Towards Quantifying Interaction Networks in a Football Match

Oliver M. Cliff, Joseph T. Lizier, Rosalind X. Wang, Peter Wang, Oliver Obst, and Mikhail Prokopenko

CSIRO Information and Communication Technologies Centre, Adaptive Systems, P.O. Box 76, Epping, NSW 1710, Australia

Abstract. We present several novel methods quantifying dynamic interactions in simulated football games. These interactions are captured in directed networks that represent significant coupled dynamics, detected information-theoretically. The model-free approach measures information dynamics of both pair-wise players' interactions as well as local tactical contests produced during RoboCup 2D Simulation League games. This analysis involves computation of information transfer and storage, relating the information transfer to responsiveness of the players and the team, and the information storage within the team to the team's rigidity and lack of tactical flexibility. The resultant directed networks (interaction diagrams) and the measures of responsiveness and rigidity reveal implicit interactions, across teams, that may be delayed and/or long-ranged. The analysis was verified with a number of experiments, identifying the zones of the most intense competition and the extent of interactions.

1 Introduction

Many team games, real and virtual, are characterised by rich interactions occurring dynamically and shaping the course of the contest both locally and globally. The interactions across the teams are created by opposing objectives of competing players and tactical schemes. The interactions within a team are usually constrained by cooperation and shared plans. Generally, the interactions are directed (e.g., a defender is marking an opponent's forward), varying in strength over time and/or space, and typically do not result from direct messaging or communications — rather they manifest some tacit correlations that often are delayed in time and/or are long-ranged over the play-field.

While a significant number of patterns emerging during a game may be evident even without an in-depth analysis, most of the interactions may appear intractable to an external observer who does not have an access to the logic and neural processing of the players. One then may formulate a general problem: how can an external observer identify most generic interaction networks that link together autonomous players, without re-constructing the players' behaviour and using only the positional data, such as planar coordinates and their changes? The problem is difficult as some of the dependencies between players are not discernible simply by correlating their dynamic locations over time — one needs to take into the account a possibly directed nature of such correlations, where dynamics of one of the players affects the positioning of another.

In general, as mentioned by Vilar et al. [1], "quantitative analysis is increasingly being used in team sports to better understand performance in these stylized, delineated,

S. Behnke et al. (Eds.): RoboCup 2013, LNAI 8371, pp. 1–12, 2014.
© Springer-Verlag Berlin Heidelberg 2014

complex social systems". One of the older examples is "sabermetrics" — the specialised analysis of baseball through objective evidence, e.g. baseball statistics measuring in-game activity [2]. Another recent example is described by Fewell et al. [3] who analysed basketball games as networks, where players are represented as nodes and passes as edges: the resulting network captures ball movement, at different stages of the game. Their work studies network properties (degree centrality, clustering, entropy and flow centrality) across teams and positions, and attempts to determine whether differences in team offensive strategy can be assessed by their network properties. Strategic networks analysed by Fewell et al. consider only explicit interactions (such as passes) within a team, and not implicit (delayed and/or long-ranged) interactions, across teams.

Another very recent investigation by Vilar et al. [1] proposed a novel method of analysis that captures how teams occupy sub-areas of the field as the ball changes location. This study was important in focussing on the local dynamics of team collective behavior rather than individual player capabilities: when applied to football (soccer) matches, the method suggested that players' numerical dominance in some local sub-areas is a key to "defensive stability" and "offensive opportunity". While the method rigorously used an information-theoretic approach (e.g. the uncertainty of the team numerical advantage across sub-areas was determined using Shannon's entropy), it was not aimed at and did not produce interaction networks, either explicit or implicit.

Construction of interaction networks for (possibly competing) teams is not unique to sport, but arguably its utility can be leveraged quite strongly in team games, such as football, basketball and so on. RoboCup 2D Soccer Simulation League is a well-known benchmark domain for Artificial Intelligence that specifically targets soccer with its realistic and challenging multi-agent dynamics, characterised by autonomous decision-making under constraints, set by tactical plans and teamwork (collaboration) as well as opponent (competition) [4,5,6,7,8,9,10,11], and so we use this domain in our study.

Information dynamics is a recent methodology for analysis of complex systems in general and swarm behavior in particular. In this paper we describe a novel application of information dynamics to the RoboCup 2D Simulation. In particular, we develop an approach to build several interaction diagrams, given data from a number of games, followed by a tactical analysis. The interaction diagrams reveal a few interesting dependencies between pairs of players that are useful for game analysis, while the tactical analysis extends these findings to formation-level interactions (e.g., between defensive line-up of team Y with the attacking line-up of team X, etc.).

2 Motivation and Approach

2.1 Information Dynamics

A recently developed framework of *information dynamics* studies the phenomenon of computation in a systematic way: it uncovers and analyses information-theoretic roots of the most basic computational primitives: *storage*, *transmission*, and *modification of information* [12,13,14,15].

The *active information storage* quantifies the information storage component that is directly in use in the computation of the next state of a process [15]. More precisely, it is the average mutual information between the semi-infinite past of the process

$x_n^{(k)} = \{x_{n-k+1}, \ldots, x_{n-1}, x_n\}$ (as $k \to \infty$) and its next state. The *local information storage* (or pointwise mutual information) is then a measure of the amount of information storage in use by the process at a particular time-step $n+1$:

$$a_X(n+1) = \lim_{k\to\infty} \log_2 \frac{p(x_n^{(k)}, x_{n+1})}{p(x_n^{(k)})p(x_{n+1})}. \tag{1}$$

In practice, one deals with finite-k estimates $a_X(n+1,k)$, as well as the finite-k estimates $A_X(k)$ of the average active information storage $A_X = \langle a_X(n+1)\rangle_n$.

Transfer entropy [16] is designed to detect asymmetry in the interaction of subsystems by distinguishing between "driving" and "responding" elements. The *local information transfer*, based on *transfer entropy*, captures information transmission [12] from source Y to destination X, at a particular time-step $n+1$. It is defined as the information provided by the source y_n about the destination's next state x_{n+1} that was not contained in the past of the destination $x_n^{(k)}$:

$$t_{Y\to X}(n+1) = \lim_{k\to\infty} \log_2 \frac{p(x_{n+1} \mid x_n^{(k)}, y_n)}{p(x_{n+1} \mid x_n^{(k)})}. \tag{2}$$

It is important to realise that information transfer between two variables does not require an explicit communication channel, it rather indicates a high degree of directional synchrony or nonlinear correlation between the source and the destination. It characterises a degree of *predictive* information transfer, i.e., "if the state of the source is known, how much does that help to predict the state of the destination?" [12].

Sometimes it is useful to condition the local information transfer on another contributing process W, considering the *local conditional transfer entropy* [13]:

$$t_{Y\to X|W}(n+1) = \lim_{k\to\infty} \log_2 \frac{p(x_{n+1} \mid x_n^{(k)}, y_n, w_n)}{p(x_{n+1} \mid x_n^{(k)}, w_n)}. \tag{3}$$

In this study we used the average information transfer $t_{Y\to X|W} = \langle t_{Y\to X|W}(n+1)\rangle_n$. One may, however, utilise local values as well in order to trace the information dynamics over time, e.g. identifying its peaks during specific moments.

2.2 Pair-Wise Information Dynamics and Interaction Diagrams

In order to estimate strength of directed coupling between two agents we compute the average transfer entropy between them during any given game. For a game g with N time steps, between two teams $\mathbb{X}$ and $\mathbb{Y}$, the local transfer entropy at each time step $n \leq N$ is calculated between each source variable Y_i (a change in the 2D positional vector of agent i from team $\mathbb{Y}$) and destination variable X_j (a change in the 2D positional vector of agent j from team $\mathbb{X}$), given the change in current 2D ball position b:

$$t^g_{Y_i\to X_j|b}(n) \ .$$

Dynamics of the ball is conditioned upon in order to compute the transfer entropy in context of the game, which is greatly affected by the ball trajectories. Then, the average transfer entropy for each source-destination pair over the entire match is calculated as

$$T^g_{Y_i \to X_j|b} = \frac{1}{N} \sum_{n=0}^{N-1} t^g_{Y_i \to X_j|b}(n) \; . \tag{4}$$

Information-Sink Diagrams. Once the game's average transfer entropy, $T^g_{Y_i \to X_j|b}$, is determined for each pair Y_i, X_j, we identify the *source* agent $\hat{Y}_i(X_j, g)$ from the opposing team that transfers maximal information to a given agent X_j:

$$\hat{Y}_i(X_j, g) = \arg\max_{Y_k \in Y} \; T^g_{Y_k \to X_j|b} \; . \tag{5}$$

Over a number of games G, we select the source agent $\hat{Y}_i(X_j)$ that transfers maximal information to X_j most frequently, as the mode of the series $\{\hat{Y}_{i_1}(X_j, 1), \ldots, \hat{Y}_{i_G}(X_j, G)\}$. Then, we consider the average information transfer between these two agents $\hat{Y}_i = \hat{Y}_i(X_j)$ and X_j across all games:

$$T_{\hat{Y}_i \to X_j|b} = \frac{1}{G} \sum_{g=1}^{G} T^g_{\hat{Y}_i \to X_j|b} \; . \tag{6}$$

Intuitively, the movement of the source agent $\hat{Y}_i = \hat{Y}_i(X_j)$ affected the agent X_j more than movement of any other agent in team $\mathbb{Y}$. That is, the agent X_j was responsive most to movement of the source agent $\hat{Y}_i$. Crucially, when we use the notion of *responsiveness* to another (source) agent, we do not load it with such semantics as being dominated by, or driven by that other agent. Higher responsiveness may in fact reflect either useful reaction to the opponent's movements (e.g., good marking of the source), or a helpless behaviour (e.g., constant chase after the source). Vice versa, generating a high responsiveness from another agent may result in either a useful dynamic (e.g., positional or even tactical dominance over the responding agent), or a wasteful motion (e.g., being successfully marked by the responding agent). In short, responsiveness captured in the maximal transfer $T_{\hat{Y}_i \to X_j|b}$ detects a directed coupling from the source agent $\hat{Y}_i$ to the responding agent X_j and should not be interpreted in general as a simple index for comparative performance. It is, however, a useful identifier of the opponents' source player that was affecting a given agent X_j most.

Given a series of games, we identify the pairs "source-responder" by finding the source agent for each of the agents on both teams (always choosing the source among the opponents). The identified pairs can be visualised in an "information-sink" interaction diagram $\hat{D}(\mathbb{Y}, \mathbb{X})$ that depicts a directed graph with 20 nodes representing players (typically excluding goalkeepers), with the edges representing all source-responder pairs, where a single edge is incoming to every agent from the corresponding source.

Figure 1a shows the information-sink interaction diagram $\hat{D}$(Oxsy, Gliders) built for several hundred games between Oxsy and Gliders (cf. Results section).

Information-Base Diagrams. Similarly, having obtained the average transfer entropy during a game, $T^g_{Y_i \to X_j|b}$ for all pairs, we identify the *responder* agent $\check{X}_j(Y_i, g)$ that "received" maximal information from a given agent Y_i. Formally, for any game g:

$$\check{X}_j(Y_i, g) = \arg\max_{X_k \in \mathbb{X}} \; T^g_{Y_i \to X_k|b} \; . \tag{7}$$

Over a number of games G, we select the responder agent $\check{X}_j(Y_i, g)$ to whom maximal information was transferred by Y_i most frequently, as the mode of the series $\{\check{X}_{j_1}(Y_i, 1), \ldots, \check{X}_{j_G}(Y_i, G)\}$. Finally, we consider the average information transfer between these two agents Y_i and $\check{X}_j = \check{X}_j(Y_i, g)$ across all games:

$$T_{Y_i \to \check{X}_j | b} = \frac{1}{G} \sum_{g=1}^{G} T^g_{Y_i \to \check{X}_j | b} \, . \tag{8}$$

The pairs $(Y_i, \check{X}_j)$ identified for each agent treated as a source are combined in an "information-base diagram" $\check{D}(\mathbb{Y}, \mathbb{X})$.

The intuition in this case is the same as in the previous subsection — the difference is that now we identify the highest responder agent, having selected a source. In general, of course, the pair $(\hat{Y}_i, X_j)$ defined for the information-sink diagrams and the pair $(Y_i, \check{X}_j)$ defined for the information-base diagrams may differ. That is, the agent Y_i may be the most informative source $\hat{Y}_i$ for the agent X_j, among all possible sources in Y, but the agent X_j may be not the best responder $\check{X}$ to the agent Y_i among all possible responders in $\mathbb{X}$, and vice versa.

While an information-sink diagram reflects more where the information tends to be transferred to, an information-base diagram tends to depict where the information is transferred from. Neither of the diagrams presents a complete "story", highlighting only a small part of the overall information dynamics. There are more comprehensive diagrams, where the edges would represent in the descending order the highest information transfers for all the pairs, retaining a given number of such links, or keeping the edges for the information amounts above a certain threshold, etc. — in these instances, some agents may have no incoming or outgoing links at all. Nevertheless, we believe that the interaction diagrams presented here are valuable, being particularly simple and easy to interpret. Specifically, for an information-sink diagram every agent has an incoming edge, and for an information-base diagram every agent has an outgoing edge.

Figure 2a shows the information-base interaction diagram $\check{D}$(Oxsy, Gliders) built for several hundred games between Oxsy and Gliders (cf. Results section).

2.3 Tactical Analysis

Building up on the information dynamics measures, it is possible to investigate group behavior in complex systems, such as swarms. For instance, recent studies by Wang et al. [17] quantitatively verified the hypothesis that the collective memory within a swarm can be captured by *active information storage*. Higher values of storage are associated with higher levels of dynamic coordination, while information cascades that correspond to long range communications are captured by *conditional transfer entropy* [12,13]. In other words, information transfer was shown to characterise the communication aspect of collective computation distributed within the swarm.

In applying information dynamics to the RoboCup 2D Simulation League we make the following conjecture:

> *a higher information transfer $t_{Y \to X|W}$ from the source Y (e.g. dynamics of player Y) to the destination X (e.g., dynamics of another player X), in the context of some other dynamics W (e.g., the movement of the ball W), is indicative of a higher responsiveness of the process/player X to the process/player Y.*

That is, the "destination" player Y responds, for example, by repositioning, to the movement of the "source" player Y. This may apply to many situations on the field, for instance, when one team's forwards are trying to better avoid opponent's defenders, we consider the information transfer $t_{\mathbb{Y}_{def} \to \mathbb{X}_{att}}$ from defenders $Y_i \in \mathbb{Y}_{def}$ to forwards $X_j \in \mathbb{X}_{att}$, where the involved probability distributions are obtained for different relative positions on the soccer field. Vice versa, the dynamics of the opponent's defenders, who are trying to better mark our team's forwards, are represented in the information transfer $t_{\mathbb{X}_{att} \to \mathbb{Y}_{def}}$ from forwards $X_j \in \mathbb{X}_{att}$ to defenders $Y_i \in \mathbb{Y}_{def}$. These two examples specifically consider a coupling between the attack line $\mathbb{X}_{att}$ of our team and the defense line $\mathbb{Y}_{def}$ of opponent's team (henceforth we keep denoting opponent's lines (attack, midfield or defense) by $\mathbb{Y}_{line}$ and our team's lines by $\mathbb{X}_{line}$).

We further contrast these two transfers in the coupled lines:

$$\Delta(\mathbb{X}_{att}, \mathbb{Y}_{def}) = t_{\mathbb{Y}_{def} \to \mathbb{X}_{att}} - t_{\mathbb{X}_{att} \to \mathbb{Y}_{def}} \ . \tag{9}$$

When our forwards are more responsive on average to the opponent's defenders than the opponents defenders are to our forwards, $t_{\mathbb{Y}_{def} \to \mathbb{X}_{att}} > t_{\mathbb{X}_{att} \to \mathbb{Y}_{def}}$, and the relative responsiveness $\Delta(\mathbb{X}_{att}, \mathbb{Y}_{def}) > 0$. It is also possible to combine relative responsiveness scores for each of the coupled lines in the overall tactical relative responsiveness (including, for example, relative scores for midfielders $\mathbb{X}_{mid}$ and $\mathbb{Y}_{mid}$):

$$\Delta(\mathbb{X}, \mathbb{Y}) = \Delta(\mathbb{X}_{att}, \mathbb{Y}_{def}) + \Delta(\mathbb{X}_{def}, \mathbb{Y}_{att}) + \Delta(\mathbb{X}_{mid}, \mathbb{Y}_{mid}) \ . \tag{10}$$

Here all the transfers to team $\mathbb{X}$ are added up, and the transfers from team $\mathbb{X}$ are subtracted away. When each of the transfers is conditioned on some other contributor W (e.g., all the dynamics are computed in the context of the ball movement), the overall tactical relative responsiveness $\Delta(\mathbb{X}, \mathbb{Y}|W)$ is also placed in this specific context, W.

In principle, competitive situations result in quite vigorous dynamics within the involved lines and overall formations, and the team that manages to achieve a higher degree of tactical relative responsiveness does often perform better. While this is not a hard rule, we may correlate the scores of relative responsiveness (e.g., line-by-line) with the game scores, and identify the lines which impacted on the games more.

Our tactical analysis also involves computation of the active information storage within the teams. We characterise team's rigidity $A_{\mathbb{X}}$ as the average of information storage values for all players of the team. We also determine the relative rigidity $A(\mathbb{X}, \mathbb{Y}) = A_{\mathbb{X}} - A_{\mathbb{Y}}$ for the teams (or their coupled lines). The hypothesis here is that

> *a higher rigidity $A_{\mathbb{X}}$ within the team is indicative of a higher dependence of players on each other, or a higher redundancy within the team's motion.*

The average information storage, or rigidity $A_{\mathbb{X}}$, is high whenever one can predict the motion of some players based on the movements of their other teammates. In these cases, the players are not as independent of each other as a truly complex or swarm behavior would warrant, making the tactics less versatile. Obviously, this may be counter-productive, since an opponent team can counteract by only partially observing the 'rigid' team's dynamics, and deducing the rest. Consequently, the relative rigidity $A(\mathbb{X}, \mathbb{Y})$ should be anti-correlated with the team $\mathbb{X}$ performance against team $\mathbb{Y}$.

3 Results

To compute the measures described in previous sections, produce interaction diagrams and correlate tactical responsiveness with team performance, we carried out multiple iterative experiments matching Gliders2013 up against some well-known teams, such as Oxsy [18] and Marlik [19]. The correlation scores (Pearson product-moment correlation coefficients) reported below were tested for statistical significance, and corrected for multiple comparisons.

3.1 Interaction Diagrams

Figure 1 presents the information-sink interaction diagram $\hat{D}(\text{Oxsy, Gliders})$ and the information-base interaction diagram $\check{D}(\text{Oxsy, Gliders})$, built over almost 500 hundred games between Oxsy and Gliders. Analogously, Fig. 2 shows the information-sink interaction diagram $\hat{D}(\text{Marlik, Gliders})$ and the information-base interaction diagram $\check{D}(\text{Marlik, Gliders})$, built over nearly 450 hundred games between Marlik and Gliders.

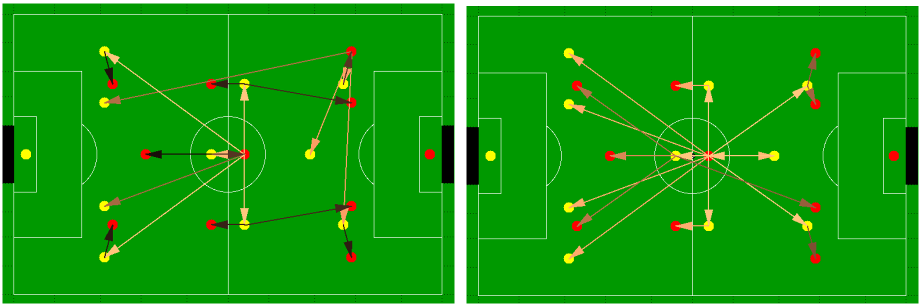

(a) Information-sink diagram for Gliders (left) and Oxsy (right)

(b) Information-sink diagram for Gliders (left) and Marlik (right)

Fig. 1. Interaction-sink diagrams. Arrows represent highest information transfer between players. MATLAB copper colormap is used to indicate the strength of transfer, varying smoothly from black (weakest) to bright copper (strongest). Example interactions: two arrows in the left diagram from Oxsy's central mid-fielder, positioned in the centre circle, to Gliders' left and right defenders indicate that these defenders respond mostly to the central mid-fielder's motion.

Several interesting observations can be made. In general, the diagrams are highly symmetric with respect to left and right wings. The diagrams represent interactions averaged over many games, and so the symmetry demonstrates that the employed methods are robust to noise present in individual games. Also, the information-sink diagrams do differ from information-base diagrams, as expected. We begin a more detailed analysis with the information-sink interaction diagrams 1a and 1b:

- Gliders' defenders mostly respond to opponent's central mid-fielder;
- Gliders' mid-fielders mostly respond to opponent's central mid-fielder;
- Gliders' forwards mostly respond to Oxsy's defenders or Marlik's central mid-fielder;

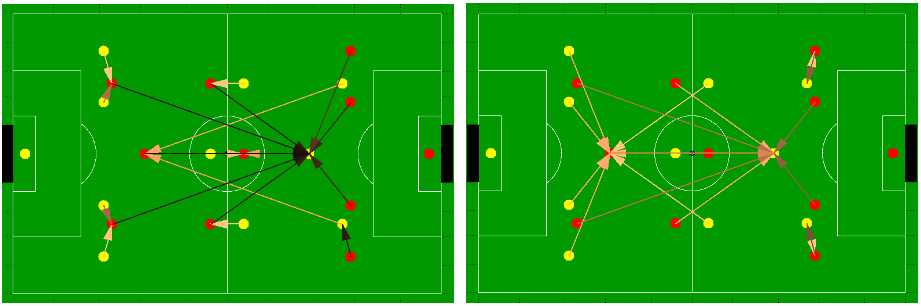

(a) Information-base diagram for Gliders (left) and Oxsy (right)

(b) Information-base diagram for Gliders (left) and Marlik (right)

Fig. 2. Interaction-base diagrams. Arrows represent highest information transfer between players. MATLAB copper colormap is used to indicate the strength of transfer, varying smoothly from black (weakest) to bright copper (strongest). Example interactions: four arrows in the right diagram from Marlik's central mid-fielder, positioned in the centre circle, to all Gliders' defenders indicate that the defenders respond mostly to the central mid-fielder's motion.

- Oxsy's wing forwards mostly respond to Gliders' side defenders, while Oxsy's centre-forward does not mostly respond to Gliders' centre-backs;
- Oxsy's side defenders mostly respond to Gliders' wing forwards, while Oxsy's centre-backs do not mostly respond to Gliders' centre-forward;
- Marlik's forwards mostly respond to Gliders' central mid-fielder;
- Marlik's defenders mostly respond to Gliders' side-wingers.

Now we turn our attention to the information-base interaction diagrams 2a and 2b:

- Gliders' defenders mostly transfer information to Oxsy's wing forwards, but not to their centre-forward;
- practically every Oxsy's player transfers infromation to Gliders' centre-forward;
- Gliders' defenders mostly transfer information to Marlik's centre-forward, but not to their wing-forwards;
- Gliders' centre-forward is transferred information from many Marlik's players, but not from their side defenders;
- Gliders's wing forwards are tightly coupled with Marlik's side defenders.

Even such a brief analysis helps to point out that in the contest with Oxsy, Gliders have a problem with their centre-backs not actively checking the opponent's centre-forward, but a similar problem also exists in Oxsy's own defense. Not surprisingly, most goals are scored in these games through the centre and not via the wing attacks and crosses. In addition, it appears that a lot of Gliders' motion is tuned to opponents' central mid-fielder which highlights a high degree of redundancy that may need to be exploited. In the games against Marlik it is evident that the opponents central mid-fielder plays a key role in both defense and attack, which again presents an opportunity to exploit such an overload. At the same time, it appears that a lot of interactions occur on the flanks of Marlik's defense (defenders mark forwards who try to evade), while Marlik's wing forwards are not marked by Gliders's side defenders.

3.2 Tactical Analysis

In this subsection, we correlate scores of relative responsiveness (either line-by-line or overall), as well as rigidity, with the game scores, and identify the lines which impacted on the games more. That is, we compute a correlation coefficient between a series of game scores and a series of information values per game.

The analysis of the games between Gliders and Oxsy shows that a sufficiently high correlation ($\rho_1 = 0.425$) exists between the game score and only one relative responsiveness $\Delta(\text{Gliders}_{def}, \text{Oxsy}_{att})$. That is, the games between these two teams are decided mostly in the opposition between Gliders' defenders and Oxsy's forwards. Specifically one may conjecture that whenever the Oxsy's forwards are more responsive in evading the defense, Oxsy tend to win, and whenever Gliders' defenders are more agile in closing on to the forwards, Gliders tend to win.

However, the main information transfer component of $\Delta(\text{Gliders}_{def}, \text{Oxsy}_{att})$, correlated with the performance, is $t_{\text{Oxsy}_{att} \to \text{Gliders}_{def}}$, at 0.553 ("our responsiveness helps our scoreline"), while the correlation with $t_{\text{Gliders}_{def} \to \text{Oxsy}_{att}}$ is just 0.089 ("opponents responsiveness does not hurt our scoreline"). This means that on average the relative agility of Gliders' defenders is correlated with the scoreline more than the response of Oxsy's forwards. This is not a causal inference, but simply a correlation observation.

The dynamic contests between Gliders' forwards and Oxsy's defenders, or between the midfield players, do not seem to be greatly correlated with the scoreline on average (the scoreline is correlated with $\Delta(\text{Gliders}_{att}, \text{Oxsy}_{def})$ at just 0.099, and with $\Delta(\text{Gliders}_{mid}, \text{Oxsy}_{mid})$ at just 0.216). The transfer components of these characteristics do not show any higher correlations either. The overall tactical relative responsiveness $\Delta(\text{Gliders}, \text{Oxsy})$ is correlated with the scoreline at a credible level of 0.310.

As expected, the relative rigidity $A(\text{Gliders}, \text{Oxsy}) = A_{\text{Gliders}} - A_{\text{Oxsy}}$ is observed to be highly anti-correlated with the scoreline: $\rho = -0.641$. The main contributing part is found to be the rigidity of the mid-fielders: the correlation of rigidity $A(\text{Gliders}_{mid}, \text{Oxsy}_{mid})$ with the performance is also quite high at -0.503, and the major component of this comes due to the rigidity of Oxsy's mid-fielders: correlation of $A(\text{Oxsy}_{mid})$ is 0.377 (it is positive as the scoreline is presented as Gliders vs Oxsy, so that higher game scores for Gliders are correlated with higher Oxsy's rigidity).

The tactical analysis of the games between Gliders and Marlik produces mostly concurring observations. In this pair, the outcome is mostly decided in the contest between Gliders' attack and Marlik's defense: the correlation between relative responsiveness $\Delta(\text{Gliders}_{att}, \text{Marlik}_{def})$ is low but statistically significant: 0.157. Interestingly, however, both individual components are anti-correlated with the scoreline: the transfer $t_{\text{Marlik}_{def} \to \text{Gliders}_{att}}$ is anti-correlated at -0.210, and the transfer $t_{\text{Gliders}_{att} \to \text{Marlik}_{def}}$ is anti-correlated at -0.366.

This poses an interesting question: why two individual components of the relative responsiveness are both anti-correlated with the scoreline, but their combination is positively correlated, albeit at a low level? One possible explanation is as follows. Both involved groups (Gliders forwards and Marlik defenders) are in almost constant interdependent motion that often confounds the players. When Marlik defenders respond to Gliders forwards' attempts to find free spots, they effectively mark and/or block the forwards, resulting in lower scores for Gliders team — hence, the negative correlation between $t_{\text{Gliders}_{att} \rightarrow \text{Marlik}_{def}}$ and the scoreline, which is seen from the Gliders' perspective ("opponents responsiveness hurts our scoreline"). However, when Gliders forwards respond to Marlik defenders' attempts to mark them, they may abandon good scoring positions, also resulting in lower scores for Gliders team — hence, the negative correlation between $t_{\text{Marlik}_{def} \rightarrow \text{Gliders}_{att}}$ and the scoreline ("our responsiveness also hurts our scoreline"). Nevertheless, when the scoreline is correlated with the relative responsiveness, rather than the individual components of the latter, the result is positive but low. This means that the remaining difference is still slightly important because of the interdependence of motion: when Gliders forwards reposition, they attract Marlik defenders again, and the 'circle' repeats, until one side gains a brief advantage ("when our responsiveness is higher than opponents responsiveness, it helps our scoreline"). In short, it is not the level of our responsiveness that is positively correlated with the scoreline, but the level of relative responsiveness.

There are no surprises with the analysis of rigidity: the relative rigidity $A(\text{Gliders}, \text{Marlik}) = A_{\text{Gliders}} - A_{\text{Marlik}}$ is highly anti-correlated with the scoreline: $\rho = -0.505$. It is interesting that (relative) rigidities of separate groups (defense, mid-field, attack) were not found to be correlated significantly with the outcomes: the dependence is detected only at the overall team level, being arguably an emergent property in this contest.

In summary, the findings demonstrate applicability of the information dynamics measures to analysis of football matches, revealing the areas of most intense competition and the extent of interactions. The latter aspect is evident when one compares the interpretations of relative responsiveness in the games against Oxsy and against Marlik. The higher responsiveness of Gliders' defenders to Oxsy's forwards was found to be positively correlated with the scoreline, while the higher responsiveness of Gliders' forwards to Marlik's defenders was anti-correlated (as was the responsiveness of Marlik's defenders to Gliders' forwards). The difference shows that in the first case responses were productive and the interaction was clearly directional, while in the second instance, the responses were strongly interdependent and the interaction was quite circular. These observations are also supported by the interaction diagrams: both information-sink and information-base diagrams 1a and 2a for Gliders vs Oxsy show that Gliders' defenders respond strongly to Oxsy's forwards, while the information-sink and information-base diagrams 1b and 2b for Gliders vs Marlik highlight the extent of cross-coupling between Gliders' forwards and Marlik's defenders. The anti-correlation of rigidity in both experimental set-ups is also encouraging: this measure can be suggested as a simple robust measure of tactical flexibility, at the emergent team level.

4 Conclusion

The paper proposed an approach for constructing interaction networks that reveal significant coupled dynamics produced during team games, or other activities that are characterised by concurrent cooperation and competition. The approach uses a novel application of information dynamics analysing pair-wise interactions and group-level tactics of RoboCup 2D Simulation League games. The input data needed for the analysis contain only positional data, such as planar coordinates and their changes, followed by computation of corresponding probability distributions and local information transfer measures. The model-free approach does not include any re-construction of the players' behaviour, being purely data-driven. Also, the method is not aimed at explicit interactions (such as passes) within a team (cf. [3]), but rather at implicit interactions, across teams, that may be delayed and/or long-ranged.

The interaction networks were exemplified with two sub-types highlighting different "slices" of the directed interactions: information-sink and information-base diagrams. In an information-sink diagram every node (every player) has an incoming edge, while in an information-base diagram every node has an outgoing edge. These diagrams were computed for two experimental set-ups that matched our team (Gliders) against two well-known teams, Oxsy [18] and Marlik [19], showing interesting player-to-player interactions, and pointing out weak spots and areas to be exploited.

The follow-up tactical analysis involved computation of information transfer and storage, and two hypotheses. The first one related positive information transfer from players Y to players X as an indication of responsiveness of the latter, suggesting to compute relative responsiveness between the opposing lines of two teams. The second hypothesis connected the information storage within the team with the team's rigidity, harming the fluidity and tactical richness of the team. This relation yielded the score for relative rigidity between the opposing teams. Both measures, relative responsiveness and rigidity, were correlated with the game results, and the obtained observations supported the hypotheses. In addition, the results pointed to important couplings that were particularly intense, and the main areas where the game outcomes were mostly decided.

This approach has been further successfully applied to opponent modelling and selecting the best available tactics in an opponent-specific way — this topic is a subject of future research. We hope that the proposed methods would be useful not only in the RoboCup leagues, but also in various analyses of team games, whether virtual or real.

Acknowledgments. This analysis was carried out using Gliders2013 [20], a follow up on Gliders2012 [11] — a simulated soccer team for the RoboCup soccer 2D simulator [21]. Gliders2012 and Gliders2013 reached semi-finals in RoboCup tournaments of 2012 and 2013. The team code is written in C++ using agent2d: the base code developed by Akiyama et al. [22], fragments of released source code of Marlik [19], as well as Gliders' in-browser basic soccermonitor (GIBBS): a log-player for viewing 2D Simulation League logs over web browser [23].

We thank Ivan Duong, Edward Moore and Jason Held for their contribution to Gliders2012 and GIBBS, as well as David Budden for his help with developing new self-localisation method. Team logo was created by Matthew Chadwick. Some of the Authors have been involved with RoboCup Simulation 2D in the past, however the code of their previous teams (Cyberoos and RoboLog, see, e.g., [10,24]) is not used in Gliders.

References

1. Vilar, L., Araujo, D., Davids, K., Bar-Yam, Y.: Science of winning soccer: Emergent pattern-forming dynamics in association football. Journal of Systems Science and Complexity 26, 73–84 (2013)
2. Grabiner, D.J.: The Sabermetric Manifesto, The Baseball Archive
3. Fewell, J., Armbruster, D., Ingraham, J., Petersen, A., Waters, J.: Basketball teams as strategic networks. PLoS One 7(11), e47445 (2012)

4. Noda, I., Stone, P.: The RoboCup Soccer Server and CMUnited Clients: Implemented Infrastructure for MAS Research. Autonomous Agents and Multi-Agent Systems 7(1-2), 101–120 (2003)
5. Riley, P., Stone, P., Veloso, M.: Layered disclosure: Revealing agents' internals. In: Castelfranchi, C., Lespérance, Y. (eds.) Intelligent Agents VII. LNCS (LNAI), vol. 1986, pp. 61–72. Springer, Heidelberg (2001)
6. Stone, P., Riley, P., Veloso, M.: Defining and using ideal teammate and opponent models. In: Proc. of the 12th Annual Conf. on Innovative Applications of Artificial Intelligence (2000)
7. Butler, M., Prokopenko, M., Howard, T.: Flexible synchronisation within RoboCup environment: A comparative analysis. In: Stone, P., Balch, T., Kraetzschmar, G.K. (eds.) RoboCup 2000. LNCS (LNAI), vol. 2019, pp. 119–128. Springer, Heidelberg (2001)
8. Reis, L.P., Lau, N., Oliveira, E.C.: Situation based strategic positioning for coordinating a team of homogeneous agents. In: Hannebauer, M., Wendler, J., Pagello, E. (eds.) Reactivity and Deliberation in MAS. LNCS (LNAI), vol. 2103, pp. 175–197. Springer, Heidelberg (2001)
9. Prokopenko, M., Wang, P.: Relating the entropy of joint beliefs to multi-agent coordination. In: Kaminka, G.A., Lima, P.U., Rojas, R. (eds.) RoboCup 2002. LNCS (LNAI), vol. 2752, pp. 367–374. Springer, Heidelberg (2003)
10. Prokopenko, M., Wang, P.: Evaluating team performance at the edge of chaos. In: Polani, D., Browning, B., Bonarini, A., Yoshida, K. (eds.) RoboCup 2003. LNCS (LNAI), vol. 3020, pp. 89–101. Springer, Heidelberg (2004)
11. Prokopenko, M., Obst, O., Wang, P., Held, J.: Gliders2012: Tactics with action-dependent evaluation functions. In: RoboCup 2012 Symposium and Competitions: Team Description Papers, Mexico City, Mexico (June 2012)
12. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Local information transfer as a spatiotemporal filter for complex systems. Physical Review E 77(2), 026110 (2008)
13. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Information modification and particle collisions in distributed computation. Chaos 20(3), 037109 (2010)
14. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Coherent information structure in complex computation. Theory in Biosciences 131, 193–203 (2012)
15. Lizier, J.T., Prokopenko, M., Zomaya, A.Y.: Local measures of information storage in complex distributed computation. Information Sciences 208, 39–54 (2012)
16. Schreiber, T.: Measuring information transfer. Phys. Rev. Lett. 85(2), 461–464 (2000)
17. Wang, X.R., Miller, J.M., Lizier, J.T., Prokopenko, M., Rossi, L.F.: Quantifying and tracing information cascades in swarms. PLoS One 7(7), e40084 (2012)
18. Marian, S., Luca, D., Sarac, B., Cotarlea, O.: Oxsy 2011 team description paper. In: RoboCup 2011 Symposium and Competitions: Team Description Papers, Istanbul (2011)
19. Tavafi, A., Nozari, N., Vatani, R., Yousefi, M.R., Rahmatinia, S., Pirdir, P.: MarliK 2012 Soccer 2D Simulation Team Description Paper. In: RoboCup 2012 Symposium and Competitions: Team Description Papers, Mexico City, Mexico (June 2012)
20. Prokopenko, M., Obst, O., Wang, P., Budden, D., Cliff, O.: Gliders2013: Tactical Analysis with Information Dynamics. In: RoboCup 2013 Symposium and Competitions: Team Description Papers, Eindhoven (June 2013)
21. Chen, M., Dorer, K., Foroughi, E., Heintz, F., Huang, Z., Kapetanakis, S., Kostiadis, K., Kummeneje, J., Murray, J., Noda, I., Obst, O., Riley, P., Steffens, T., Wang, Y., Yin, X.: Users Manual: RoboCup Soccer Server — for Soccer Server version 7.07 and Later (February 2003)
22. Akiyama, H.: Agent2D Base Code (2010), `http://www.rctools.sourceforge.jp`
23. Moore, E., Obst, O., Prokopenko, M., Wang, P., Held, J.: Gliders2012: Development and competition results. CoRR abs/1211.3882 (2012)
24. Obst, O., Boedecker, J.: Flexible coordination of multiagent team behavior using HTN planning. In: Bredenfeld, A., Jacoff, A., Noda, I., Takahashi, Y. (eds.) RoboCup 2005. LNCS (LNAI), vol. 4020, pp. 521–528. Springer, Heidelberg (2006)