

MOTIFS IN EVOLVING COOPERATIVE NETWORKS LOOK LIKE PROTEIN STRUCTURE NETWORKS

DAVID HALES

Technical University of Delft
Dept. EWI/PDS/Tribler, PO Box 5031, 2600 GA Delft. Netherlands

STEFANO ARTECONI

University of Bologna
Mura Anteo Zamboni 7, 40127 Bologna, Italy

ABSTRACT. The structure of networks can be characterized by the frequency of different subnetwork patterns found within them. Where these frequencies deviate from what would be expected in random networks they are termed “motifs” of the network. Interestingly it is often found that networks performing similar functions evidence similar motif frequencies. We present results from a motif analysis of networks produced by peer-to-peer protocols that support cooperation between evolving nodes. We were surprised to find that their motif profiles match closely protein structure networks. It is currently an open issue as to precisely why this is.

1. Introduction. In previous works we presented a copy and re-wire peer-to-peer (P2P) protocol that promoted cooperation between connected nodes even when they had incentives to behave selfishly [2, 5]. The protocol structures the population into competing ‘tribes’ that, through a group-like selection process, lead to socially beneficial strategies even when the individual nodes behave in an essentially evolutionary way - copying the behaviours and links of other nodes that outperform them.

We tested the protocol by having nodes play the Prisoner’s Dilemma game (PD) - a canonical game for exploring situations in which collective and individual interests diverge. We called this protocol SLACER (Selfish Link Adaptation for Cooperation) and applied it to a simulated file-sharing scenario demonstrating it had the ability to control the outbreak of selfish strategies by nodes (downloading without uploading, so called “leeching”) [6].

Our previous analysis of the network structure indicated that SLACER networks were fully connected with cooperative paths between almost all nodes and that the network followed a small-world topology (high clustering and short average path lengths).

In order to examine the structure of SLACER networks further we applied a motif analysis technique which produces subgraph ratio profiles (SRP) [10]. This approach allows for the classification of networks into similar functional groups -

2000 *Mathematics Subject Classification.* Primary: 91A22, 91A43; Secondary: 91D30.

Key words and phrases. Evolutionary games, games on graphs, social networks.

Partially supported by the EU within the 6th Framework Programme under contract 001907 “Dynamically Evolving, Large Scale Information Systems” (DELIS). .

```

Periodically for node  $i$ :
  select a random node  $j$  from the population
  if utility  $U_i \leq$  utility  $U_j$  then
    copy strategy from  $j$ 
    drop each link from  $i$  with prob( $W$ )
    copy each link from  $j$ 
    link to  $j$ 
    with prob( $M$ ) mutate strategy of  $i$ 
    with prob( $MR$ ) mutate links of  $i$ 

```

FIGURE 1. The SLACER protocol pseudocode. For an overview of the protocol see the text. For more detail see [5].

so called “superfamilies” of networks. We found that SACER produced networks which fit within a superfamily of networks derived from protein structure networks. We report this work here.

It is not currently understood why networks with similar functions have similar SRPs. In this paper we do not attempt to answer why our P2P protocols have similar features to protein structure networks although this could be an interesting topic for further study.

We introduce the SLACER protocol in section 2. We describe the PD cooperation game in section 3. The technique of motif analysis using SRPs is presented in section 4. Finally we show the results of the analysis and give some brief observations and conclusions.

2. SLACER protocol. The SLACER protocol follows a link based incentive approach. That is, nodes make and break links in the network in order to minimise the effect of selfish nodes. Hence the topology itself comes to reflect a network of cooperation.

The protocol models node “behaviour” as following an evolutionary approach. Nodes periodically compare their utility - derived by some functional interaction with their neighbours - with another randomly selected node. If the utility of the other node is higher then nodes copy the interaction strategy and neighbour links of the other node. With low probability they apply “mutation” to their strategy and links. In this sense the protocol does not model individual selfish behaviour but rather evolutionary learning. Nodes do not have the ability to change their strategies based on individual learning or assessment.

Figure 1 shows the pseudocode. Over time, nodes engage in some application task and generate some measure of utility U . This utility is a numeric value that each node needs to calculate based on the specifics of the particular application domain. For example, this might be number of files downloaded, jobs processed or an inverse measure of spyware infections over some period. The higher the value of U the better the node believes it is performing in its target domain.

Periodically, each node i compares its performance against another node j , randomly selected from the population. If $U_i \leq U_j$ then i drops each of its current links to other nodes with high probability W , and copies all j ’s links and adds a link to j itself. Additionally i then copies j ’s strategy - the strategy codes some application level behaviour. After such a copy operation has occurred, then, with

low probability M , node i adapts its strategy and with probability MR adapts its links. Adaptation involves the application of a “mutation” operation. Mutation of the links involves removing each existing link with probability W and adding a single link to a node randomly drawn from the network. Mutation of the strategy involves applying some form of change in application behaviour with probability M - the specifics of strategy mutation are dictated by the application domain (see later). After the periodic utility comparison, the node resets its utility to zero. Previous “tag” models, on which SLACER is based [6] have indicated that the rate of mutation applied to the links needs to be significantly higher than that applied to the strategy by about one order of magnitude hence $MR \gg M$.

Each node is limited to a maximum number of links or neighbours (view size). If any SLACER operation causes a node to require an additional neighbour above this limit then a randomly selected existing link is removed to make space for the new link. Links are always undirected, hence symmetrical, so that if node i links to node j , then j must also maintain a link to node i and conversely if node i breaks a link to node j then node j also breaks its link to node i . In the work presented here each node has a maximum view size of 20 links.

When applied in a suitably large population the algorithm follows an evolutionary process in which nodes with high utility replace nodes with low utility. However, as will be seen, this does not lead to the dominance of selfish strategies, as might be intuitively expected, because a form of social incentive mechanism results from the emergent network topology. This means that high utility but anti-social strategies, even though favoured by the individual nodes, do not dominate the population. The topology therefore guides the adaptation of the strategy away from anti-social selfish strategies.

When the link drop probability $W = 1$ (see figure 1) SLACER produces highly cooperative yet disconnected networks. However, when W is slightly reduced, SLACER produces networks in which almost all nodes are members of a giant connected and cooperative component. SLACER networks are also small-world, with a low average path length between nodes, growing logarithmically with network size, and a high clustering coefficient (remaining constant with respect to network size). We found, previously, that varying the main parameters of the model W and M and MR did not effect these topological characteristics greatly so long as $0 < M < 0.01$ and $0.1 > MR \gg M$ and $W > 0.5$. Here, however, we use typical parameter values from previous work that promote high levels of cooperation between nodes while keeping the network connected - see [5, 2]. We did not vary the parameters to produce any kind of fitting with the protein structure networks.

SLACER networks promote cooperation between nodes in a number of tested scenarios. In addition they have desirable properties of scalability, self-organisation and robustness. This means that if nodes are inserted or removed, or links broken, the network quickly readjusts back into a cooperative state.

Although we have tested the protocols in P2P-like scenarios, such as file sharing [6] and cooperative job sharing [4] we evaluated them initially with an abstract cooperation game called the Prisoner’s Dilemma (PD). This is a canonical test for the emergence of cooperation between individuals when there is an incentive for non-cooperation. Results presented here are based on nodes playing the PD with neighbours in their network. In the next section we describe the two-player PD game.

3. Prisoner’s Dilemma. The two players single-round Prisoners Dilemma (PD) game captures a situation in which there is a contradiction between collective and individual self-interest. Two players interact by selecting one of two choices: to “cooperate” (C) or “defect” (D). For the four possible outcomes of the game, players receive specified payoffs. Both players receive a reward (R) or a punishment (P) payoff for mutual cooperation and defection respectively. However, when individuals select different moves, different payoffs of temptation (T) and sucker (S) are awarded to the defector and the cooperator respectively (see table 1). Assuming that neither player can know in advance which move the other will make and wishes to maximise her own payoff, the dilemma is evident in the ranking of payoffs: $T > R > P > S$ and the constraint that $2R > T + S$. Although both players would prefer T , because its the highest payoff, only one can attain it in a single game. No player wants S because its the lowest payoff. No matter what the other player does, by selecting a D move a player always gets a higher score than it would have obtained if it had selected C . D is therefore the dominant strategy, hence an ideally rational player would always choose D , and the only Nash equilibrium in the game is given by the pair of strategies (D, D)

	C	D
C	R,R	S,T
D	T,S	P,P

TABLE 1. Prisoner’s Dilemma payoff table.

Therefore, the dilemma is that if both players select a cooperative (C) move they are jointly better off (getting R each) than if they both select D , but selfish players will select mutual defection, getting only P each, because of the individual incentive to select defection. We select this game as a minimal test that captures a range of possible application tasks in which nodes need to establish cooperation and trust with their neighbours but without central authority or external mechanisms that enforce it.

In order to apply the game to test the SLACER P2P protocol we set application to each node playing the PD with randomly selected network neighbours. A node can only choose one out of the two pure strategies: cooperate or defect. The utility value required by each node in the protocols is set to the average payoff the node received from recent game interactions. The SLACER algorithm then adapts the links and strategy of the nodes as discussed previously. The specific payoff values chosen for T, R, P and S do not significantly effect performance as long as they meet the PD constraints discussed above. Figure 2 shows the typical evolution of the network over time.

4. Motifs and subgraph ratio profiles. Networks are often characterised using average global measures, such as average path length and clustering coefficient. Although valuable such measures rarely give a picture of the detailed structure of the networks. This means that networks with different topologies on a local scale, can have identical global average measurements. Hence, in order to further understand and classify natural and artificial networks new methods have been proposed.

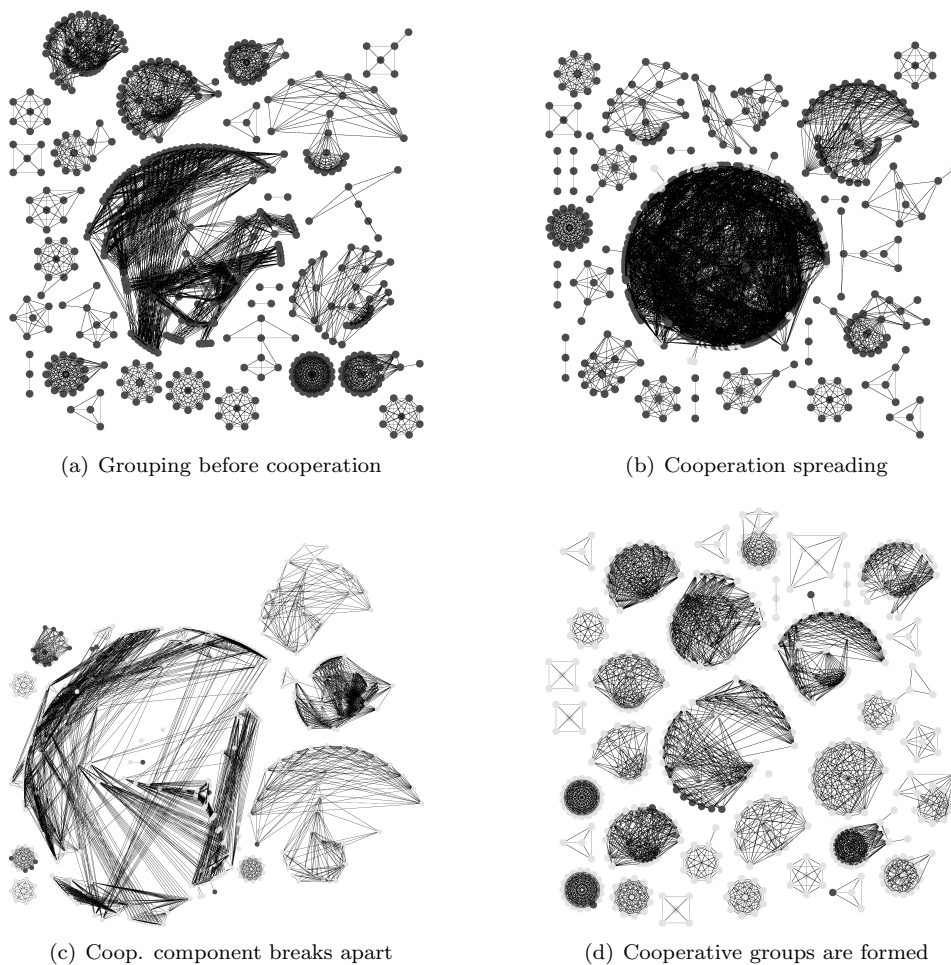


FIGURE 2. Evolution of a SLACER network with nodes playing the Prisoner's Dilemma. From an initially random topology composed of all nodes playing the defect strategy (dark shaded nodes), components quickly evolve, still containing all defect nodes (a). Then a large cooperative component emerges in which all nodes cooperate (b). Subsequently the large component begins to break apart as defect nodes invade the large cooperative component and make it less desirable for cooperative nodes (c). Finally an ecology of cooperative components dynamically persists as new components form and old components die (d). Note: the cooperative status of a node is indicated by a light shade. For diagrammatic clarity we set $W = 1$ which leads to disconnected components but the dynamics for lower W are essentially the same but with components as cohesive clusters or cliques.




FIGURE 3. All six possible four-node undirected subgraphs (tetrads). Nodes are not shown but should be assumed at the end of each straight line. The adjacency matrix derived id's are not shown but the tetrads are ordered by ascending value of them.

Recently, researchers working with complex networks (both natural and artificial) have begun to analyse and characterise them using more sophisticated topological techniques and one of these approaches is so called “motif analysis” [11].

By breaking a network down into all possible n -node subgraph patterns and counting them it is possible to compare those counts against randomly generated networks with the same characteristics (number of nodes and in / out degree links). Then, where certain n -node subgraph patterns are significantly more prevalent than in the random case, these are considered motifs of the network. Additionally, although less discussed in the literature, n -node subgraph patterns that are under-represented in the network have been termed anti-motifs [10] and are of equal value in characterising network structure.

Obviously, for large subgraph sizes the number of possible motifs becomes large but for smaller sizes (3 and 4 nodes) it is computationally feasible to search, even large, networks for all occurrences.

The P2P networks produced by SLACER are undirected in the sense that all links are bidirectional. So for the purposes of analysis we search for all undirected four-node subgraphs (tetrads). Figure 3 shows the six possible undirected tetrads. We analysed the networks using the subgraph ratio profile (SRP) method [10]. This approach is particularly useful for undirected networks since traditional motif analysis methods using z-scores are not network size invariant for tetrads and this makes comparison with networks of different sizes difficult. For a given network N the SRP is a normalised vector of Δ_i values:

$$SRP_i = \frac{\Delta_i}{\sqrt{\sum_i \Delta_i^2}} \quad (1)$$

The vector elements, one for each of the six tetrads, are calculated based on the abundance of each tetrad i relative to randomly generated networks, To avoid large values as an artefact of very small occurrences of tetrads in both the real and random networks the value $\varepsilon = 4$ is added to the the denominator.:

$$\Delta_i = \frac{N_{real_i} - \langle N_{rand_i} \rangle}{N_{real_i} + \langle N_{rand_i} \rangle + \varepsilon} \quad (2)$$

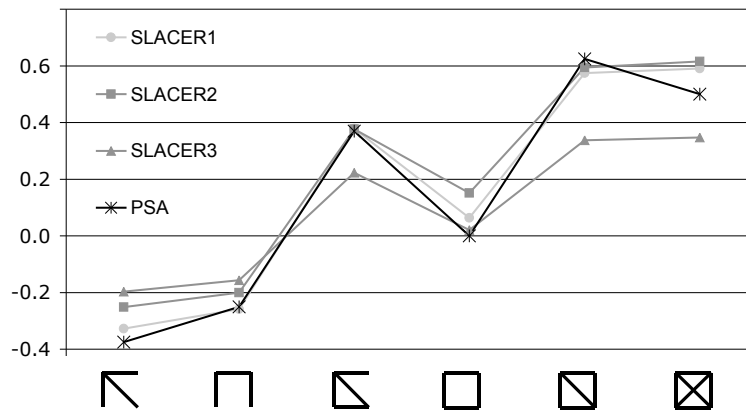


FIGURE 4. The subgraph ratio profiles (SRP) of SLACER P2P networks at different stages in their evolution to cooperation - network size $N = 500$, edges $E \approx 3500$. In each case three time ordered network snapshots are shown. The networks were initialised with random topologies and all nodes as non-cooperative. The first snapshot is taken immediately before any cooperation has been attained, the second during the rapid outbreak of cooperation and the final snapshot is taken after stable cooperation is attained. The results show a similarity to the protein-structure superfamily (PSA). PSA shows the average SRP of the protein structure results given in [10]. These results are from networks of secondary-structure elements adjacency for several large proteins [structure based on the PDB database (www.rcsb.org/pdb/); the proteins (and their PDB ID) were 1A4J, an immunoglobulin ($N = 95$, $E = 213$); 1EAW, a serine protease inhibitor ($N = 53$, $E = 123$); and 1AOR, an oxidoreductase ($N = 99$, $E = 212$)]

A given SRP can be graphed producing a curve which characterises the tetrad motifs and anti-motifs visually.

5. Results of analysis. Figure 4 show SRP's for SLACER P2P networks at different stages in their evolution. In each case three time ordered network snapshots are shown (SLACER1, SLACER2 and SLACER3). The application task that generated the node level utility required for the protocols was to periodically play the Prisoner's Dilemma (PD) game with a randomly chosen neighbour and accumulate average payoff. The period of game playing was one order of magnitude higher than the period used by the SLACER protocol as shown in figure 1. This means that on average ten games of PD would be played between SLACER invocations.

As previously discussed, SLACER is parameterized over that the link drop probability (W). In this case we used $W = 0.9$ since this has been found to produce high-levels of cooperation while retaining a highly connected network [5]. All nodes in the network stored a pure PD strategy of either to cooperate or defect - representing the application behaviour. The mutation probabilities were set to $M = 0.001$ and $MR = 0.01$. We fixed the maximum degree of each node to 20 links. All links are undirected and hence symmetric. The protocol perseveres the symmetry

of links between nodes at all times. It can be noted that our networks therefore have a higher average degree than the protein networks and it would be of interest to reduce the degree and see if the results we find here still hold (this is possible future work).

The networks were initialised with random topologies and all nodes as non-cooperative (defect strategy). The first snapshot is taken immediately before any cooperation has been attained, the second during the rapid outbreak of cooperation and the final snapshot is taken after stable cooperation is attained. We do not show the SRP for the initial network topology since this would be a flat line along the x-axis (being a random network). However, we have found that from *any* initial topology the network evolves to the same topology immediately before cooperation breaks-out [5] so we would expect the shown results to look the same no matter what the starting topology.

Notice that the curves in figure 4 follow a time evolution. Immediately before cooperation (SLACER1) the curve already has a very similar shape to the final curve, during the outbreak of cooperation (SLACER2) the curve tends to move upward slightly (less anti-motifs 1 and 2, but more motifs 3 to 6), then, after stable cooperation is attained (SLACER3) the curve tends to flatten (with all points moving towards the x-axis). Motifs 1 and 2 are under represented (anti-motifs) and motifs 3 to 6 are over represented (motifs) but with a large dip for motif 4 - almost close to zero (identical to the random occurrence) when cooperation has stabilised (SLACER3).

The SRP curves shown in figure 4 were directly compared to those given in [10]. There a number of curves were given for different networks. Similar curves were clustered into similar groups (or superfamilies). We have copied the average of the protein structure curves onto our figures (labeled as PSA). The protein structures are non-directed networks of protein structure in which nodes are secondary-structure elements (α helices and β strands) and two nodes are connected if their distance is smaller than 10\AA . The PSA line shows a characteristic anti-motif representation in 1 and 2 and dip at 4. These characteristics are reproduced in the P2P networks. However, notice that motif 6 dips slightly in the protein-structure networks (PSA) but this is not reflected in the P2P networks. Here the P2P networks look more like the BA models [1] or the power-grid networks shown given in [10]. Notice also that the early snapshots of SLACER look most like the protein-structure networks - interestingly this is before cooperation has stabilised. The SLACER1 networks are taken before *any* cooperation has formed in the networks and can therefore only be the result of random copy and rewire since all utilities are equal, at this stage, because all nodes use the defect strategy in the PD.

This suggests some rather intriguing new possibilities *and* reinforces existing findings. Firstly, that the motifs and anti-motifs evident in the P2P SRP's initially result from a randomised process rather than the specific function of cooperation *but* that this randomised process is an artefact of an evolutionary algorithm designed *for* cooperation.

SLACER appears to work *because* a randomised re-wiring process draws the network towards a topology that eventually supports cooperation. We have discussed in detail this multi-stage process elsewhere [5]. What we see is a complex interplay of function and topology formation leading to a phase transition in cooperation with negative scaling properties (i.e. the bigger the network the quicker it happens). Essentially, what SLACER does is structure the population into cooperative cliques

- from any initial network topology. Cooperation is then produced by a kind of group selection over those cliques because cliques that contain defectors are quickly dissolved as nodes move away to better (high utility producing) clusters. This is a *kind of* group selection because actual selection is only operating at the individual node level but the emergent clustering promotes group pro-social behavior - even when there are incentives to cheat.

Another interesting result is that the different snapshots show significant differences in the SRP, such that for a given (P2P network) SRP it would be possible to predict if it was in a stable cooperative mode or not - without considering the strategies stored in the nodes. This could possibly be very useful, since it indicates a way to detect if the network is functioning cooperatively or which stage in the evolution to cooperation has been reached based purely on structural characteristics. Since we have observed SLACER to support other kinds of cooperative tasks (not just playing the PD) and produce networks with similar topologies we speculate that potentially the final snapshot 3 SRP curve could characterise many possible application domains that were functioning under SLACER cooperatively (correctly), *even if the specifics of the applications themselves were complex or unknown or unknowable from a snapshot*. If this were the case, then a SRP curve could provide a very powerful fingerprint of the global collective functioning of a SLACER supported application. A deviation from the cooperative fingerprint could indicate malfunctioning or network attacks. This could be a topic for future work along with the design of a mechanism to perform motif analysis in a distributed fashion requiring all nodes to interact locally and aggregate their knowledges on local structures.

We have not performed an extensive comparison of other network statistics such as degree distributions, clustering coefficients and average path lengths. The networks currently used resist close comparison due to the differences in average degree k between the protein networks and SLACER networks. One aspect of future work could be to produce SLACER networks that more closely mirror the protein structure networks.

5.1. Conclusions. We have briefly summarised a P2P protocol that we have designed to promote cooperation - system level utility - when nodes behave in an evolutionary way with local information. We performed a motif analysis by calculating and plotting subgraph ratio profiles as curves. We compared this to existing work showing curves for various natural and engineered networks. We did not expect to find the P2P networks to match any of the existing graphs but were intrigued to find that the P2P networks appeared very similar to a family of protein-structure curves. However, should we be surprised? The P2P algorithms were inspired by an evolutionary tagging algorithm [3] that produces robust and desirable life-like properties such as scalability, robustness and self-repair. It therefore may not be so surprising that the resulting topologies come to resemble naturally occurring networks with similar properties.

Rather than design an algorithm for constructing a particular network *form* [1] we designed the P2P protocol for a specific *function* - to suppress selfish strategies of nodes for the collective good of the network without the need for complex or sophisticated reputation systems, repeated reciprocal strategies or centralised control. Perhaps this function has some kind of universal applicability - if so, could this throw any light on protein networks and their role in cell level computation? Is there really any linkage at all or just superficial similarity? It may be the case that

given the low fidelity of this kind of motif analysis (i.e. there are only six possible 4-node undirected subgraphs) almost any kind of “copy and rewire” operation would lead to similar profiles. We therefore seriously qualify any conclusions we draw at this stage of our work.

Since SLACER works via repeated copying of links between nodes then it would appear that the over expression of highly connected motifs would be an inevitable by-product of the process. This is evidenced by the fact that the initial SRP’s, before cooperation had occurred, already look like the protein structure networks. Further, it has been demonstrated that simple geographically constrained randomised models produce very similar topological properties to protein structure networks [10, 12]. In such geographical models, nodes are distributed randomly in unit space and pairs are selected randomly with some distance based probability that they are connected. In this sense perhaps our contribution might be to show how a simple evolutionary process - starting without any particular geography - emerges a kind of geographical network.

One interesting idea that this work suggests is that although networks may appear to be constructed from a random process of copy and re-wire this could be an *artefact of an underlying functional process* - which in certain phases reduces to a random process but nevertheless harnesses the functionality of the properties produced by that process. In some sense, this is what evolution is. If we were to make a grandiose conjecture we might say that the kind of SRP profiles we have found could be considered as evidence for (though not proof of) an evolutionary process in a network, specifically where robustness to noise is desirable. Recent work may lend some support to the conjecture. Motif distributions of designed networks, obtained through an evolutionary optimization process, were compared with those of the biological networks of a biological cell and a good agreement was found [9, 8].

Another potentially very useful line of work (for P2P engineering) would be to explore our hunch that the cooperative fingerprint shown in the SRP will hold over many possible application domains, hence providing a method for *monitoring the network for serious malfunctions or malicious attacks that degrade collective performance, generically, without having to know the application specifics*. This would be consistent with the thesis of [10] which is that perhaps the super-families identify some very general underlying functional characteristics even though we might not know the specifics. However, given we know that randomised behaviour can lead to similar profiles this kind of approach would only make sense if we knew that nodes were following at least some aspects of the specified protocol. This could be claimed as begging the question.

Our simple analysis has raised a number of questions, rather than providing answers, that could be an interesting topic of future work. P2P networks are increasingly employed to provide functionality over the Internet. Although designed by human engineers their emergent behaviours and structures are rarely understood due to their scale, dynamically and open nature. Perhaps application of network structure analysis could provide new insights.

Acknowledgements. The P2P protocol simulations were implemented in Peersim¹. The motif analysis was carried out using the Mfinder software² and the graph visualisations in figure 2 were produced using Graphviz³. We would like to thank Sergi Valverde and Ricard Solé for their collaboration - introducing us to motif analysis techniques and tools and hosting us briefly at UPF, Barcelona. We thank the attendees of a number of workshops, and the European Conference on Complex Systems (2007) at which we presented this work, for their comments, interest and encouragement. Finally we thank the anonymous reviewers who commented on earlier drafts, pointed us to relevant work and made us aware of omissions and errors.

REFERENCES

- [1] A. Barabási and R. Albert, *Emergence of scaling in random networks*, Science, **286** (1999), 509–512.
- [2] D. Hales, *From selfish nodes to cooperative networks – emergent link-based incentives in peer-to-peer Networks*, in “Proceedings of The Fourth IEEE International Conference on Peer-to-Peer Computing,” IEEE Computer Society Press, (2004).
- [3] D. Hales, *Cooperation without space or memory: tags, groups and the prisoner’s dilemma*, in “Multi-Agent-Based Simulation, Lecture Notes in Artificial Intelligence 1979” (eds. S. Moss, P. Davidsson), Springer-Verlag, (2000).
- [4] D. Hales, *Emergent group-level selection in a peer-to-peer network*, Complexus, **3** (2006), 108–118.
- [5] D. Hales and S. Arteconi, *Slacer: a self-organising protocol for coordination in P2P networks*, IEEE Intelligent Systems, **21** (2006), 29–35.
- [6] D. Hales and B. Edmonds, *Applying a socially-inspired technique (tags) to improve cooperation in P2P networks*, IEEE Transactions in Systems, Man and Cybernetics - Part A: Systems and Humans, **35** (2005), 385–395.
- [7] M. Jelasity, A. Montresor and O. Babaoglu, *Gossip-based aggregation in large dynamic networks*, ACM Trans. Comput. Syst., **23** (2005), 219–252.
- [8] P. Kaluza, M. Ipsen, M. Vingron and A. S. Mikhailov, *Design and statistical properties of robust functional networks: A model study of biological signal transduction*, Phys. Rev. E, **75** (2007), 015101.
- [9] P. Kaluza and A. S. Mikhailov, *Evolutionary design of functional networks robust against noise*, Europhys. Lett., **79** (2007), 48001.
- [10] R. Milo, S. Itzkovitz, N. Kashtan, R. Levitt, S. Shen-Orr, I. Ayzenshtat, M. Sheffer and U. Alon, *Superfamilies of designed and evolved networks* Science, **303** (2004), 1538–1542.
- [11] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii and U. Alon, *Network motifs: simple building blocks of complex networks*, Science, **298** (2002), 824–827.
- [12] B. Stacey, *On motif statistics in symmetric networks*, in “Proceedings of the Sixth International Conference on Complex Systems” (eds. Ali Minai, Dan Braha, Yaneer Bar-Yam), NECSI, (2006).

Received August 2007; revised January 2008.

E-mail address: dave@davidhales.com

E-mail address: arteconi@cs.unibo.it

¹The open source Peersim environment and the code for the SLACER protocol, in addition to tutorial materials are available at: <http://peersim.sf.net>

²The Mfinder network motif detection tool software is freely available at: <http://www.weizmann.ac.il/mcb/UriAlon/groupNetworkMotifSW.html>

³Graphviz graph visualisation software is freely available at: <http://www.graphviz.org/>