Analysis & Generative Model for Trust Networks

Pranav Dandekar
Management Science & Engineering
Stanford University
Stanford, CA 94305
ppd@stanford.edu

ABSTRACT

Trust Networks are a specific kind of social network where edges in the network have positive and negative signs connoting friendship/trust and antagonism/distrust respectively. While the last few years have seen a rich body of work on generative models for social networks, there hasn't been much work on understanding in what ways trust networks differ from social networks and to create generative models for them. In this work, we analyze the network structure and properties of two online trust networks – Epinions and Slashdot – using theories of structural balance and social capital from social psychology and propose a generative model for *undirected graphs* based on these theories that produces synthetic graphs which agree with observed data on key network properties. A generative model for directed graphs is part of ongoing work.

1. INTRODUCTION

Over the last few years, network analysis has attracted a lot of interest in the computer science, sociology and economics research communities. Unsurprisingly, this interest has coincided with widespread adoption of social networking technologies such as email, IM, SMS, etc. among the general population. This has resulted in the availability of large scale real world network data which has heretofore been hard to find for sociologists and economists studying networks. It has also resulted in a number of interesting questions being asked about the structure and dynamics of networks. Along with efforts to understand the structure of networks, there have been concomitant efforts to create realistic and tractable mathematical models that explain the evolution and structural properties of networks observed in data. These mathematical models help create synthetic networks that are similar to real networks when real network data is hard to obtain. This is especially true in large-scale machine learning applications that work on data collected off the web where it is hard to find labeled "ground truth" that can be used to train and/or evaluate algorithms.

A specific kind of social networks are *trust networks*. These are graphs whose edges have positive and negative signs denoting friendship/trust and antagonism/distrust respectively. A rich part of social network theory is devoted to studying trust networks and the notion of *structural balance* which governs the evolution and stability of trust networks. Trust networks have begun to proliferate in the last few years as an increasingly larger number of websites rely upon some notion of ratings/testimonials for facilitating interactions between strangers. Some examples of such websites are Ebay, Amazon, Epinions, CouchSurfing, etc.

This project aims to bring together these two branches of social network theory – study of structure and evolution of social networks, and the study of structural balance and trust networks – in trying to understand whether trust networks are structurally different from other social networks studied heretofore. I will analyze the structural properties of two (potentially three) trust networks and compare and contrast them with results on similar-sized social networks. I will then use ideas from structural balance theory to either extend existing generative models to account for positive and negative edges or come up with a new generative model that will explain the properties observed in the data.

The next section overviews the rich literature on generative models from the last decade as well as concepts from structural balance theory and some recent work on propagation of trust in networks. Section 3 will contain a detailed analysis of the two trust networks which we had access to, as well as a smaller Wikipedia who-votes-for-whom trust network whose analysis we gleaned from an unpublished manuscript. Finally, we will present our generative model and compare synthetic graphs generative using the model with the real-world networks that we analyzed.

2. BACKGROUND

Since the scale of real-world networks has exploded over the last two decades or so, it has become harder to visualize what the network looks like in its entirety and to reason about individual nodes and their properties (e.g degree, centrality, etc.). This has given rise to an active area of research identifying the macro-properties of networks that best enable us to completely and accurately characterize them. Some of the properties that have been especially useful in distinguishing social networks from other naturally occurring networks have been the degree distribution (which is heavy-tailed in most social networks), the diameter (longest shortest-path between two nodes in a network, which is small for social networks), clustering coefficient (a measure of the locality of edges, which is significantly higher in social networks than in random graphs), and densification (average degree per node increases with the size of the network).

2.1 Overview of Generative Models

One of the first network models that explained heavy-tailed degree distributions observed in data was *Preferential Attachment* by Barabási and Albert [2]. In this model, arriving vertices create an edge with existing vertices with probability proportional to the degree of the existing vertex. This leads to a rich-get-richer phenomena which the authors observed in a wide variety of large-scale networks such as citation networks, actor-collaboration networks, power-grids, etc. A number of variations of this model appeared subsequently such as the "coping model" in [5] and the "forest-fire model" in [7]. While the original model does not have the "small-diameter" property exhibited by real-world networks, the forest fire model does exhibit a small (and shrinking) diameter, densification as well as the power law distribution for in-degree and out-degree.

Another line of research started with the "small-world" model in [8] which successfully explained the small diameter and high clustering coefficient observed in real networks. The *Kronecker graph model* in [6] uses a recursive graph construction model that has the small-diameter property, the heavy-tailed degree distribution and the densification property. The paper analyzes the model on a static graph of Internet Autonomous Systems and a temporal graph of research collaborations in High-Energy Physics.

2.2 Overview of Trust Networks

One of the first analyses of structural balance in social networks was done by Cartwright and Harary in the 1950s [3] where they noted that a triad that has one or three negative edges is unstable and went on to define the *Structural Balance property* for a complete graph as: a complete graph is structurally balanced if all triads have either zero or two negative edges. In subsequent work they also defined a weak structural balance property after observing that in real networks it is much more common to see triads with one negative edge which leads to changes in network topology (e.g. conversion of the negative edge to positive) and therefore they define complete graphs as being weakly balanced if none of the triads have exactly one negative edge.

Note that the structural balance theory is a theory of undirected edges, whereas in most real-world online applications the edges are directed. Moreover, unlike regular (unsigned) social networks, the existence of a directed signed edge from u to v cannot be assumed to mean that an edge from v to u of the same polarity also exists. In other words, directed signed networks cannot be converted to undirected signed networks without some (and potentially critical) loss of semantics. A competing theory in sociology is the theory of status which is better suited to understand signed directed edges. This theory says that edges in trust networks impute expertise, i.e., a positive edge from u to v implies that u puts v on a "higher pedestal", or considers v a greater expert. A direct consequence of this is that if there is a negative edge from u to v and another negative edge from v to w, then the theory predicts that there would be a negative edge from u to w (in other words, distrust is transitive). This is contrary to the theory of balance which says that "the enemy of my enemy is my friend".

More recently, [1] studies the evolution of trust networks under specific dynamics. In *Local Triad Dynamics* (LTD) at each time step they pick a triad at random and flip a sign on it if it is unbalanced (i.e. has one or three negative edges). In *Constrained Triad Dynamics* the number of imbalanced triads cannot increase with any update. They show that as the propensity p for friendly links in an update increases, the network undergoes a dynamic phase transition from steady state to one where all links are friendly. They also show that a finite network always reaches a steady-state where none of the links are negative.

	Epinions	Slashdot
Nodes	131,828	82,164
Edges	841,372	549,981
+ edges	85.3%	77.4%
– edges	14.7%	22.6%
Triads	13,372,109	1,514,209

Table 1: Dataset Statistics

There has also been some work on prediction of missing polarities in trust networks based on the polarities and/or magnitudes thereof of the surrounding edges. One of the first works on this was [4] where they hid some of the labels in an Epinions trust network and tried to propagate trust values from neighboring nodes/edges. They couldn't accurately predict values on edges (u, v) where (u, w) and (w, v) were both negative, which can be explained by the structural balance theory.

However, there hasn't been much work on trying to analyze the structural properties of trust networks with a view to creating a generative model for them either by extending existing generative models or using ideas from theories of balance and/or status from the social psychology literature.

3. DATASET DESCRIPTION

We study the properties of two online trust networks – Epinions and Slashdot. Epinions is a online product review website where users can express trust or distrust of other users. Slashdot is a technology/science news aggregator website where users can state that they are a "friend" or a "foe" of another user. A third trust network that we planned to analyze but could not get access to in time was the Couchsurfing network. Couchsurfing is a non-profit that allows users to host other users that might be visiting from out-of-town for a few days. Given the sensitivity of the real-world interaction involved with Couchsurfing (i.e. the trust expressed by a user for another user on the website is a manifestation of their interaction in the real-world), it would be interesting to study its properties and compare them to those from Epinions and Slashdot that exist mostly in the online world. We plan to include Couchsurfing data as part of extending this work (ongoing with Juré).

Both the datasets we analyze are directed signed networks (i.e. the edges are positive or negative). The dataset statistics of the undirected graph (i.e. ignoring direction) are listed in Table 1. Both the networks are very sparse with the average degree of a node being less than 10. The edges are overwhelmingly positive, and perhaps interestingly, the ratio of positive edges p is close to 0.8 in both cases. This is consistent with observations made with respect to diverse settings such as fraction of positive votes on Amazon and fraction of positive ratings on eBay that on the web, users are much more likely to express positive opinions (of products, reviews, or people) than to say bad things.

Both the positive and the negative degree distributions of both networks follow a power law (see fig 1 and fig 2) but the positive degree distribution has a smaller α ($\alpha = 2.03$ for Epinions and $\alpha = 1.79$ for Slashdot). The overall degree distribution is also a power law with an exponent roughly equal to the respective exponent for the positive subgraph.

An important difference between the two networks is the degree

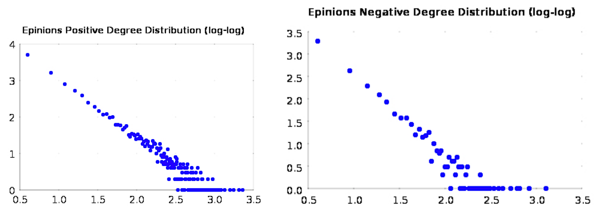

Figure 1: Epinions Positive Degree Distribution: Power law with $\alpha = 2.03$; Negative Degree Distribution: Power law with $\alpha = 2.40$.

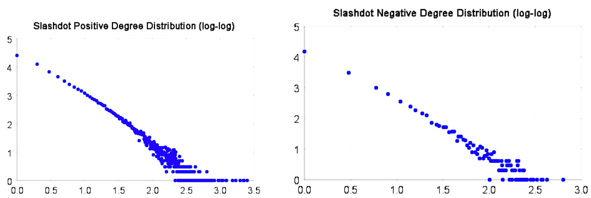

Figure 2: Slashdot Positive Degree Distribution: Power law with $\alpha = 1.79$; Negative Degree Distribution: Power law with $\alpha = 2.04$.

of clustering as evident by the number of triads. Epinions has a significantly higher clustering coefficient compared to Slashdot. The frequency of triads observed in the data (see table 3) is only partly explained by the theory of balance. While the all-positive (T_0) triads were significantly more frequent than predicted by the sign distribution, so were the all-negative (T_3) triads (which are considered unbalanced by the theory of balance, but consistent with the theory of status). Therefore, it is clear even from the analysis of these networks as *undirected* graphs (i.e. ignoring the direction of the edges) that the theory of balance does not fully explain the observed data.

Taking into account the direction of the edges, we find that expectedly, the indegree and outdegree distributions of both networks also follow power laws (see fig 3 and fig 4).

We also looked at how many edges are reciprocated (i.e. where there is an edge from u to v and from v to u) and a surprisingly high number of edges are reciprocated. On Epinions, 259,752 (30%) of the edges are reciprocal whereas on Slashdot 97,482 (18%) of the edges are reciprocal. An overwhelmingly large fraction (compared to chance) of the reciprocal edges are $(+, +)$ which is consistent with observations in other datasets that people feel obligated to reciprocate praise/trust leading to high incidence of "mutual admiration societies". At the same time, the fraction of $(-, -)$ is in line with that predicted by chance which could be explained by the fact that most trust networks (including Epinions and Slashdot) do not expose a user's distrust ratings (i.e. they do not list all the users that a given user distrusts). If this data were displayed on a user's profile along side trust ratings, it would likely lead to a "tit-for-

Symbol	Description
$\#T_i$	Number of triads of type i
Δ	Total number of triads in the network
p	fraction of positive edges in the network
$p(T_i)$	Fraction of triads T_i, $p(T_i) = \#T_i/\Delta$
$p_0(T_i)$	A priori prob. of T_i (based on sign dist.)

Table 2: Table of Symbols

Triad T_i	$\#T_i$	$p(T_i)$	$p_0(T_i)$
Epinions			
T_0	11,658,420	0.872	0.621
T_1	926,945	0.069	0.054
T_2	698,951	0.052	0.322
T_3	87,793	0.007	0.003
Slashdot			
T_0	1,271,899	0.839	0.464
T_1	109,584	0.072	0.406
T_2	116,428	0.077	0.119
T_3	16,298	0.012	0.011

Table 3: Triad Counts (see table 2 for description of symbols)

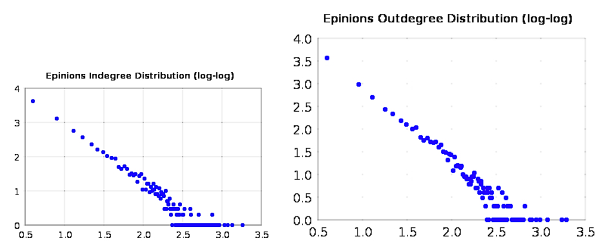

Figure 3: Epinions In-degree and Out-degree distributions follow a power law.

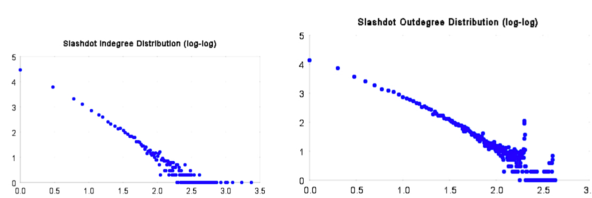

Figure 4: Slashdot In-degree and Out-degree distributions follow a power law.

	Epinions	Slashdot
Reciprocal Edges	259,752	97,482
++ edges	96.09%	86.6%
+– edges	2.09%	4.0%
– – edges	1.82%	9.4%

Table 4: Reciprocity

tat" behavior resulting in a significantly higher fraction of $(-,-)$ edges.

4. GENERATIVE MODEL & RESULTS

In this section we present a generative model for *undirected* signed graphs and compare graphs produced by the model with the Epinions and Slashdot datasets along some of their undirected graph properties. The properties that we are particularly interested in is the power law distribution of positive and negative degrees, the fraction of positive and negative edges and the numbers and fractions of the four types of triads. A generative model for directed graphs is part of ongoing work with Juré.

4.1 What Doesn't Work

A natural first attempt at a generative model for trust networks is to take an existing model for *unsigned* networks and extend it for signed networks by labeling edges + or – in post-processing. We tried this approach using the preferential attachment model by Barabási and Albert [2]. We constructed a graph using their model and then used some heuristics to label the edges + or – (first randomly labeling and edge + with probability p and later, using the embeddedness of the edge to determine its sign). While this model gives us the power law degree distributions for positive and negative edges it does not give us the number of triads. The datasets have a significantly larger number of triads (and consequently a higher clustering coefficient) than that produced by the basic PA model. Therefore, we abandoned this approach in favor of one where we construct the graph from scratch and label edges as we go based on principles from the theories of balance and status.

4.2 Our Model

We base our model loosely on the copying model in [5]. This gives us the power law degree distributions since it has a "rich-get-richer" flavor. Additionally, it also gives us a higher clustering coefficient than the simple PA model.

4.2.1 *Intuition*

The intuitive idea behind the model is that when a user i joins the network, he begins interacting with other users. If their interaction with a user j is positive, then over time i discovers the users that j knows. If i trusts j then i will trust anyone that j trusts. However, if i trusts j, he won't automatically start distrusting everyone that j distrust. However, the likelihood that i will trust user k will go down (to 50% in our model) if j distrusts k. On the other hand, if i does not think highly of j (i.e. if (i, j) is negative), then he will automatically form negative edges with everyone that j distrusts. However, i will disregard people that j trust because i doesn't have a high opinion of j to begin with. This model is based in spirit on the theory of status, although there is an asymmetry in creating edges to friends of a user that you don't trust versus the enemies of your friend. The probability r_p determines the fraction of neighbors of j that i discovers given that (i, j) is positive whereas the probability r_n determines the fraction of neighbors of j that i discovers given that (i, j) is negative. We impose the constraint on the model that $r_p \geq r_n$ which intuitively corresponds to the fact that i is much more likely to interact with j's neighbors if (i, j) is positive.

The parameter p along with the parameter r_p controls the fraction of edges in the generated graph that get labeled +, whereas r_p and r_n control the number of triads and their relative frequencies.

Algorithm 1 Trust-Copy Model (n, p, r_p, r_n)

```
for i = 1 to n do
  Pick a node j uniformly at random from {1, ..., i − 1}
  Create edge (i, j)
  Label the edge + with prob. p, and – with prob. (1 − p)
  for all neighbors k of j do
    if (i, j) is + then
      Create edge (i, k) with prob. r_p.
      if (j, k) is + then
        Label edge (i, k) +.
      else
        Label edge (i, k) + or – with prob. 1/2.
      end if
    else
      if (j, k) is – then
        Create edge (i, k) with prob. r_n.
        Label edge (i, k) –.
      end if
    end if
  end for
end for
```

	Epinions	G(119287, 0.58, 0.75, 0.23)
Nodes	119,287	119,287
Edges	841,372	915,864
+ edges	85.3%	84.5%
– edges	14.7%	15.5%

Table 5: Graph Statistics: Epinions v/s Synthetic Graph

4.3 Results

In this section we compare the graphs generated using our model with the Epinions and Slashdot trust networks using key network properties.

Table 5 compares the basic statistics of the Epinions graph with that of a graph consisting of the same number of nodes generated using our model with parameters ($p = 0.58, r_p = 0.75, r_n = 0.23$). The graph generated by our model closely resembles the Epinions network in terms of the basic graph statistics. The generated graph also exhibits a power law distributions of positive and negative degrees (see fig 5). The negative degree distribution has a higher α than the positive degree distribution just like that for Epinions. Table 6 compares the graph generated by the model with the Epinions networks in terms of triad density and the frequencies

	Epinions	G(119287, 0.58, 0.75, 0.23)
T_0	0.872	0.901
T_1	0.069	0.059
T_2	0.052	0.030
T_3	0.007	0.010
Number of Triads	13,372,109	11,170,091

Table 6: Triad Frequencies: Epinions v/s Synthetic Graph

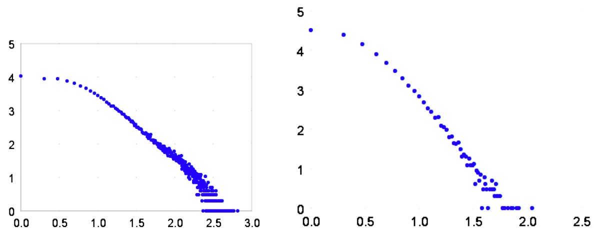

Figure 5: Synthetic Epinions-like graph: Positive ($\alpha = 2.02$) and Negative ($\alpha = 2.57$) Degree Distribution

	Slashdot	G(82164, 0.67, 0.58, 0.25)
Nodes	82,164	82,164
Edges	549,981	418,367
+ edges	77.4%	84.2%
– edges	22.6%	15.8%

Table 7: Graph Statistics: Slashdot v/s Synthetic Graph

of various triad types, and again finds a high degree of consistency between the two sets of numbers.

Next we generated a graph consisting of the same number of nodes as the Slashdot network with parameters ($p = 0.67, r_p = 0.58, r_n = 0.25$). Table 7 compares the basic statistics of the Slashdot graph with those of the graph generated by our model. We note that the Slashdot graph has a much lower clustering coefficient compared to the Epinions graph, and our model does not resemble the Slashdot graph as closely as it does the Epinions graph (with somewhat fewer number of edges, our graph has many more number of triads).

Figure 6 shows the positive and negative degree distribution of the generated Slashdot-like graph. While the α values of the two distributions are higher than the corresponding values for the Slashdot graph, they are not significantly different. A potential fix to this might involve parameterizing the prob. that (i, k) is positive given that (i, j) is positive and (j, k) is negative (that probability is currently fixed at 1/2).

5. CONCLUSION & FUTURE WORK

This work analyzed the structural properties of the Epinions and Slashdot trust networks and showed how the structural balance theory and the theory of status don't fully explain the sign patterns in the data. We described a model that does not work well and explained the reasons for it. We also proposed a model that uses a combination of the theory of balance and status and produces graphs that closely resemble the Epinions and Slashdot networks in terms of degree distribution of positive and negative edges, the density of triads and the fraction of various triad types.

	Slashdot	G(82164, 0.67, 0.58, 0.25)
T_0	0.839	0.867
T_1	0.072	0.063
T_2	0.077	0.049
T_3	0.012	0.019
Number of Triads	1,514,209	2,038,774

Table 8: Triad Frequencies: Slashdot v/s Synthetic Graph

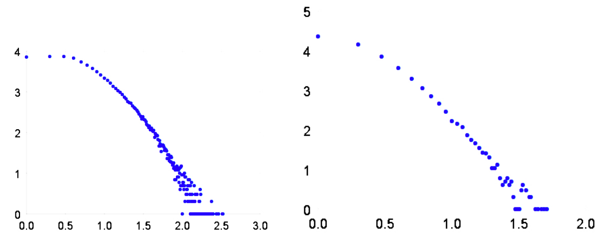

Figure 6: Synthetic Slashdot-like graph: Positive ($\alpha = 2.02$) and Negative ($\alpha = 3.57$) Degree Distribution

There are three directions to extend this work in. The first involves analyzing the CouchSurfing dataset (once we obtain it) and validating our model against the Couchsurfing network. That validation will be significant because the Couchsurfing network is much closer to a real-world trust network compared to Epinions and Slashdot since it is based on real life interactions of users, and the social psychology theories that form the basis of our model also stem from real-world social interactions. The second direction is to extend our model to account for the direction of edges and in particular account for the reciprocity patterns that we observed in the datasets. Lastly, we would like to obtain theoretical results about the degree distributions (positive, negative, in-degree and out-degree), clustering coefficient, and triad frequencies. All these directions will be explored as part of ongoing work with Juré.

6. REFERENCES

[1] T. Antal, P. L. Krapivsky, and S. Redner. Dynamics of social balance on networks, 2009.

[2] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[3] D. Easley and J. M. Kleinberg. *Networks, Crowds and Markets: Reasoning about a highly connected world*. Cambridge University Press, Cambridge, England, 2010.

[4] R. Guha, R. Kumar, P. Raghavan, and A. Tomkins. Propagation of trust and distrust. In *WWW '04: Proceedings of the 13th international conference on World Wide Web*, pages 403–412, New York, NY, USA, 2004. ACM.

[5] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic models for the web graph. In *FOCS '00: Proceedings of the 41st Annual Symposium on the Foundations of Computer Science*, page 57, 2000.

[6] J. Leskovec, D. Chakrabarti, J. M. Kleinberg, C. Faloutsos, and Z. Ghahramani. Kronecker graphs: An approach to modeling networks, 2009.

[7] J. Leskovec, J. M. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *KDD '05: Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 177–187, 2005.

[8] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.