

On Network formation,
(Sybil attacks and Reputation systems)
(Position Paper)

George Danezis and Stefan Schiffner

January 12, 2007

Abstract

We propose a model of network formation in peer-to-peer networks, that allows us to observe their susceptibility to sybil attacks against routing security. Peers try to selfishly fulfill their communication needs, by connecting directly to communication partners ('friends') or indirectly through stranger nodes. We assess the strategies nodes will follow depending on the topology of the friendship graphs, and the number of links nodes are allowed. We show that it is common to connect to friends, therefore automatically foiling exogenous attacks. A roadmap of further work, including realistic networks, adversaries and using reputation systems is discussed.

1 Introduction

Peer-to-peer systems and generally systems that are distributed across multiple trust domains present a unique challenge to security designers and engineers. The disparate entities that come together to form such systems cannot realistically be expected to behave according to a pre-determined set of protocols, in particular at times when following such protocols would conflict with their own objectives. Many studies have appeared [4] on the problem of free riding in content or resource sharing networks – which is a typical example of selfish (yet rational) behaviour.

Aside from otherwise honest nodes behaving rationally (and selfishly), external attackers with objectives that are different from honest players may also try to influence the functioning of the systems. The most usual objectives of such attackers would be surveillance, to gather as much information as possible about actions of other nodes; or disruption, preventing nodes from carrying out actions within the distributed system. In peer-to-peer systems such external adversaries can have orders of magnitude more power than any individual node, and may be able to masquerade their identity and appear as multiple nodes. This is called a sybil attack [3].

The aim of our research is to uncover the fundamental mechanisms that allow such sybil attacks or can be used to defend against them, in the context of rational distributed and selfish nodes. So far there has been a separation between research on security or efficiency problems resulting from selfish behaviour, and

the problems of sybil attacks. Yet the two are intimately interconnected as the topology and strategies that rational nodes will choose affects parameters of the distributed system, such as topology, that are key to the success or not of sybil attacks. In turn the knowledge that a sybil attack may be possible is bound to influence nodes in their choices of strategies: we expect them to balance their need to extract maximal utility from the network, with the needs to (personally) not be the victims of a sybil attack.

The key tools we use to study the interactions of rational strategic nodes trying with sybil attacks are:

- Game theory: allows us to model choices of strategic players, according to the utility of the outcomes that different strategies would lead to.
- Network formation: the strategies we will consider will have an impact on the connectivity of the nodes, and which other nodes in the network they rely on to reach their objectives.
- Social network theory: to model reality it is a good idea to move away from the assumption that nodes have random needs, and model communication needs that are more likely to be observed in real networks. These include a power law distribution of degrees, and cliques, and easy of routing.
- Simulation: it is rather difficult to find satisfactory analytical answers to all the question we put forth, so we have to resort to simulating networks with multiple nodes.

We will discuss in the next sections how we combined those techniques, and our (preliminary) results.

1.1 The Sybil Attack and Reputation Systems

The Sybil Attack [3] is a powerful way to get more benefits from a cooperative system than other users, without contributing the corresponding work. A sybil adversary creates a large amount of cheap nodes or pseudonyms (sybils) that act in the systems as separate entities, vouching for each other if necessary to fool the reputation system.

For such an attack to be possible some conditions must hold:

- Creating new nodes is cheap.
- New nodes are equal.
- New nodes are not discriminated.

How cheap introducing nodes is depends on the adversary and many reputation systems like eBay's try to make creating new accounts somehow expensive, by way of validating emails and CAPTCHAS. Equality means that new nodes are indistinguishable for the reputation system. Non-discrimination means that a new node can give ratings.

Reputation Systems A Reputation systems in general is a filter system for content. There is a wide range of reputation systems some provide qualitative information (like books reviews) some only assign a value of ‘goodness’. Since we want to study systems for automated filtering, we are more interested in the latter ones.

Googles pageRank [1] sorts automatically web pages in terms of likelihood of importance for the searcher. It is based on the random surfer model: A surfer starts with a random page and follows random links on this page to the next one after a (random) time he gets bored and he starts on an new (random) page. The rank of a page is the likelihood that the random surfer visit this page.

EigenTrust [6] Is a system to rate users in a system. It assumes that all users have the same preferences and Trust is arbitrary transitive. Hence the reputation of a user is the same independent of the querying user. The authors define a trust matrix and show that the eigen vector is equivalent to their trust vector. They also give a probabilistic interpretation which is very similar to the random surfer of PageRank.

A sybil attack on PageRank also known as link farming. The attacker creates a huge amount of web pages which all reference the page of the user. An attack on eigenTrust works in a similar way.

In general no symmetric reputation system can be sybil proof [2] A adversary which can create arbitrary many nodes can easily create as many nodes as in the original network and build an exact copy of all trust relations. The reputations in the system will be the same as in the original network. Hence the attacker is owner of one node with the highest possible trust value. The trust network of the adversary is completely disconnected to the original one. that means he does not need the help of other nodes to get the high reputation.

Before dipping further into sybil attacks and reputation systems, we will study the process by which networks form to accomplish the most basic task: routing messages amongst them.

2 A simple model

Game theory is conceptually a powerful tool that allows us to make predictions on how strategic players would behave, when all strategies interact with each other to dictate the final outcome. Sadly most games are too complex (in the complexity theory sense) to reason about, or to solve using today’s computing technology. Generally, a game with N players, having each M possible strategies, requires an effort of about $\mathcal{O}(M^N)$ to ‘solve’ using brute force. Slightly more efficient algorithms exist for simple games, e.g. where all players do not influence the utility of all others. For those reason we tried to capture the essence of what we are looking for for, i.e. what makes networks susceptible to the sybil attack, in a simple minded model.

2.1 The model

The key parameters of our model are as follows:

- We assume we have a set of *nodes* N that are to be connected in a network.

- Each node n has a set of *friends*, or cardinality say F_n , that he wishes to talk to. Friendship is symmetric so if A is friends with B, then B is also friends with A.
- Each node also has a *link budget* of allowed links he can use, of say L_n , for each node $n \in N$. As we shall see we will require $L_n < F_n$. Links are symmetric and (unlike friendship) consume from the link budget of both nodes at the ends of the link.
- Given a graph of links between nodes the *utility* of each node is defined as the negative sum of the length of the shortest paths to all his friends. This means that the objective of nodes is to use the network to talk to their friends, and the shortest the path to each of them, the better. (We use the negative sum, so that utility increases as paths lengths decrease.)

We have to pause before considering on one hand the strategies being offered to nodes (which will affect how the link graph is formed), and the introduction of an adversary.

As we stated before the objective of nodes is to communicate with their friends, in the minimum number of hops possible. It is clear that if nodes had a link budget that was at least as great as their number of friends, the game would have a straightforward dominant strategy (or graph to be exact), which would be for each node to connect to their friends. Each node n then would achieve a utility of $-F_n$, i.e. connect to all their F_n friends in one hop. What is even more interesting in this case is that no node relies on any other node to ‘transit’ its communications to their friends, since there is a direct link. It is therefore hard to see how one could model an adversary to disrupt such a network. This leads us to our first remark:

Remark 1. *If nodes have the ability to connect directly to everyone they want to talk to, there is no possible adversary.*

In order to find more ‘interesting’ link topologies we require each node to have a shortage of links. The key intuition behind this is that nodes will be forced to relay communications over each other, making the introduction of an adversary possible. This case is also more realistic: computers have a limited number of independent connection points to networks (the Internet say), yet they communicate to more than the connection points – it is a rule that in the Internet communications are relayed over others. This is also true for overlay and peer-to-peer networks.

Ideally each node should have the freedom to dispose of their link budget as they wish, in order to maximize its utility. There are though two key problems with this approach: link symmetry, and again complexity. First it is only fair to assume that a link between two nodes can only be established if both parties agree to establish it. This is an established assumption in network formation [5], and does not seem to pose any further problems. Second and more problematic is the number of games that are possible if each node has full freedom to choose who to connect to. Assuming a one-shot game with perfect information, where all nodes bid for links (up to their budget), and links that have a bid from both concerned nodes get established. The number of possible games ($\prod_{n \in N} \binom{N-1}{L_n}$) is extremely large, even for moderately sized networks.

An alternative is to use a restricted set of strategies and use them to seed a deterministic (non-strategic) network formation algorithm. The small set of strategies on offer should encapsulate the decisions of nodes concerning what we are interested in researching, while the network formation algorithm should mimic as much as possible a realistic process of network formation.

Our key interest is the study of the effect of sybil attacks on networks, the conditions under which they arise, the influence they have on nodes and possible defences. Sybil attacks are by their very nature exogenous, since they require nodes to be talking and using ‘strangers’ to provide some network service. In case those strangers are sybils, they can in this way subvert the functioning of the network. It seems appropriate to model this aspect of a nodes strategic behaviour: whether it connects to strangers, therefore enabling the sybil attack, or only to friends, making network infiltration harder (or impossible).

As a result we allow nodes to chose amongst two strategies: either they only connect to friends, or they splits their link budget in half between friends and strangers. It is clear why nodes have incentives to talk to friends, since it increases their utility directly by decreasing the length of the path to those friends. On the other hand, given the limited link budget, nodes cannot directly connect to all their friends, and may find it beneficial to relay their communications through a stranger that is closer to two or more friends of theirs.

Given a strategy, amongst the two available, we use a deterministic network formation algorithm. We select pseudo-randomly a candidate link amongst all possible links in the network, and offer each of them in order to the two concerned nodes that would become linked. If the utility of both nodes increases or remains stable, the link is accepted and becomes part of the network. Otherwise the link is rejected. Nodes can of course only accept links as their link budget permits. For this reason when they have already spent their budget they consider the new link under the assumption that they will have to give up a (pseudo-random) existing link. We borrow this (rather myopic) strategy from [5].

At this stage we are only left with defining a sybil adversary, that would no doubt attempt to infiltrate the network, by providing shorter paths between friendly nodes. Before doing this we seek to characterise the networks resulting from our model so far, without an adversary.

2.2 Analysis, experimental results and limitations

We looked for the Nash equilibrium of the simple game without an adversary both analytically and by simulation. A Nash equilibrium is a set of strategies, where each player has no incentive to change their strategy if they assume that other players will also not change their strategy.

Our second interesting remark is that we can find a Nash equilibrium analytically.

Remark 2. *The set of strategies where each player only connects to friends is a Nash equilibrium.*

Proof. Assume that a player deviates from the strategy to connect only to friends, and splits its link budget between connecting to friends and connecting to strangers. This node will find no stranger that would be willing to connect to him, since all other nodes only connect to friends. As a result its utility would

at best be as good as if it was only connecting to friends. We conclude that the all-friends strategy is a Nash equilibrium. \square

In the case all nodes only connect to friends it is rather difficult to introduce an exogenous sybil attacks. The full link graph will be a subset of the friendship graph, leaving an exogenous adversary, seen by nodes as a stranger, little opportunity to connect.

It would therefore be interesting to find other equilibria, that include some nodes that find it beneficial to connect to strangers. We used two models to simulate network formation, and attempt to find another Nash equilibrium:

- **Random Model.** We considered a random friendship graph, where all nodes initiate a set number F of friendships with other random nodes in the network. All nodes have the same link budget $L < F$.
- **Unbalanced Random Model.** The friendship graph is as before, but one node has a very large link budget ($L_0 > 2 \cdot F$).

Using the random model we have failed to consistently find another Nash equilibrium aside from the expected all-friends set of strategies. (Often a Nash equilibrium appears because of the particularities of the schedule in which links are offered in the Network formation stage. The disappear as soon as a different or longer schedule is offered.) We conjecture that nodes always have incentives to spend their limited link budget to connect to friends, and use friends to relay communications to other friends.

The unbalanced random model provides us with further insights. If the link budget of most nodes is comparable with their number of friends (while still lesser) most still choose to only connect to friends. The intuition behind this is that in a random graphs shortest paths will be $\mathcal{O}(\log \mathcal{N})$, and the probability a friend or a stranger is closer to another friend or stranger is roughly equal. As a result nodes will prefer to get access to other friends by connecting to their friends rather than strangers. On average this provides better utility.

A special case of the unbalanced random model is worthy of attention, when most nodes are only allowed one link $L_n = 0, n \in N$. In this case the Nash equilibrium is a star topology, with the link rich node at its centre. The nodes that are natural friends of the rich node chose to only connect to their rich friend, while others chose the mixed strategy (that means in this case that they can connect to the rich stranger). It is important to note that the rich node is indifferent about his strategy since in both cases he can connect directly to his friends reaching the same utility.

In case nodes have a rather small link budget ($L=2$) we start seeing interesting topologies emerging, where some nodes choose to connect to the central hub, but others prefer to connect directly to others. Such an Equilibrium is represented in figure 1.

3 A Roadmap for Future work

In this position paper we have discussed a simple game for network formation, where nodes can chose between relaying information over only friends or also strangers. The second strategy open the way to an exogenous adversary that pretends to shorten paths in order to capture the nodes' links, and then disrupt




Figure 1: Nash Equilibrium of unbalanced random model with parameters $N = 20, F_i = 2, L_i = 2, (L_0 = 20)$.

or observe their actions. The next key step in this research is to formalise and introduce such an adversary, and observe the effect they will have on the nodes' choice of strategies.

As second important open problem is mapping our simple model to reality in a more convincing way. In particular real work networks are far from random: nodes want to talk to clusters of other nodes, and both friend ship and link capacity is distributed according to a power law. Furthermore what matters in real world networks is not a short path merely existing, but also being able to find it within some reasonable time (social networks in particular are navigable [7]). Modifying the friendship graph and link budgets to fulfil those requirements is an important next step.

Finally we have mostly considered routing security as the security property, i.e. the inability of an attacker to control the route honest nodes' messages take in the network. This is also closely related to surveillance or other cryptographic attacks, where an adversary node seeks to become a man-in-the-middle. The holly grail of such research would of course be a framework in which proposed and new reputation systems can be evaluated given nodes and attackers with various objectives.

Despite the presented line of research being in its infancy we still can provide some interesting insights for designers of peer-to-peer systems. The first is that in the absence of any restriction in the number of links it is safer to connect directly to whoever a nodes needs to talk to. Special classes of security properties such as anonymity do not permit this. Yet connecting to strangers opens you to exogenous not just endogenous attackers.

The second key finding is that given a limited, but not tiny, link budget only talking to friends is a Nash Equilibrium, and also only equilibrium strategy we have found. As a lesson it is therefore important to use friends as much as possible as the infrastructures to route information on peer-to-peer networks. The difficulty we had to define models in which it is rational to talk to strangers is

in sharp contrast with the established peer-to-peer [8] paradigms, that relay exclusively on strangers to route. Re-aligning those systems to make better use of high level relations and needs amongst nodes would probably also automatically increase their routing security.

References

- [1] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117, 1998.
- [2] Alice Cheng and Eric Friedman. Sybilproof reputation mechanisms. In *P2PECON '05: Proceeding of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, pages 128–132. ACM Press, 2005.
- [3] John R. Douceur. The sybil attack. In Peter Druschel, M. Frans Kaashoek, and Antony I. T. Rowstron, editors, *IPTPS*, volume 2429 of *Lecture Notes in Computer Science*, pages 251–260. Springer, 2002.
- [4] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. *Selected Areas in Communications, IEEE Journal on*, 24(5):1010–1019, 2006.
- [5] M.O. Jackson. A Survey of Models of Network Formation: Stability and Efficiency. *Group Formation in Economics: Networks, Clubs and Coalitions*, 2003.
- [6] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks.
- [7] Jon M. Kleinberg. The small-world phenomenon: an algorithm perspective. In *STOC*, pages 163–170, 2000.
- [8] Ion Stoica, Robert Morris, David Liben-Nowell, David R. Karger, M. Frans Kaashoek, Frank Dabek, and Hari Balakrishnan. Chord: a scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Trans. Netw.*, 11(1):17–32, 2003.