The Impact of Network Topology on Trade in Bartering Networks – Devising and Assessing Network Information Propagation Mechanisms

David Cabanillas[1] and Steven Willmott[1]

Technical University of Catalonia,
Software department,
Campus Nord, Omega building
Jordi Girona Salgado, 1-3
Barcelona (08034), Spain
{dconrado, steve}@lsi.upc.edu

Abstract. Resource allocation in distributed systems is an exciting area of research. Inherent properties in this environment, such as strategic users acting selfishly and the structure of the environment within which exchanges occur, are relevant challenges to study. This paper proposes a market–based resource allocation in a distributed environment. Our model has as it is objective the exploration of the effects of network structure in the allocation of performance together with mechanisms to improve the performance of the market. The proposed model, as well as mechanisms to maximize the allocation of objects/goods have been implemented and studied experimentally. The results obtained shows that topology affects the performance of the market. Using information propagation mechanisms clearly contributes to its improvement.

Key words: Market–based, Network Topology, Peer–to–Peer Computing

1 Introduction

Network exchanges and markets consist of a environment containing many interconnected agents interested in buying and selling objects. In many realistic situations agents are only connected to a limited number of other agents and unable to directly trade with all others present in the environment. Relevant challenges in the performance of the market have to do with:

- The fact that agents are goal–oriented – *quid pro quo*: Often trades cannot happen without mutual interest ([11], [2] and [3]).
- Agent needs are not detected: An efficient market requires that information about potential sellers/buyers is dispersed among participants [4].

Therefore, an efficient market relies upon two main interrelated factors: competition and information ([15], [9]). *Competition* arises due to the individual and

conflicting objectives amongst the members of the market. The available *information* that helps the decision–making process. These challenges can be seen in multiple markets. From P2P networks which share content amongst peers to people networks which share goods.

In this context, this paper is concerned with analyzing how network topologies affect the bartering processes. The analysis in particular focuses on systems with:

- Sparse topologies (with few links per node).
- Self–interested agents (which only trade for immediate gain).
- Random networks.
- A variety of simple information propagation mechanisms ([1], [10]).

1.1   Problem Model

The model of the problem discussed here, has the following characteristics:

- The market studied is composed of nodes which offer objects.
- Nodes are located in a network and are linked to a small quantity of other nodes.
- The links amongst nodes are static, but the nodes that form the network have periods of being switched on and off (to simulate variable up-times of nodes).
- Each object has a unique level of satisfaction associated with it for each agent in the system (Level of Satisfaction - *l.o.s.*).
- Initially the objects are randomly assigned to nodes.
- Trades are conducted by means of bartering.
- Trades modify the global *l.o.s.*.
- Members take only local decisions.
- Information about available objects is only available from local connections.
- Members can only trade with directly connected neighbors.
- Members only decide to trade when the trade is immediately beneficial.
- Global performance is measured as the sum of *l.o.s.* over all nodes.
- A steady state is achieved when no more trades are possible.

The model economy studied is captured is a Java program and with repeated simulations we explore the effects of the degree in the performance of the market and how the propagation of preferences and its extensions modifies the market performance. Table 1 shows an overview of the parameters that are used in the simulator. The simulations start with a *forced step* which ensures that a certain number of nodes offers goods initially (i.e. to ensure trading begins) – without this forced step in certain scenarios no trading would begin. [13]

1.2   The topology

The relationship amongst nodes in a market can take many different forms. The first aim is to demonstrate that the topology, specifically the quantity of links

Name of the parameter	Value
Quantity of nodes	500
Quantity of objects	3,000
Initial objects x node	6
Average initial *l.o.s.*	7,000

Table 1. Simulators parameters

in the network, influences the performance of the market. In our case, we will work with random topologies introduced by Erdös and Renyi ([6], [7]).

Nodes follow rules at the instant of offering objects to the market. These rules decide which available objects the node should offer to its neighbors. They are influenced by the objects offered in the market. If the market, concretely the nodes that are direct neighbors, is offering valuable objects, the node is willing to offer its objects. When a node is connected to the rest of the nodes in the system (i.e. fully connected topology) it has more opportunities than a node that has a reduced number of connections. In this second case, the quantity of neighbors will be relatively lower in relation to the first case. Thus, the fewer neighbors that a node has, the greater the reduction in the possibility of useful objects being offered.

1.3 Metrics Related to the Topology

The goal of the experiments is to investigate the importance of the variations in the topology with respect to the quantity of links. Looking for this goal we vary the topology of the network with respect to degree. Table 2 gives an overview of the measures related to the scenarios studied. This range of scenarios allows us to recognize when the market is affected by the lack of links. The scenarios share the same set up. The only altered parameter has been the quantity of links. In these scenarios the quantity of links decreases from scenario 1 to scenario 6. From scenario 1, that has a fully connected structured with 124,759, to the scenario 6 that only has 779 links. For each scenario, a set of different network topologies have been tested in order to verify that the results are not dependent on a concrete wired network.

To determine the quantity of links accurately, W, required by the market, the diameter, the average path length and the average degree in the network are measured by each scenario. Also, in all of the scenarios, the number of unreachable pairs is equal to zero to ensure a graph is not disconnected. The *diameter* D of a network is defined as the maximum distance amongst all distances between any pair of nodes in the network (i.e. the longest shortest path between any pairs of nodes). The *average path length* L of a network is defined as the mean distance between two nodes, averaged over all pairs of nodes (i.e. average distance amongst reachable pairs of nodes). Finally, R is the *average degree* of the network.

In order to contrast the effect that the quantity of links has on the performance of the market, a set of scenarios where the quantity of links has been varied

has been simulated: from a fully connected topology to a quasi non-connected topology [14]. Table 2 shows the different scenarios and figure 1 shows the different scenarios using the force–directed graph layout proposed in [8].

Scenario	W	L	D	R
s1	124,759	1	1	500
s2	5,457	2.72	4	16,07
s3	3,897	2.90	4	11,47
s4	2,338	3.02	5	6,88
s5	1,559	3.57	6	4,59
s6	779	5.34	11	2,29

Table 2. Network measures

2   Simple Bartering in Random Networks

Figure 2 shows the progression of the *l.o.s.* in a simple bartering environment. In this case, from time 0 to 15, the *l.o.s.* increases and then stops growing. This is the period of time where the nodes have randomly assigned objects that probably are of little interest. This could be the same situation for the rest of the members. Many nodes with little desired objects is a fruitful interchange scenario. But when nodes get good object, it becomes harder to offer more advantageous objects (since what is on offer must be better than what the nodes already have).

In scenarios 1, 2 and 3 the *l.o.s.* gets a value near to 7,800 points. In the rest of the scenarios, the value obtained is far from the optimal *l.o.s.*. But in scenarios where L is greater than two, some simulations have shown that no trades are made in the market.

## 3   Effects of Information Propagation on Trade Diffusion

Up until now, each node has used knowledge of their neighbors to evaluate possible exchanges of object with its neighbors. However, while some trade can occur in which objects travel multiple hops, it is clear that sparse topologies significantly restrict market activity. In order to explore what happens when more information is available, applying propagation of preferences, every node will now have a propagation list. This list contains the objects that neighbors desire and the information will be used to make trades. The algorithm used by the propagation of preferences can be summarized as follow:

    The propagation takes the following steps:

1. Select the neighbors of a node.
   (a) Send their preferences to these neighbors
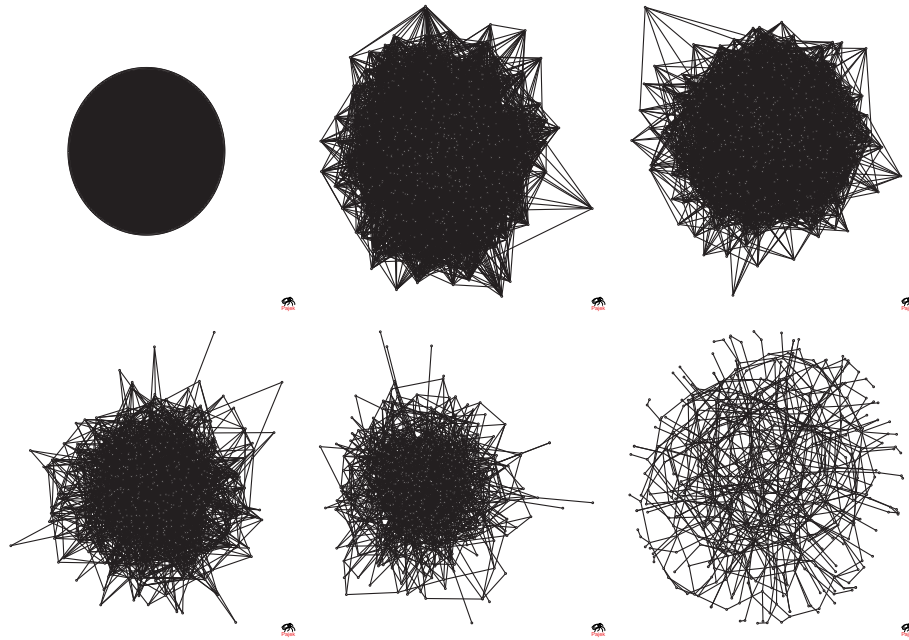
Fig. 1. Structure of the trade network from scenario 1 (i.e. fully connected) to scenario 6 (i.e. quasi non-connected).
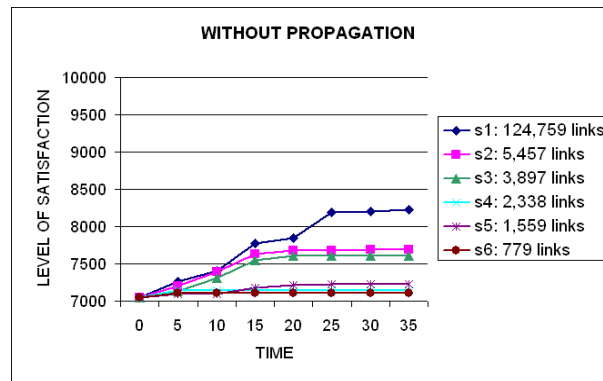
Fig. 2. Results working without propagation. In the graphs, x–axis represents the past of time and y–axis the *l.o.s.*.

```
    (b) Arrange the list of preferences (i.e. in top positions the most
        desired objects)
2. Go to 1 until all nodes have been visited in the market
```

Results of these simulations are shown in figure 3. In the initial time (i.e. from time 0 to 15) behavior is similar to that of the previous case. In scenarios 1, 2 and 3, the *l.o.s.* is above 9,000 points. In scenario 4 the *l.o.s.* is near to 8,500 points. When the network has 1,559 links (i.e. scenario 5) the *l.o.s.* is around 7,500 points and in the scenario 6 is where the value obtained is farthest from the optimal *l.o.s.*. In this case, the worst results appear when L is greater than 3. This shows that propagation begins to extend the range of trade, yet only in a limited way.
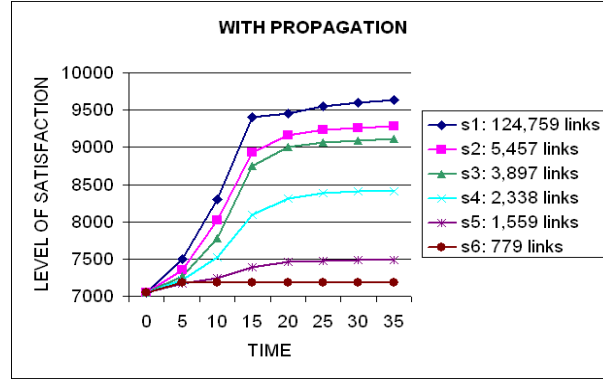
Fig. 3. Results working with propagation. In the graphs, the x–axis represents the passing of time and y–axis the *l.o.s.*.

It can be seen from the simulations that using propagation of preferences, the whole *l.o.s.* in the market increases. The results show that in scenarios with a small number of links the *l.o.s.* is substantially reduced.

In order to improve the current results, two extensions have been implemented and the results obtained from the experiments have been compared with the original propagation. Two different extensions have been studied in relation to the propagation of preferences:

– Extension 1: Avoiding the re–sending of preferences. In order to avoid a neighbor re–sending duplicate preferences to the original owner, this extension re–uses that otherwise wasted space.
– Extension 2: To promote the propagation of preferences in nodes with few links and to promote the propagation of neighbors in nodes with many links. The idea behind this extension is to make the nodes with few links put more wishes in the list of propagation. It gives more emphasis on the desires of this node. With respect to nodes with many links, the extension gives more emphasis to the preferences of its neighbors.

Figure 4 shows the *l.o.s.* when the market is using propagation of preferences based on extension 1 and extension 2. The experiments shown here, focus on scenarios 4, 5 and 6 as theses are the scenarios where there are significant variations in the *l.o.s.*. The first column of the set is with the market using the original propagation, the second column is the *l.o.s.* obtained with extension 1 and finally the last column is related to the extension 2. The results reveal that with 779 links neither extension 1 nor extension 2 work properly. The low quantity of links makes this market a sterile market, as much for the original approach as well as for the proposed extensions. In scenario 5 with 1,559 links, the results using extension 2 are better than in the original and extension 1 approaches. In scenario 4, both extensions improve the results of the original approach.

4   Discussion and Related Work

Whilst simple, the results show two things. The first is that the quantity of links clearly has an impact on the performance of the market. In fully connected networks, all nodes are free to trade with any node in the market. Optimal allocations are possible amongst nodes with few trades. The drop in one–to–one trades is less than for instances where a chain of one–to–one is necessary trades - as occurs in very spares networks. The second is that market performance is affected by the use of propagation of preferences. Even simple propagation can significantly change the efficiency of the market.

The information propagated by the nodes with respect to their preferences is used by other nodes which effectively act as mediators. A mediator is a node that trades objects useful to others with the hope that these others accede to make trades with the mediator. By means of these mediators the expectation that the total quantity of trades and the *l.o.s.* increase are true, as is revealed using the mechanism of propagation of preferences proposed in this paper.

Comparing the results working with and without the propagation mechanisms, we can establish that the asset using propagation is relevant. To use a propagation list has a computational cost and can also be a handicap that in topologies with a high density of links since propagation and analysis of preference lists costs more than the relative gain. The model applied limited the length of the preference list propagated, hence processing was still relatively rapid - however comparison studies are needed to identify the tradeoffs between propagation mechanisms and different topologies.

In summary, topology has a direct effect on the performance of the market. The quality and quantity of information from one individual to another, in our case propagation of preferences, has a positive influence on the performance of the market.

Information diffusion is an important area with strong similarities to this work ([12], [1], [5]). The innovation/knowledge spreads throughout a social network as objects allocated in our market. But the non–rivalry property is commonly assumed in the information diffusion process. During the spread of knowledge, the source of the information does not loose its information. In contrast, in
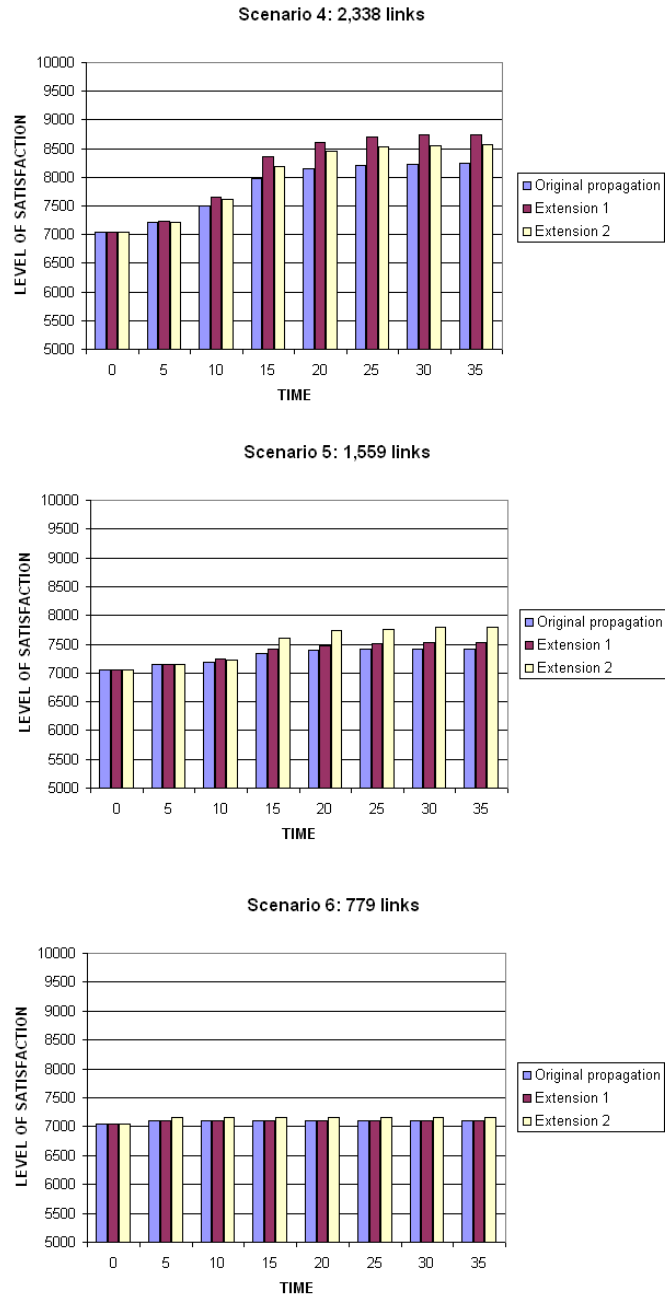
Fig. 4. *L.o.s.* with the passing of time working with original propagation, extension 1 and extension 2.

our approach, the object only belongs to unique members in the social network at the same time. For this reason, when an object changes hands, the owner of the object loses the *l.o.s.* associated with the object and this does not happen in the information diffusion approach.

5   Future Work

The mechanisms studied here provide a baseline for further work and represent simple examples of propagation techniques and lay down a starting point for more complex mechanisms. Future work focuses on extending the two different issues studied in this paper – topology and propagation of preferences:

- With respect to topology, the next steps are to work with other topologies, such as scale–free and small–world topologies which are known to support different types of propagation properties. In these cases, the effect of the bartering need for "double coincidences of wants" will likely have different types of effects on the market dynamics.
- In terms of extending the propagation of preferences, work now focuses on propagation which has longer and more directed reach across multiple nodes in a network. An important first step that is the propagation of ownership (who owns which objects). This provides a counterpart to demand information which allows agents to direct their demand propagation mechanisms in the network.

References

1. Floortje Alkemade and Carolina Castaldi. Strategies for the diffusion of innovations on social networks. *Comput. Econ.*, 25(1-2):3–23, 2005.
2. David Cabanillas and Steven Willmott. Self-organization amongst non–altruistic agents for distribution of goods: Comparing bartering and currency based exchange. In Barbara Dunin-Kęplicz, Andrea Omicini, and Julian Padget, editors, *4th European Workshop on Multi-Agent Systems (EUMAS 2006)*, Lisbon, Portugal, December 2006. CEUR.
3. David Cabanillas and Steven Willmott. Studying viable free markets in peer–to–peer file exchange applications without altruistic agents. In *Fifth International Workshop on Agents and Peer-to-Peer Computing (AP2PC)*, 2006.
4. Thomas E Copeland and Daniel Friedman. The market value of information: Some experimental results. *Journal of Business*, 65(2):241–66, April 1992. available at http://ideas.repec.org/a/ucp/jnlbus/v65y1992i2p241-66.html.
5. Robin Cowan and Nicolas Jonard. Network structure and the diffusion of knowledge. *Journal of Economic Dynamics and Control*, 28(8):1557–1575, June 2004.
6. P. Erdos and A. Renyi. On random graphs i. *Publ. Math. Debrecen*, 6:290, 1959.
7. P. Erdos and A. Renyi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 7:17, 1960.
8. Thomas M. J. Fruchterman and Edward M. Reingold. Graph drawing by force-directed placement. *Software - Practice and Experience*, 21(11):1129–1164, 1991.

9. Rachel E. Kranton and Deborah F. Minehart. A theory of buyer-seller networks. *American Economic Review*, 91(3):485–508, June 2001. available at http://ideas.repec.org/a/aea/aecrev/v91y2001i3p485-508.html.

10. Abhishek Kumar, Jun Xu, and Ellen W. Zegura. Efficient and scalable query routing for unstructured peer-to-peer networks. In *Proc. of IEEE Infocom*, 2005. to appear.

11. K. Ranganathan, M. Ripeanu, A. Sarin, and I. Foster. To share or not to share' an analysis of incentives to contribute in file sharing environments, 2003.

12. Everett M. Rogers. *Diffusion of Innovations*. Fifth edition, 2003.

13. Mario T. Schlosser, Tyson E. Condie, and Sepandar D. Kamvar. Simulating a p2p file-sharing network, 2002.

14. C.-L. Tang, W.-X. Wang, X. Wu, and B.-H. Wang. Effects of average degree on cooperation in networked evolutionary game. *European Physical Journal B*, 53:411–415, October 2006.

15. Allen Wilhite. Bilateral trade and 'small–world' networks. *Comput. Econ.*, 18(1):49–64, 2001.