine Enemy: A Champion RoboCup Coach Agent. In: Proceedings of the Twenty-First National Conference on Artificial Intelligence, pp. 1463–1468 (2006)