

**THE ASYMMETRIC DIFFUSION OF TRUST BETWEEN COMMUNITIES:
SIMULATIONS IN DYNAMIC SOCIAL NETWORKS**

Luca Allodi

University of Trento
Via Sommarive, 14
Povo (TN), I-38123, ITALY

Luca Chiodi
Marco Cremonini

University of Milan
Via Bramante 65
Crema (CR), I-26013, ITALY

ABSTRACT

In this work, we present a model of social network showing non-trivial effects on the dynamics of trust and communication. Our model's results meet the characteristics of a typical social network, such as the limited node degree, assortativeness, clustering and communities formation. Simulations have been run first to present some of the most fundamental relations among the main model's attributes. Next, we focused on the emerging asymmetry with which trust develops within different communities in a network. In particular, we considered categories of nodes differing for their communication profiles and a specific example of bridge between two communities. The results are discussed to provide insights about the dynamic formation of communities based on trust relations. These results are the basis for future works with the aim of better understanding the dynamics of the diffusion of trust and its influence on a growing social network.

1 INTRODUCTION

Simulating the dynamics of social networks is a research topic that is increasingly addressed for different reasons. The most recent is the success of some well-known systems, Facebook to mention the most cited, but other applications of the social network model have proved to be extremely interesting. In the literature, there exist many works that have applied models of social networks to real case studies. The exchange of the emails in a community of people, for instance, represents a relevant case study (Tyler, Wilkinson, and Huberman 2003), as well as the dynamics showed by individuals joining and leaving groups of interests, which may stem from leisure (e.g. the case of online games) to scientific research or corporate projects (Newman 2004). In many cases, the research trail is that of increasing the complexity of models and of simulations to address more elaborate research goals and analyze network's properties. The aim of our work is twofold:

- To present our model of dynamic social network based on knowledge exchange between nodes and the simulation results of the stochastic behaviors and emergent properties;
- The analysis of a case study based on the effects of trust and of communication efficiency on the formation of communities and on the flow of knowledge.

More precisely, our social network model includes the selection of partners and the action choice, mimicking the diffusion of knowledge between actors, in the form of questions and answers. Each actor knows a variable number of *topics*, each one characterized by a degree of knowledge and a degree of interest in knowing more. As a result, some not trivial, recurrent behaviors have been observed and analyzed.

The background for our work comes from the large, interdisciplinary literature related to dynamic social networks, which exhibit peculiar characteristics with respect to non-social networks, the most notable of which are related to degree correlations of adjacent nodes and to clustering (Jin, Girvan, and Newman 2001,

Newman and J. 2003, Skyrms and Pemantle 2000). Social networks are typically *assortative*, meaning that the degree correlations of adjacent nodes is positive, i.e. nodes of high degree tend, on average, to be connected with other nodes of high degree. This observation has fostered some relevant studies about the special structure and behavior of social networks, which distinguish them from other non-social networks. The second peculiar characteristic, *clustering*, has been extensively studied and contribute to explain the assortativity of social networks. Clustering has been defined in term of network transitivity, that is, given an edge between a pair of nodes A and B and another edge between nodes A and C , a network is said to be high transitive if it is likely that there will also be a connection between nodes B and C (Newman and J. 2003). For social networks, it has been observed how the clustering coefficient is typically greater, possibly orders of magnitude greater, than in the corresponding random graph (Watts and Strogatz 1998, Newman 2003, Newman 2001). The clustering effect of social networks has been observed since long by sociologists, which have called it “triadic closure bias”, i.e., the tendency that individuals have of meeting a friend of a friend rather than maintaining relations with reciprocally disconnected friends or meeting strangers (Hanaki, Peterhansl, Dodds, and Watts 2007).

The clustering effect is key to the study of social networks and consequently for the work we are presenting in this paper, because it implies that the network’s dynamic behavior is nonrandom and that communication of a node with new nodes is mainly driven by information made available by already known “friends” or “acquaintances”.

2 RELATED WORK

Recent research has dealt with models of social networks more complex than the ones studied in the past, which were mostly concerned with studying the fundamental mechanisms of social network growth. More recent studies have, instead, focused on advanced features like trust, recommendation, cooperation and information exchange, as we did in this work.

Walter *et al.* presented the work most closely related to ours, although different in the research goal (Walter, Battiston, and Schweitzer 2008). They considered a model of trust-based recommendation system on a social network, which assumed the transitivity of trust along a chain of relationships connecting agents, while the trust dynamics is restricted to neighbors only. Similarly, in our work trust is local because a node maintains trust relationships with friends only but, differently to them, we admit only a limited degree of trust transitivity (which is restricted to the best friend-of-friends) and did not model trust chains. However, in many other aspects the two models have similar features, such as the correlation between trust and node similarity, node’s preferences heterogeneity, and the dependence on knowledge sparseness. On one side, differently from our work, their modeling of the trust function is more complete, including a fast negative dynamics that we have not yet included. On the other side, while their model is of a static network with no rewiring of edges, ours is a model of an evolving social network, providing node choice rules reflecting the typical assortativity of social networks and a rich behavior dynamics that we simulated.

Brzozowski and Romero studied different features for recommending people in a directed social network (Brzozowski and Romero 2010). Although not directly related with our work, relevant for us is their analysis of structural closures in directed social networks, along with observations regarding forbidden triad and the relative relevance, under certain conditions, of similarity.

Hanaki *et al.* provided relevant observations with regard to the emergent cooperative behavior in social networks (Hanaki, Peterhansl, Dodds, and Watts 2007). In particular, they examined the correlation between the strength of ties and the level of cooperation of the network. With respect to our work, the interaction dynamics driving the evolution of the network is based on different assumptions and rules for the node selection. We have verified their findings in our context and found a confirmation under certain configurations of the simulated social network. Trust in their work has been modeled as a weighted sum of past trust’s average and depends on a factor measuring past experiences. Differently, we have not introduced a weighted average to model trust but, similarly to their work, in our model nodes trust others based on a self-declared expertise of the counterparts and on the frequency of past interactions.

Important for the analysis of mixing patterns and community structures in networks is the work by Newman and Girvan (Newman and Girvan 2003). This research analyzed most of the characteristics that our model of social network presents and that we have tested and discussed in this work, from the assortative mixing to the formation of communities, from the relevance of friend-of-friend relationships to the dynamics of the growing network.

3 MODEL DESCRIPTION

We consider a set of N nodes, n_1, n_2, \dots, n_N each one characterized by a *Personal state* PS_{n_i} and a *Friend state* FS_{n_i} .

Personal State. The Personal state PS represents which topics a node knows, how well it knows them, and how much interested it is in each one of them. In the model these information are described as tuples having the form $(topic, quality, interest)$. We consider a set of topics T representing all distinct topics that the population of nodes knows; each node n_i knows a variable subset of them, $T_i \subseteq T$. Therefore each node n_i has a Personal state having the form $PS_{n_i} = (\bigcup_{j \in T_i} (topic_j, quality_{i,j}, interest_{i,j}))$.

Friend State. The Friend state FS represents the connections a node has with other nodes (i.e., “friends”). A connection is a reciprocal relation between two nodes and is established when a question is asked by one and a valid answer is provided by the other (details about the interaction mechanics will be described in the following). When the first interaction between two nodes occurs, both nodes update their Friend states by adding a new pair composed by the friend’s identifier n_i and a counter *answers* keeping track of the answers received by another node. The reason for this choice is that friendship is considered reciprocal in our model of social network, therefore both nodes establish a reciprocal connection. On the contrary, trust, which in our model is a function of the number of answers received, is directed, therefore the parameter *answers* is increased only by the receiver node when an interaction takes place. More formally, each node n_i has $N_i \subset N$ friends, and a Friend state having the form $FS_{n_i} = (\bigcup_{j \in N_i} (n_j, answers_{i,j}))$.

3.1 Trust

In our model, the meaning of trust is “*an expectation of technically competent role performance*”, as for Barber’s subtypes of trust (Barber 1983, Thomborson 2010). Although relatively general, this definition reflects the dynamics of nodes in our social network. A node interacts with another based on the expectation of increasing its knowledge by establishing a relationship with a more competent counterpart. Trust tends to be local in our model, because a node interacts with unknown ones only when neither a friend nor a best friend-of-friends nodes can answer to its question.

More specifically, in our model, the notion of trust is key to the behavior of a node in two different actions: the choice of the peer node to interact with and the knowledge it gains from the selected peer. In other words, a node trusts another one when it chooses it and when subsequently it learns from it. Operationally, the two attributes that control these actions are the *quality* associated to the topic for which an interaction takes place and the *answers* recording the past history of interactions between the two nodes. In particular:

- Attribute *quality* is used for node selection among friends, best friend-of-friends or randomly chosen nodes, while attribute *answers* for selecting the best friend-of-friends.
- The difference between values of attribute *quality* owned by the respondent and the requester node represents the nominal amount of learning of the requestor node.

In general, trust should be time dependent; typically, a repetition of interaction within a time period reinforces trust, while the absence of interaction reduces trust (Burt 2001). Such dependency has been modeled in our work. The quality gain is discounted, that is recent friendships (those with a small number of *answers*) would result in smaller gains (with respect to the difference of quality between the respondent node and the requester) with respect to older friendships.

Therefore, trust has a dynamic based on the history of interactions. For this reason, we operationally measure trust by the number of interactions (i.e. attribute *answers*) between nodes.

As for Burt (Burt 2001), this assumption could be seen as the “baseline hypothesis” for trust setting the benchmark for future improvements and, by analyzing the dynamics of interactions, some interesting, qualitative insights could be derived about the dynamics of trust in the evolution of a relatively complex social network.

3.2 Node Setup

The Personal and the Friend state of nodes are initialized as follows:

Topics. A random set T_i of topics is defined for each node. The maximum number of topics assigned to the nodes can be limited by setting the maximum rate $\lambda_T \in (0, 1]$, so that $|T_i| \leq \lambda_T \cdot |T|$.

Quality and Interest. The quality associated to each topic of a node’s Personal state is set to a random value in $[1, 100]$. For the interest, the initial value is equally distributed among all topics, and is calculated as $100/|T_i|$.

Topic 0. A dummy topic called $topic_0$ is always present and, when chosen, a topic that does not belong to the node’s Personal state is selected. Forcing its presence in the Personal state means that each node has always a chance of requesting an unknown topic. The *quality* associated to $topic_0$ is always zero, while the *interest* is calculated as for the other standard topics during the network evolution.

Friends. All nodes have no friends at setup, making the evolution of the network fully stochastic. At start up, the selection mechanism is the random choice, then the preferential selection mechanism increasingly rely on local connections. As for the topics, a maximum number of friends per node can be configured by setting a maximum rate $\lambda_N \in (0, 1]$, so that $|N_i| \leq \lambda_N \cdot |N|$.

3.3 Node Choice

In our model, interactions between nodes represent a flow of knowledge, with nodes knowing less about a topic that ask to those knowing more and with the answers that increase the topic’s knowledge of the requesters. The node’s choice is preferentially bound to nodes already belonging to the requester’s Friend state and to “*best friends-of-friends*”. Only when these two options fails, then a random choice is executed.

“Best friend-of-friend” node. A “best friend-of-friend” node is a node belonging to a friend’s Friend state, owning the selected topic in its Personal state (i.e., “a friend of a friend who knows about the topic that is requested”) and having the higher value of the *answer* attribute (“the most reliable friend of a friend”). It is worth noting that the inclusion of “best friends-of-friends” among the nodes that could be chosen fosters network transitivity and the triadic closure bias.

More specifically, a node chooses another node as follows:

1. For a given node n_i , a topic ($topic_{j^*} \in T_i$) is selected among those in the node’s personal state PS_{n_i} . The selection is made by chance, with the chances being determined by the relative weights, i.e., the value of the interest associated to the topic ($interest_{i,j^*}$).
2. The node $n_{i'}$ to interact with is chosen based on the quality associated to the selected topic. Node $n_{i'}$ is the one owning $topic_{j^*}$ and having the maximum $quality_{i',j^*}$ among node n_i ’s friends and “best friends-of-friends”.
3. Node $n_{i'}$ must know more than node n_i about $topic_{j^*}$, that is $quality_{i',j^*} > quality_{i,j^*}$.
4. If no node satisfies the previous conditions, the selection of a counterpart for node n_i is made randomly over the entire population N (i.e., this is the case at start up).

3.4 Personal State Update

The update of the Personal state after a successful interaction is the key mechanism of our model. The Personal state is updated by the requestor only, because an assumption of our model is that the behavioral dynamics of nodes is driven by the new knowledge that a requestor node gains when an interaction is completed. The idea is that the new knowledge depends on the difference between the *quality* associated to the *topic* known by the responding node and the quality associated to the same topic known by the requestor. That difference increases both the requestor's topic quality ("the node learns something about that topic") and its interest in that topic ("the node wants to learn more about it"). In this way, the interaction between nodes based on a topic that both know enhances the similarity of connected nodes, which tends to form clusters as a consequence of the triadic closure bias.

3.4.1 Quality Increase

The quality increase does not depend only on the respondent's topic quality, but also on the number of answers a node has already received from the counterpart, which is where the notion of *trust* is more relevant. Intuitively, a node distrusts another one when they interact for the first time and this distrust progressively diminishes as interactions occur, with a negative exponential dynamics. This means that the knowledge gain of a node from the interaction with another one is discounted by a term that starts at a given value (i.e. ρ) and goes to zero exponentially.

More formally, we call $n_{i'}$ the requestor node, $n_{i''}$ the respondent node, and $topic_{j^*}$ the topic for which the interaction take place. The *quality gain* obtained by $n_{i'}$ is calculated as

$$\delta quality_{j^*} = \frac{quality_{i'',j^*} - quality_{i',j^*}}{\gamma + \rho e^{-\frac{x}{\theta}}} \quad (1)$$

with:

- $\gamma \geq 1$: the nominal fraction of $\delta quality$ that $n_{i'}$ could learn from another node for $x \rightarrow \infty$, i.e. $\lim_{x \rightarrow \infty} \rho e^{-\frac{x}{\theta}} = 0$;
- x : the value of the attribute *answers* of node $n_{i'}$ with respect to node $n_{i''}$;
- ρ : the initial discount applied to learning for $x = 0$. Intuitively, it represents the *distrust* a node has towards another one before any interaction occurs;
- θ : the parameter controlling at which rate node $n_{i'}$ increases its trust towards node $n_{i''}$.

3.4.2 Interest Increase and Reduction

The dynamics of a topic's interest could be either positive (*interest gain*) or negative (*interest reduction*).

Interest Gain. The interest associated to the topic that a node requested increases when the interaction successfully completes, otherwise it remains unmodified. The new interest depends on the $\delta quality$ value calculated in (1). The function used to calculate a new value of the interest has the exponential form:

$$\delta interest_{i',j^*} = \alpha \left(1 - e^{-\frac{\delta quality_{i',j^*}}{\beta}} \right) \quad (2)$$

with $\alpha > 1$ and $\beta > 1$ the two parameters we use to control, respectively, the scale and the slope of the function. This function has $\frac{\partial interest}{\partial quality} > 0$ and $\frac{\partial^2 interest}{\partial quality^2} < 0$, meaning that the interest is increasing, but with diminishing marginal gains. This way nodes tend to exhibit preferences among the topics of their Personal state, but the polarization is mitigated and the emergent behavior is not trivial. During simulations we have consistently seen heterogeneous behaviors among nodes, with some strictly polarized, while others not showing strong preferences.

Parameter β is important, because changing the slope of the interest growth modifies the tendency of interests to polarize. Nodes strongly polarized on few topics (i.e., high interest values) tend to acquire

knowledge more quickly and so be more likely to act as respondents in following iterations. In general, the result is a tendency to form isolated communities connected with strong ties. With interests less polarized, instead, nodes tend to choose a wider range of topics, then forming larger communities connected with weaker ties.

Interest Reduction. As a result of a successful interaction and the increase of the chosen topic's interest, a proportional *interest reduction* is applied to all the other topic's interests belonging to the node's Personal state. The reason for decreasing interests associated to topics known by a node except the one chosen, intuitively, is that it seemed to us not reasonable to model actors of a social network as having an always increasing total amount of interest. On the contrary, interests in different topics should balance each other, within certain limits.

For this reason, when an interaction completes successfully, all topics known by node $n_{i'}$, except $topic_{j^*}$, have their corresponding interest decreased by $\delta interest_{i',j \neq j^*}(t_k, t_{k-1}) = \delta interest_{i',j^*}(t_k, t_{k-1}) / (|T_{i'}| - 1)$, that is the value of the interest gain for $topic_{j^*}$ at t_k divided by the number of topics $|T_{i'}|$ minus one. As a side note, the interest reduction applies to $topic_0$ as well, which is included in the total number $|T_{i'}|$ of topics known by node $n_{i'}$.

3.5 Metrics

For the analysis of the network's dynamics that are presented in this work, we consider some relevant metrics: the *Clustering Coefficient* and the *Communication Efficiency*.

Clustering Coefficient. For the Clustering Coefficient we adopted the definition introduced by Watts and Strogatz: the Clustering Coefficient in social networks measures the cliquishness of a typical friend circle (Watts and Strogatz 1998).

Communication Efficiency. With Communication Efficiency we want to measure how often nodes are able to successfully interact, i.e. to receive an answer to a request, with respect to the number of requests they made during a simulation. The number of requests made by the population of nodes equals the number of ticks of the simulation, shown as Γ , because for each tick a node is selected and a request is produced. Formally, we define:

$$Communication\ Efficiency = \frac{Total\ No.\ of\ Answers}{Total\ No.\ of\ Requests} = \frac{\sum_{i=1}^N \sum_{j \in N_i} answers_{i,j}}{\Gamma}$$

4 NETWORK SIMULATIONS

All simulations have been run with the same basic configuration:

- Number of nodes $N = 150$;
- No limitation to the number of friends (out-degree);
- Duration of the simulation/No. of ticks $\Gamma = 150000$.

Number of friends. The first result regards the number of friends, which is a critical attribute because it represents the average outdegree of a network. It influences the topology and, as already mentioned, in social networks it should be limited. In our simulations we have consistently observed that the limited outdegree is an emergent behavior and does not need to be set by design. The typical average rate of friends was between 2% and 10% of the population.

Number of topics per node. The number of topics per node is a key factor for determining the emergent behavior of the network. We have analyzed: (i) the dependence of the average number of answers per edge, thus, qualitatively, the trust of a requesting node towards a respondent node; (ii) the dependence of the Communication Efficiency. We have observed a positive correlation in both cases: the more topics a node knows the more answers it receives and the better is the Communication Efficiency of the network. In short, heterogeneity of interests improves the flow of knowledge and the growth of the social network.

This result is intuitive, because nodes with higher number of topics tend to be less polarized, then to ask questions on more topics and to interact with more nodes. This permits them to enlarge their friend circle and communicate more. Figure 1(left) shows these results.




Figure 1: Results by varying the number of topics per node (*left*) and the number of topics in the network for different values of the parameter β (*center* and *right*).

Slope of the interest function. In this case we have tested the variation of the average number of answers per node and the Communication Efficiency, setting different values of the parameter β defined in equation (2). We recall that β modifies the slope of the exponential function that determines the new value of *interest* after a successful interaction. Increasing the value of β has the effect of decreasing the slope, hence the value of *interest* increases more slowly as well as the tendency of nodes to polarize their interests. The results of Figure 1 (*center* and *right*) are not trivial, in particular for what regards the average number of answers per node. The value *Number of topics*=50 is an exceptional case, because due to the scarcity of communication, varying β makes practically no difference. For higher values of the number of topics, the effect of β on the communication is visible and the correlation emerges. It is, however, also evident how for very high values (e.g. for $\beta = 500$ and $\beta = 1000$), meaning very small increases of the interest value, there is not a clear difference and the stochastic behavior is dominant. The same does not hold for the Communication Efficiency that is clearly dependent from β .




Figure 2: Model's efficiency with $\rho = 10$ (*left*), $\rho = 50$ (*center*), $\rho = 100$ (*right*).

Learning rate. Communication Efficiency have shown to significantly vary for different configurations of the model. In particular, the rate at which nodes *learn* seems to strongly affect this metric: learning speed is controlled directly by two parameters, ρ and θ , and indirectly by $\delta quality$, the difference between the topic's quality of the requester node and that of the respondent. We recall that ρ and θ control how fast a respondent node becomes trustworthy to a requesting node, with respect to the number of answers received by the latter. The higher the ρ , the lower the initial trust between two nodes will be; the higher the θ , the slower a node will trust another one.

To better understand these dynamics, we tested a standard configuration (150 nodes, $\gamma = 10$, $\beta = 50$) varying ρ and θ between simulations, and testing each setup through different numbers of topics in the network. These results are shown in Figure 2. As expected, higher values of ρ and θ result in higher

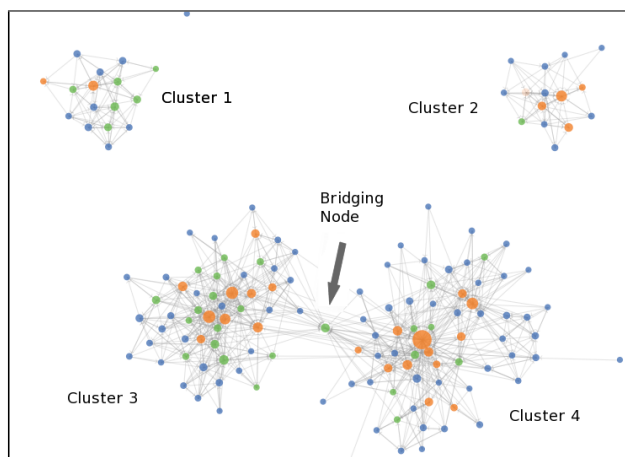


Figure 3: Case Study.

Communication Efficiency, because, intuitively, slowing the rate at which nodes trust each other, slows the rate at which they learn from their interactions, so that they need more answers before the quality associated to a topic becomes close to that of the best respondents.

We can observe that increasing the value of θ , initially, it sensibly improves the Communication Efficiency, while at higher values its increments are not as much significant. To clarify, we can consider $\rho = 100$ as showed in Figure 2. By changing θ from 1 to 20, the efficiency increases across all the simulations by almost 17%, while changing it from 20 to 100 provides much lesser benefits, in the order of 3-5%. This behavior is constant whatever the value of ρ is set to, as the distances between the curves clearly picture. Finally, we note how for low values of θ (e.g., $\theta = 1$), the effect of distrust dissipates quickly with few interactions, without any relevant difference on the Communication Efficiency for all values of ρ .

5 CASE STUDY

The case study we present is of a particular network configuration as showed in Figure 3. In that simulation, which we stopped at $\Gamma = 60793$, four clear communities of nodes have emerged, with two of them (i.e. *Cluster 3* and *Cluster 4* in Figure 3) loosely connected by few links and through a single node bridging most communication between the two clusters.

The network structure is of particular interest for studying how trust is distributed in these communities. All of them are formed by some core nodes (i.e. *proxies*, pictured as *orange* dots) having a rate of answers produced exceeding of more than 100% the answers they get from other nodes; some *peripheral* nodes (pictured in *blue*) that get more answers than those they produce; and finally some intermediate nodes (i.e., *ex-proxies*, pictured in *green*) that have played the role of proxies, but whose rate of answers is now below the threshold. For all the clusters we have studied how trust and interest values are distributed, showing that there is a clear relationship between the *topology*, the *trust distribution* and the *distribution of interests*.

Furthermore, the network structure is also interesting for the presence of that bridging node between the two clusters, which permits to study how communication is established between two separate communities, how it flows and how trust drives the dynamics of the attraction or repulsion between them.

Asymmetries in trust distribution. First we investigated the distribution of trust among nodes in the clusters. Table 1 shows the results calculated for the three node categories (i.e., proxy, ex-proxy and peripheral nodes) in each cluster. For each category, trust is represented as directed, that is a node category trusts other ones (labelled as “trust *from*” in Table 1) or a node category is trusted by other nodes (labelled as “trust *towards*” in Table 1).

Figure 4 provides a graphical representation of the distribution of trust among the different node categories. Trust is pictured as a gradient of color, from a low level of trust (*light blue*) to a high level (*dark*

Table 1: Trust directed to and from each type of node (*rows*) within the four clusters (*columns*). In **boldface** the highest trust towards nodes, while underlined is trust from nodes. Clusters are numbered from left to right starting from the the top-left corner.

	1	2	3	4
from PROXY nodes	15.5	17.9	8.3	8.5
towards PROXY nodes	39.5	41.6	19.0	23.5
from EX-PROXY nodes	23.1	22.1	11.9	12.4
towards EX-PROXY nodes	28.8	21.1	15.4	15.5
from PERIPHERAL nodes	<u>24.2</u>	<u>29.1</u>	<u>19.7</u>	<u>21.0</u>
towards PERIPHERAL nodes	14.6	11.7	6.9	9.9

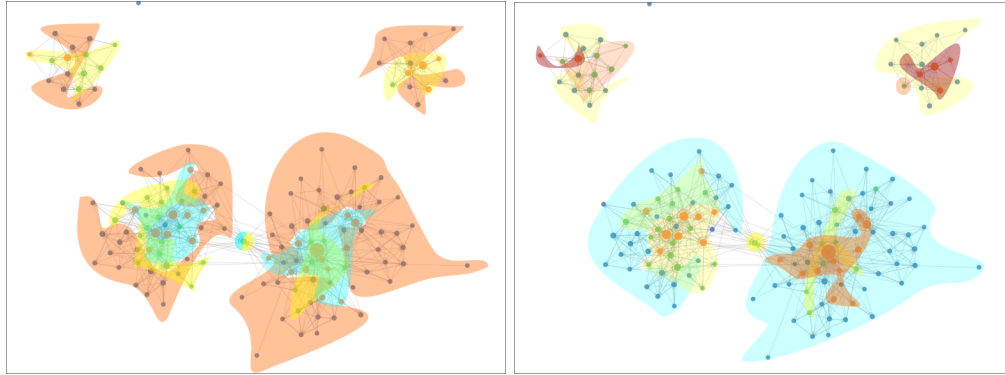


Figure 4: Trust *from* (left) and *towards* (right) nodes.

red). The picture on the *left* of Figure 4 shows the distribution of trust *from* the nodes (dark red nodes trust much others, light blue ones don't trust much, and yellow nodes are more balanced), whereas the picture on the *right* represents the distribution of trust *toward* the nodes (dark red nodes are highly trusted from others, light blue ones aren't trusted much, and yellow nodes are balanced). The results are described in Table 1. Proxy nodes act as highly trustworthy nodes (red in Figure 4 (right)) and on the opposite not trusting much other nodes (light blue in Figure 4 (left)); peripheral nodes, conversely, strongly trust other nodes (red in Figure 4 (left)), but are not trusted much (light blue in Figure 4 (right)). It is clear from these results that in our social network model trust seems to be partitioned between *trusting* and *being trusted* with few intermediate cases, which, broadly speaking, means that *who trusts a lot is, typically, not trusted much* and *vice versa*.

To better understand this behavior, we further analyzed if there was a relationship between nodes' trust and the distribution of interests associated to their topics. In other words, we made the hypothesis that there should be a skewed distribution of interests among nodes reflecting the different trust.

To this end, for each cluster, we considered the different node categories, and the interests associated to each category. We calculated the sum of the k interests with highest values, with k selected as the smallest number of different topics owned by proxies, ex-proxies or peripheral nodes of the cluster (i.e. in this way, we are sure that for all nodes categories, there are at least k interests to sum). Figure 5 shows how the distribution of interest values is indeed skewed and strongly localized among peripheral nodes, which are the less trusted (as previously showed in Figure 4). The opposite happens for more central nodes (proxies) and ex-proxies, which exhibit lower interest values. It seems then confirmed that *trustworthiness* is concentrated in nodes with low interest and high quality values, while *untrustworthiness* is a characteristic of nodes with high interest and low quality values.

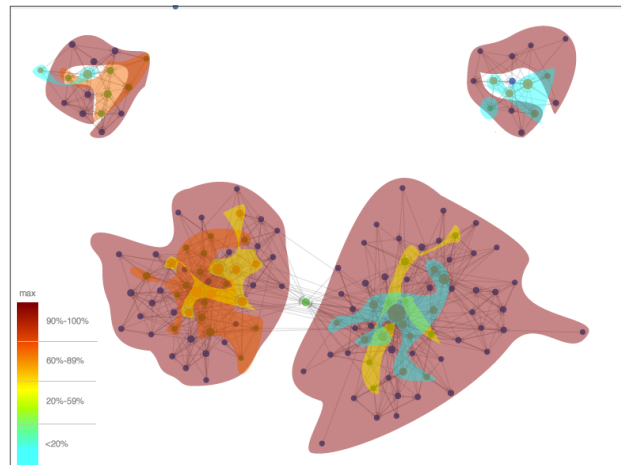


Figure 5: Interest distribution. High interest values are concentrated among *red* nodes, while lower values are located in the *light blue* nodes. Scales are evaluated locally for each cluster.

Unbalanced communication from trusted zones to untrusted zones. With respect to the bridging node and the communication between Cluster 3 and Cluster 4, we can observe that the communication flow between the two clusters through the bridging node is *asymmetric*. This is represented by trust gradients. In particular:

- In Figure 4 (*right*), the bridging node is uniformly yellow, meaning that its trustworthiness is greater than that of peripheral nodes. With respect to communication this means that it is similarly queried from Cluster 3 and from Cluster 4, and then it produces more answers than a typical peripheral node (which makes it more similar, then, to an ex-proxy);
- In Figure 4 (*left*), the bridging node is not uniformly colored, being yellow with respect to communications towards Cluster 4 and light blue with respect to Cluster 3. This means that the bridging node trusts more Cluster 4 nodes than those of Cluster 3. The reason is that it receives more answers from Cluster 4 than from Cluster 3.

We have further inspected the asymmetric communication flowing between the two clusters through the bridging node. Given the previous observation, that the bridging node typically queries Cluster 4's nodes and is queried by Cluster 3's nodes, we made the hypothesis that the two effects were not disjoint, rather there could be a communication flow from Cluster 4 to Cluster 3 mediated by the bridging node.

To verify the hypothesis, we firstly checked which topics were involved in communication between the bridging node and Cluster 3's nodes and between the bridging node and Cluster 4's nodes.

Figure 6 (*left*) shows the results. We can see that the bridging node communicates with Cluster 3, therefore it likely responds to queries, mostly about topics 17 and 8 (topics 9 and 1 are the remaining). With Cluster 4, instead, it typically interacts, therefore it likely queries, about topics 17, 8, 3 and 19 (with topics 9 and 1 with only a single interaction). Topics 17 and 8 are then good candidates to represent the knowledge that flows from Cluster 4 to Cluster 3. Figure 6 (*right*) adds a new evidence for our hypothesis by showing which nodes have been involved in communication regarding topics 17 and 8. We see again the asymmetry between the two clusters. With Cluster 4, the bridging node interacts with only three different nodes, querying them about topic 8 and 17. Differently, from Cluster 3, many more nodes interact with the bridging node, querying it.

The explanation for this effect is that the bridging node, in its Personal State, has a high value of the *interest* attribute associated to topics 17 and 8. Nodes of Cluster 4 are the ones that have been selected as the best node choice and at each interaction the bridging node increases the quality associated to topics 17

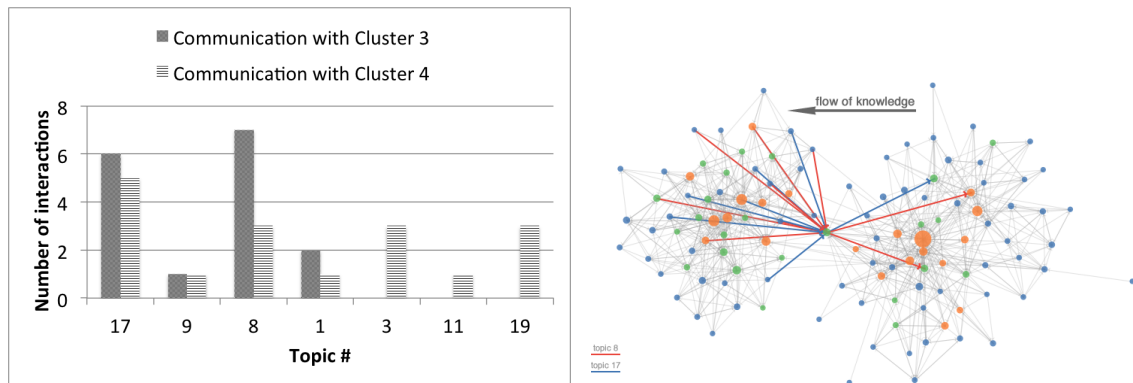


Figure 6: Flow of knowledge. On the arcs on the right of the bridging node pass the increase in quality that is shared over the edges on the left of the bridging node.

and 8 and the corresponding interest. Therefore, it has progressively become more expert on those topics, therefore more likely to be selected as the peer for an incoming query. On the other side, some nodes of Cluster 3 own either topic 17 or topic 8 and they select them for a new interaction. The bridging node has been selected by chance by one of those Cluster 3's nodes, becoming friend. From that first connection, the triadic closure mechanism has fostered other Cluster 3's nodes to select the bridging node as the best friend-of-friend and this way more communication have been established.

Communication about topic 17 and 8 then flows from Cluster 4 to Cluster 3 passing through the bridging node: The *asymmetry* between direction of trust and direction of communication is evident, being the latter directed from a *trusted* zone to an *untrusted* zone with respect to the bridging node.

6 CONCLUSIONS

In this work, we have presented a dynamic social network model based on the exchange of knowledge among nodes and on trust, both as a factor driving the network evolution. Results of the simulation showed several emergent properties of the system, in which trust evolves asymmetrically among group of nodes.

A case study has been analyzed focusing on a particular network configuration in which the effects controlling the flow of knowledge among different community of nodes were particularly clear. We speculate that, even if our model is yet untested with respect to real-case scenarios, it could be considered as an interesting and potentially meaningful model to analyze real dynamics, because the emergent behavior and characteristics of the model presented are widely studied and supported in literature, and because it permits to evaluate and possibly manipulate even very localized dynamics of group formation and knowledge exchange. There are some extensions that should be done. One of the most important is to further improve the modeling of trust, for instance by including a negative effect able to explicitly model distrust events. Another is the definition of roles to be applied to nodes with different mode of interaction. Finally, one of our goals for future works is to find a real case study and analyze the ability of our model to describe its dynamics. Nevertheless, the results already achieved do appear as a promising and interesting research direction and worthy of further development.

REFERENCES

- Barber, B. 1983. *The Logic and Limits of Trust*. New Rutgers University Press.
- M.J Brzozowski and D. M Romero 2010, August. "Who Should I Follow? Recommending People in Directed Social Networks". Accessed September 9, 2011. <http://www.hpl.hp.com/research/scl/papers/follow/>.

- Burt, R. S. 2001. "Bandwidth and Echo: Trust, Information, and Gossip in Social Networks". In *Networks and Markets: Contributions from Economics and Sociology*, edited by A. Casella and J. E. Rauch, 30–74. Russell Sage Foundation.
- Hanaki, N., A. Peterhansl, P. Dodds, and D. Watts. 2007, July. "Cooperation in Evolving Social Networks". *Management Science* 53 (7): 1036–1050.
- Jin, E. M., M. Girvan, and M. E. J. Newman. 2001, September. "Structure of growing social networks". *Physical Review E* 64 (4): 1–8.
- Newman, M. E. J. 2001, January. "The structure of scientific collaboration networks". *Proceedings of the National Academy of Sciences of the United States of America* 98 (2): 404–409.
- Newman, M. E. J. 2003. "The Structure and Function of Complex Networks". *SIAM Review* 45 (2): 167–256.
- Newman, M. E. J. 2004, April. "Coauthorship Networks and Patterns of Scientific Collaboration". *Proceedings of the National Academy of Sciences of the United States of America* 101 (Suppl 1): 5200–5205.
- Newman, M. E. J., and M. Girvan. 2003. "Mixing Patterns and Community Structure in Networks". In *Statistical Mechanics of Complex Networks*, edited by R. Pastor-Satorras, M. Rubi, and A. Diaz-Guilera, Volume 625 of *Lecture Notes in Physics*, 66–87. Springer Berlin / Heidelberg.
- Newman, M. E. J., and P. J.. 2003, September. "Why social networks are different from other types of networks". *Physical Review E* 68 (3).
- Skyrms, B., and R. Pemantle. 2000. "A dynamic model of social network formation". *Proceedings of the National Academy of Sciences of the United States of America* 97 (16): 9340–9346.
- Thomborson, C. 2010. "Axiomatic and behavioural trust". In *Trust and Trustworthy Computing*, edited by A. Acquisti, S. Smith, and A.-R. Sadeghi, Volume 6101 of *Lecture Notes in Computer Science*, 352–366. Springer Berlin / Heidelberg.
- Tyler, J. R., D. M. Wilkinson, and B. A. Huberman. 2003. "Email as spectroscopy: automated discovery of community structure within organizations". In *Communities and technologies*, edited by M. Huysman, E. Wenger, and V. Wulf, 81–96. Kluwer Academic Publishers.
- Walter, F. E., S. Battiston, and F. Schweitzer. 2008, February. "A model of a trust-based recommendation system on a social network". *Auton Agent Multi-Agent Syst* 16 (1): 57–74.
- Watts, D. J., and S. H. Strogatz. 1998, February. "Collective dynamics of 'small-world' networks". *Nature* 393:440–442.

AUTHOR BIOGRAPHIES

LUCA ALLODI is a Ph.D. student in the Computer Science Department of the University of Trento, Italy. He received his master degree in Computer Science from the University of Milan, Italy. His research interests are in security economics and the perception of security risks in society; as a master degree student he also worked on various research projects concerning the evolution of network dynamics and the diffusion of information in social networks. His email address is luca.allodi@disi.unitn.it

LUCA CHIODI is a master student in Information Security at the University of Milan. He received a bachelor degree in Systems and Networks Security from the University of Milan. His research interests are on modeling and simulation of social networks, and on the analysis of the flow of informations between social agents. His email address is luca.chiodi@studenti.unimi.it

MARCO CREMONINI is an Assistant Professor at the Department of Information Technology of the University of Milan, Italy. He got his master degree and Ph.D. in Electronic and Information Technology Engineering at the University of Bologna, Italy. He previously worked as a Research Assistant at the Institute for Security Technology Studies (ISTS) of the Dartmouth College, NH, USA. His research activity is focused on dynamic social networks, risk analysis and decisions under uncertainty, security economics, privacy, and security in ubiquitous computing. He is member of the Editorial Board of *Infosecurity Magazine* UK (Elsevier). His email address is marco.cremonini@unimi.it