Exploring Social Networks in Request for Proposal Dynamic Coalition Formation Problems

Carlos Merida-Campos and Steven Willmott

Universitat Politècnica de Catalunya, Software Department, E-08034 Barcelona, Spain
{dmerida,steve}@lsi.upc.edu

Abstract. In farsighted MAS, every agent is aware of each other, and so they can evaluate the outcomes of their interaction in an efficient way. However, this farsighted knowledge becomes an issue in large scale systems, leading to combinatorial explosion. Limited awareness of agents can be modeled as a sparse social network solving this way the combinatorial explosion issue. In this paper, a model of MAS dynamic coalition formation is adopted and experiments with fixed underlying social networks exhibit different exploratory behaviors such that *Small World* and *Scale Free* topologies, that have shown their beneficial properties in many other MAS environments, turn out to have several drawbacks in the coalitional scenario.

1 Introduction

In Multi-Agent Systems (MAS), agents need to interact in order to fulfil their common or individual goals. In some cases, an agent does not need to interact with every other agent, and social knowledge can be spread in the basis of some neighboring concept such as functional neighboring (as in supply chain models [18]), or geographical neighboring (as in sensor networks [5]). In other cases, agents might benefit from being aware of every other agent of the population (*farsighted* social knowledge). This is the case of organisational systems [11, 14, 15] where agents have to explore a search space of agent group combinations in order to improve data flow, allocate resources efficiently, solve a problem in a coordinated manner, or to improve their outcomes by creating alliances. However, farsighted social knowledge is not possible in large scale MAS, because it might lead to combinatorial explosion and computational intractability. To address this issue in the concrete organisational environment of *Coalition Formation*, the use of a dynamic formation mechanism (Iterated RFP [14]) where social awareness of agents is limited to a fixed social network is proposed.

The work presented here, covers results on the properties that specific network topologies exhibit under different exploratory behaviors. Instead of having agents that dynamically adjust their social connections [13, 9] or that propagate/contaminate their neighbors through a social network [6, 17], agents in the present model are constrained by a unmovable social network where no social capital is transmitted. However, agents can iteratively form coalitions with their neighbors, leading to an evolutionary process that converges into a Nash equilibrium. This dynamic gives shape to coalitions by putting together agents that can be far away from each other in the network. Results obtained show that network parameters like node-degree distribution or clustering coeficient in concrete nodes, determine the exploratory behavior in the social network of the population. One of the main findings presented is the disadvantage of highly clustered topologies. This result coincides with what has been recently reported in [13] and in [7], where it is explained why certain human-civilisations with scattered connections between their regions, triumphed over others optimally connected. As in that model, the present work shows how heterogeneously distributed connectivity fosters higher exploration in the system. However, the model used in the current work, has a significant difference with the previous work, that is the coalitional nature of the protocol used.

Coalition Formation area have been traditionally studied under the assumption of farsightedness, focusing on the problem of finding stability concepts . MAS research introduced the possibility of experimenting with coalitional systems with a limited number of possible interactions [1, 16], and more recently this myopic sight of agents have been shaped with concrete knowledge network topologies in team formation [9, 10] as well as in firm formation models [2]. In this line, the current work presents results on experimentation with different underlying social network topologies on an specific type of electronic market allocation mechanism called Iterated Request For Proposal (RFP from now on). This model was first studied in [12], and further explored in [14] and [15]. In this environment, an entity regularly issues a call for tender to provide specific goods or services with certain characteristics. Providers compete amongst themselves (either individually or in consortia – *coalitions*). Their bids are ranked according to an evaluation of their skills (or aggregation of skills) for the task, and receive a payoff according to their placement in the ranking. Structures created are based on complementariety of their members. The more complementary they are, the better outcome they obtain, assuming this way a social structure with proved real properties reported in [4]. There are many existent real systems that follow the RFP type procedures such as public building projects, competitive tender for government contracts or even collaborative research project grants. RFP environments can also be seen as emerging market opportunities in an economy, with individual calls for tender representing new opportunities for profit.

Section 2 presents a formalisation of the RFP mechanism. Section 3 describes an specific metric that records the amount of dynamism that a certain topology generates as well as other metrics used to perform the analysis and the experimental setup. Results are analysed in subsection **??**. Apart from examining the drawbacks of the Small World topology, section 3.2 analyses results on the importance of positioning of agents with specific individual properties (*versatility* and *competitiveness*) in certain parts of the social network. Finally, section 4 discusses the results summarising the main conclusions and explaining future lines of work.

2 Iterative RFP Coalition Formation Model

A population $I = \{1, 2, \ldots, n\}$ consists of a finite number of n individuals or *agents*. Agents compete for creating the best solution proposal to a given task *task* T. A *partition* $\sigma = \{\sigma_1, \sigma_2, \ldots, \sigma_p\}$ of the population I is a specification of p *coalitions* of agents $\sigma_i = \{\sigma_{i1}, \sigma_{i2}, \ldots, \sigma_{im}\}$, where σ_{ij} represents an agent from population I forming part of coalition σ_i. By forming coalitions, agents are able to increase their competitiveness towards the specified task. Agents have heterogeneous capabilities, thus having different performance levels in different skills. A finite number of k *skills*, indexed from 1 to k is set for which each agent σ_{ij} has a fixed value: $\sigma_{ij} = \langle \sigma_{ij}^1, \sigma_{ij}^2, \ldots, \sigma_{ij}^k \rangle$. In this way it is possible to define a continuum of possibilities between agents that are *specialised* in the performance of a certain skill being unskilled for the rest of them, and agents that are *versatile*, being averagely apt for the performance of all the skills defined. A Task T is specified by a set of k skill requirements. Each one of the k skills have a degree of requirement. These requirements are modelled in the form of a number $T = \langle T^1, T^2, \ldots, T^k \rangle$. In a coalition, skills of agents are aggregated in such a way that each agent gives the best of itself in a join effort to create a group as competitive as possible under the requirements of the Task. The coalition has a value in each skill representing the aggregated effort of its members. The aggregation for every skill $l : 1 \leq l \leq k$ in the coalition is modelled in the following way:

$$\sigma_i^l = \max(\sigma_{ij}^l) : 1 \leq j \leq m \tag{1}$$

Each skill is considered as a necessary subtask for performing task T. By using the aggregation function shown in equation 1, the agent in a coalition which is the best fit for performing a certain subtask will be the one that performs it. The aggregated effort of agents in equation 1 is used to measure an score $scr(\sigma_x, T)$ that tells how well the agents in coalition σ_x perform together for accomplishing a task specification T. The score of a coalition is computed as the scalar product between σ_i and T:

$$scr(\sigma_i, T) = \sum_{l=0}^{k} (\sigma_i^l * T^l) \tag{2}$$

Agent Choices and Strategies Each player's strategic variables are his coalition choice to join σ_j and a set of agents in this coalition $\phi_{jk} \subset \sigma_j$ to eliminate. ϕ_{jk} can be empty. The possibility of optimisation responds to the change of value that certain agents can experiment when in their coalition they are out-skilled by a new member and so they become redundant. The new membership, together with the optimisation proposal is not accepted straightforward, it is evaluated by the members of the target coalition. Just those actions accepted by a majority (more than the half) of members in the affected coalition, are performed. An agent that is requested to take an action can submit a finite number of requests in an specific order, in such a way that if an action is not accepted, next action is checked and so on. If none of its action proposals is accepted, the agent stays in the same coalition where it was.

All the agents in the population follow the *Competitive Strategy* [14], that consists on proposing a set of actions that contain every proposal that either improves the score of the coalition the agent is in, or keeps the same score while reducing the size. When they receive a proposal from an outsider, they accept if they are not in ϕ_{jk}'s proposal, and if the proposal improves the score or keeps it while reduce the coalition size. Every agent j has a fixed social network α_j that is a non-empty set of agents. When considering to join a different coalition, agents are limited to just evaluating coalitions of agents in α_j.

RFP Iterated Model At time 0, every agent is a coalition of just one element ($\sigma_i = \{\sigma_{ii}\}$). A task T is issued and a run of negotiation starts in which every agent, sequentially and following a random order, is asked about an action to take. Agents decide on the basis of their strategy (see previous subsection), and the action taken is performed before the next agent is asked. Agents have no knowledge on the order in which they are asked, but by the time they are asked they have perfect information of the state of the coalitions where the agents in their social network are. The negotiation process is repeated as many times as necessary until no agent is willing to leave the state is in or none of its choices to leave its current state is accepted by the hosting coalition[1]. Is in this state of equilibrium when the system is observed, and the data of coalitions is captured for further analysis.

3 Experiments

Metrics used in the experiments: In order to measure how competitive an agent is in a simple way, every agent in the population has the same total number of skill capabilities, but distributed differently across skills. This way we can define a simple metric of *Competitiveness*: $com(\sigma_{ij})$, by just measuring the standard deviation in its skill values weighted by the task values. Analogously, *Versatility* is defined as the inverse of *Competitiveness*: $(ver(\sigma_{ij}) = 1/com(\sigma_{ij}))$.

In order to test the dynamism that a certain underlying social network permits, a new metric called *Historical Average Degree* (HAD from now on) is defined. This metric measures the distance that exists in the social network between agents in a coalition. What makes different this metric from the classical distance definition (minimal path length between two nodes), is that the distance is measured just using the links between the coalition members. The HAD value between two members does not change through the addition or abandon of partners. When an agent A joins a coalition there is at least 1 member (B)

[1] In [15] it was shown how the system always converge to an stable state when the population follows an score maximizing function.

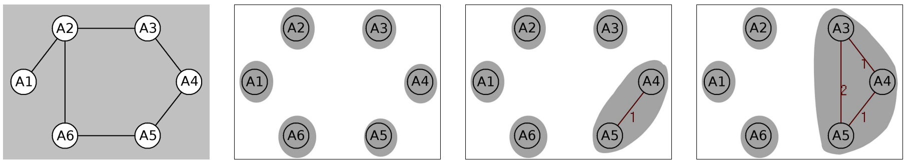

(a) Social network (b) Initial partition (c) A5 joins A4's coalition (d) A3 joins A4's coalition

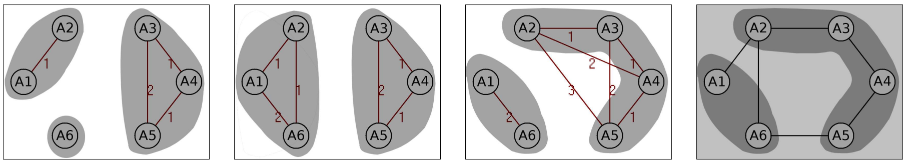

(e) A1 joins A2's coalition (f) A6 joins A2's coalition (g) A2 joins A3's coalition (h) social net with coalitions

Fig. 1. Example of iterated coalition formation process. Edge's values reflect the HAD between nodes. Shaded areas denote coalitions.

with path-length equal to 1, (otherwise this agent would not know directly any of the members of the coalition, hence it could not consider that coalition). The distance to the other members corresponds to 1 (distance from A to B) plus the minimal path length from B to the member of the coalition through the complete (HAD valued) network between members of the coalition. The final value of the metric at the end of the process is the mean of each HAD value between each member of the coalition. This mean is computed for every coalition when the system has reached equilibrium. Figure3 exemplifies how the metric is computed during the process of coalition formation.

At the end of an experiment, coalitions are compared with the optimal coalitional structure σ^+; that is the structure that creates the coalitions with the maximum possible score[2]. If top coalitions of an experiment have low score versus those in the optimal structure is because the best endowed agents are stuck in suboptimal coalitions. This way, a certain experiment ends up with better or worse results depending on how efficient the best agents are in getting to know each other. To measure this efficiency, the aggregated score of an experiment is obtained by computing the weighted sum of every coalition with an exponential weight: $Sc(\sigma) = \sum_{\mathrm{rank}=0}^{p} (\mathrm{scr}(\sigma_{\mathrm{rank}}, T) * 1/2^{\mathrm{rank}})$. To know how suboptimal a coalitional setup is the following difference is computed: $\mathrm{Sub} = \mathrm{Sc}(\sigma^+) - \mathrm{Sc}(\sigma)$

3.1 experimental set-up

In order to investigate the effects of networks properties in the dynamics of exploration of agents, a significant number of experiments have been performed

[2] This reference structure is computed by exhaustive exploration using farsighted agents

	node degree (k)	avg. distance (ad)	clust. coefficient (cc)
SF (Scale Free)	1.4	2.571	0.072
Rdm (Random Ërdos)	1.4	2.644	0.028
Rdm2 (Random Watts)	1.4	2.643	0.024
SW (Small World)	1.4	3.541	0.565

Table 1. Average properties across the set of 1000 networks used in each topology

varying the most significative variables. These are the concrete social network topology and the concrete mapping between agents and nodes in the network. A total of 4000 networks with 500 nodes and average undirected connectivity degree of $k = 10.3$ have been tested, in 4 different topologies: 1000 *Small World* networks (using the Watts-Strogatz model [19] with $p = 0.07$), 1000 *Random* networks (Using Ërdos-Reni model [8]), 1000 *Random2* networks (Using Watts-Strogatz model with $p = 0.07$ and rewiring randomly the nodes while keeping every node's degree) and 1000 *Scale Free* networks (using Barabassi model [3]).

For each one of the networks, 3 different mappings have been created in the following way: in one hand, for every network, nodes are ordered decreasingly by connectivity degree k. In the other hand, 2 different orders of agents are created by computing *com*, and *ver* metrics in every agent (see section 3) and ordering agents decreasingly by each metric's values. Every experiment maps every agent in one of the 2 specific orders (*ver-ord* or *com-ord*), to a node according to its k order in the network. An additional random mapping *rdm-ord* has been used to create the third tested order. That creates a total setup of 12000 experiments: for each 1000 networks of each of the 4 topologies tested, 3 different sets of experiments are performed, each one with a concrete mapping between node degree and agent characteristics. Experiments are run until the system has converged to an equilibrium (see section 2).

For space restriction reasons, not all the experiments performed are shown. Some of the variables (those that just produce scale effects but do not change the global effects) have been fixed. These variables are: the population composition (500 agents with 10 skills and, for every agent a total value of 100, heterogeneously distributed amongst the skills. The stdev. in the skill distribution is homogeneously distributed from 5 to 20), the task used in all the experiments ($T = \langle 1, 2, 3, 4, 5, 6, 7, 8, 9, 10 \rangle$ [3]) and the connectivity degree k (10.3 for each one of the 4000 networks).

3.2 Experiments Results

experiments with rdm-ord mapping: This setup does not prioritise any specific type of agent in the experiments this way the effect of the topology under

[3] the difference between values favors the diversity of Competitiveness and Versatility degrees (see section 3). None of the values is 0, so that all the skills are required (see section 2)

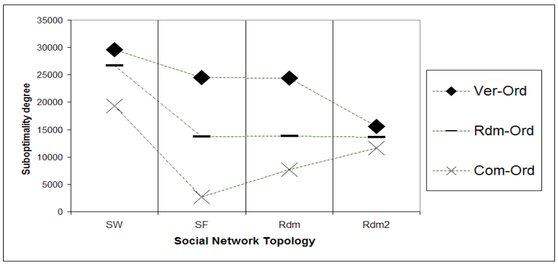

Fig. 2. Suboptimality for each topology and mapping. Topologies with higher node degree heterogeneity show higher sensibility to specific agent mappings. Small World proves to be the most suboptimal topology.

even conditions for every agent can be observed. The average suboptimality of each topology is summarised in Figure 2. There is a clear disadvantage in the SW networks. Those networks fail to provide the necessary conditions to put in contact the most competent agents. There are two main reasons that explain this fact. First, the highest average distance in the nodes of the SW networks (see Table 1), is an indicator on how short the path that interconnects two nodes can be. The longer this path is, the higher can be the required number of jumps in the network to group two agents. However this indicator is not the only fact that explains the suboptimality of SW networks, as this metric accounts for the average shortest distance, but agents do not necessarily follow the shortest path that interconnect them. Paths can be blocked when intermediate agents have an stable coalition that do not adhere new members, this way, in order to join specific agents, a network topology requires to provide not only *optimal paths* but also as much *alternative paths* as possible. The second reason that explains the inneficiency of SW networks has to do with the provision of alternative paths. The structured nature of SW networks make that, initially, agents form coalitions with other agents within their cluster; coalitions iteratively are expanded because structured clusters are interconnected. That rises HAD in coalitions to similar levels than the SF networks (see figure 3(c)) but the process of exploration is heavily dependant on agents with high *betweenness* (rewired agents in the Watts model), those that create shortcuts between clusters. If those agents are stuck in a certain coalition, they are blocking the path that interconnects distant clusters, hence just permitting agents to explore through neighbor clusters.

Another interesting fact captured in Figure 3(c) is that Rdm and Rdm2 networks show higher average HAD compared to SF. However, SF networks prove to have similar performance than Rdm and Rdm2 (see Figure 2). Again, the average distances partly explains this fact; as SF networks have, in average, shortest optimal distances, some agents can find the way to each other in a shorter number of steps than for the case of Rdm and Rdm2 networks. However there is a structural reason than makes SF networks create coalitions with less HAD. As in the case of SW networks, some nodes with a crucial network role might block the process. While for SW networks these were the shortcut agents, in SF, are the agents with higher degree (hubs). When these agents form an stable coalition, all the connections that were concentrated in them are lost and all the paths that crosses them are be blocked, so the exploration might be affected when the stabilisation is early reached. This way, the advantage of having a shorter avg. distance is compensated by the disadvantage of having nodes accumulating many connections.

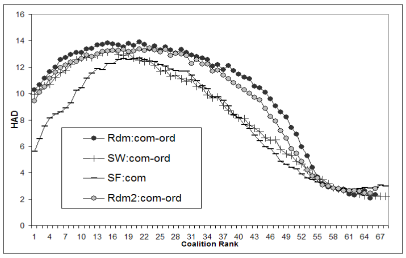

(a) HAD by ranking in com-ord

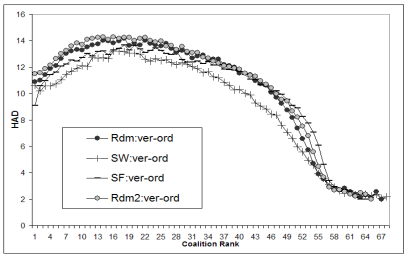

(b) HAD by ranking in ver-ord

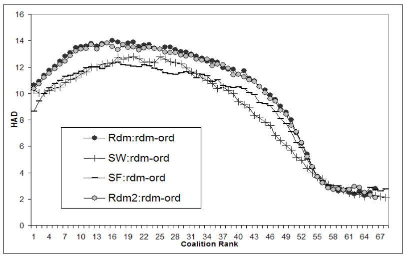

(c) HAD by ranking in rdm-ord

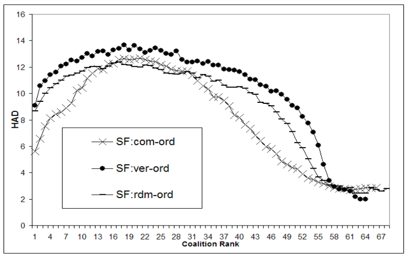

(d) HAD in Scale Free networks

Fig. 3. Avg. HAD of the 67 best coalitions through 1000 experiments showing the amount of exploration that each topology permits.

experiments with ver-ord and com-ord mapping: Network structures with high inequalities in the nodes degree are more affected by an specific mapping. SF topology is the one with higher unequal distribution, followed by Rdm. In SW and Rdm2, the degree distribution in nodes is almost even, and so they are not very sensitive to the different mappings.

An interesting result drawn from experiments shows that when the more competitive agents (those that are ment to be part of the leading coalitions) have the lower number of connections in the network, the general exploration performed by agents rises significantly (ver-ord fosters higher HAD than other mapping, see Figure 3(d)). This happens because competitive agents have higher potential of attraction. As these agents are valuable parts of many possible coalitions, others are attracted to join them. This attraction fosters the mobility of many agents that improve the score of the coalition of the competent agent by attracting each time more competent agents. However, in many cases, this incremental attraction proves to be not enough for letting the most competitive agents to get to know each other. This is why ver-ord experiments produce the highest suboptimality degree (Figure 2).

Inversely, when competitive agents are highly connected, good quality coalitions are formed in few jumps, having near optimal score. Their members are highly connected hubs in the social network that easily find a path to get in touch. However, in the lower positions the HAD analysis shows a deficient dynamism because there are not many attractive agents to foster movility between coalitions, hence the coalitions at the lower positions end up having poor performance being limited not to explore much further than their close neighbors. The lack of dynamism in those positions can be clearly detected in Figure 3(d) in the SF:com-ord series.

4 Conclusions and Future Work

In the rush to connectedness towards structured organisational topologies, the recent work by [13] suggested a potential downside of highly efficient topologies with short average path lengths. In the present work, more arguments are provided in favor of unstructured networks but using a completely different MAS model, based on a well known organisational paradigm called Request For Proposal. As in the case of [13], RFP model also proved that low levels of homogeneity in the population capabilities reduce the exploration at the social network. When highly compatible agents are closely connected in the social network, diversity in the system is squeezed out. And agents are unable to find out optimal solutions. Additionally, RFP model suggests that Small World networks are potentially inefficient because they concentrate in a reduced set of agents a set of key connections between different clusters. When those agents block the exploration, the system end up having worse global results.

In spite of the attention that dynamic coalition formation area is attracting in the last years, the use of social networks to map limited awareness of agents has not been deeply explored (for an exception see [10, 9] in the area of team formation). The present work represents an step further in the combination of social network analysis and coalition formation, proving that coalitional models can benefit from the research in the area solving issues related to farsightedness, and also highlighting some properties of well known network topologies that are not beneficial in any case.

The work in progress is focused on testing the effects of placing specific agents in nodes of the network with other network properties such as page-rank centrality, betweenness or clustering coefficient. The conclusions drawn from these work will be applied in the design of the more appropriate social network adaptation mechanism that agents should have in order to reduce the suboptimality degree of coalitions.

References

1. R. Axtell. The emergence of firms in a population of agents: Local increasing returns, unstable nash equilibria, and power law size distributions. Working Paper 3, Center on Social and Economic Dynamics, Brookings Institution, 1999.
2. R. Axtell. Effects of interaction topology and activation regime in several multi-agent systems. In *Proceedings of the Second International Workshop on Multi-Agent-Based Simulation (MABS00)*, pages 33–48, London, UK, 2001. Springer-Verlag.
3. A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
4. R. Burt. *Structural Holes: The Social Structure of Competition.* Belknap Press, 1995.
5. D. Culler, D. Estrin, and M. Srivastava. Guest editors' introduction: Overview of sensor networks. *Computer*, 37(8):41–49, 2004.
6. J. Delgado, J. M. Pujol, and R. Sanguesa. Emergence of coordination in scale-free networks. *Web Intelli. and Agent Sys.*, 1(2):131–138, 2003.
7. J. Diamond. *Guns, Germs, and Steel: The Fate of Human Societies.* W.W. Norton and Co. New York, 1999.
8. P. Erdos and A. Renyi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6:190–297, 1959.
9. M. E. Gaston and M. desJardins. Agent-organized networks for dynamic team formation. In *Proceedings of AAMAS '05*, pages 230–237, New York, NY, USA, 2005. ACM Press.
10. M. E. Gaston, J. Simmons, and M. desJardins. Adapting network structures for efficient team formation. In *Proceedings of AAMAS-04, Workshop on Learning and Evolution in Agent-based Systems*, 2004.
11. B. Horling and V. Lesser. A Survey of Multi-Agent Organizational Paradigms. *The Knowledge Engineering Review*, 19(4):281–316, 2005.
12. S. Kraus, O. Shehory, and G.Tasse. Coalition formation with uncertain heterogeneous information. In *Proceedings of AAMAS'03*, pages 1–8, 2003. Melbourne, Australia.
13. D. Lazer and A. Friedman. The hare and the tortoise: the network structure of exploration and exploitation. In *dg.o2005: Proceedings of the 2005 national conference on Digital government research*, pages 253–254. Digital Government Research Center, 2005.
14. C. Merida-Campos and S. Willmott. Agent compatibility and coalition formation: Investigating two interacting negotiation strategies.
15. C. Merida-Campos and S. Willmott. The effect of heterogeneity on coalition formation in iterated request for proposal scenarios. In *Proceedings of the EUMAS Workshop*, 2006. Lisbon.
16. I. Milchtaich and E. Winter. Stability and Segregation in Group Formation. *Games and Economic Behaviour*, 38:318–346, 2001.
17. J. M. Pujol, A. Flache, J. Delgado, and R. Sangüesa. How can social networks ever become complex? modelling the emergence of complex networks from local social exchanges. *Journal of Artificial Societies and Social Simulation*, 8(4):12, 2005.
18. H. P. Thadakamalla, U. N. Raghavan, S. Kumara, and A. Albert. Survivability of multiagent-based supply networks: A topological perspective. *IEEE Intelligent Systems*, 19(5):24–31, 2004.
19. D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, (393):440–442, 1998.