The effects of observation errors on the attack vulnerability of complex networks

Lashon B. Booker

The MITRE Corporation, 7515 Colshire Drive, McLean, VA 22102
booker@mitre.org

Abstract. Identifying the key nodes to target in a social network is an important problem in several application areas, including the disruption of terrorist networks and the crafting of effective immunization strategies. One important issue that has received limited attention is how such targeting strategies are affected by erroneous data about network structure, This paper describes simulation experiments which investigate that issue.

Keywords: Centrality measures, robustness, network attack

1 Introduction

One of the most important applications of social network analysis is the key player problem (Borgatti, 2003). Identifying the critical nodes in a network is often the first step in understanding who to target in order to disrupt a terrorist network or who to immunize in order to halt the spread of an epidemic. The most critical issue in this regard is determining how vulnerable the network structure is to the removal of nodes and links.

The importance of network topology in determining the vulnerability of networks to errors and attacks was first pointed out by Albert et al. (Albert, et al., 2000). They studied two network models – an exponential network and a scale-free network – and observed how the characteristic path length and the size of the largest cluster changes as nodes are removed. When nodes are removed randomly, as in the case of node failure or errors, the exponential network becomes more fragmented as errors increase while the scale-free network is only minimally affected. This error tolerance in scale-free networks is due to the fact that random node removal is most likely to take out nodes with small connectivity, thereby having a small impact on the network topology. On the other hand, when node removal is done to deliberately inflict damage by deterministically removing the highest degree nodes, there is a drastically different outcome. The scale-free network quickly becomes fragmented, while the exponential network shows much more resistance to fragmentation.

Subsequent research has extended our understanding of how complex networks respond to attacks and failures. Holme et al. (Holme, et al., 2002) considered a variety of strategies for attacking both nodes and edges. They used betweenness centrality values as well as node degree to select which nodes to remove, and also varied when those values were calculated: either once in the initial network as done by Albert et al., or repeatedly as each node is removed. Crucitti et al. (Crucitti, et al., 2003) examined scale-free graphs with high clustering properties and measured both global and local aspects of the network response to an attack. Gallos et al. (Gallos, et al., 2005) studied situations where nodes are removed based on a probability distribution that depends on node degree, rather than a deterministic strategy that removes the highest degree nodes. Wu et al. (Wu, et al., 2007) provided a theoretical analysis of intentional attacks in scale-free networks, focusing on scenarios involving incomplete information.

One issue that has not been addressed in previous research on network attacks is the impact of erroneous information. In real applications, network data is likely to include incomplete or mistaken data about the network structure. This paper addresses that issue by using simulation experiments to investigate the impact of errors on the effectiveness of network attacks.

2 Methodology

The simulation experiments used in this research were structured along the lines of previous studies that examined the effects of noisy observations on the measurement of properties of social networks (Kossinets, 2003) (Borgatti, et al., 2006). The experimental procedure starts with a collection of randomly generated synthetic graphs which are assumed to be complete. These graphs represent the "true" networks to be attacked. For each true network, observation errors are introduced in a controlled manner (nodes/edges are added/removed) to generate a corresponding "observed" network.

In this study, we add an additional step to model network attacks. Two attack strategies are considered: a random attack where nodes are selected for removal randomly[1]; and, a targeted attack where nodes are selected for removal deterministically based on centrality scores. The centrality scores, and a list of nodes to remove, are computed in the observed network. Since we are interested in the effects of the attack on the true network, we do not execute the attack on the observed network or make any measurements comparing properties in the true and observed networks. Instead, the list of targeted nodes is deleted from the true network where the effect of the attack is measured.

In more detail, to construct a true network we select a topology (erdos-renyi (Erdos & Renyi, 1959), scale-free (Barabási & Albert, 1999), small world (Watts & Strogatz, 1998), or scale-free communities (Lancichinetti & Fortunato, 2009)), a size (25, 50 or 100 nodes) and a density (0.01, 0.02, 0.05, 0.1, 0.3, 0.5). To construct a corresponding observed network, we introduce exactly one of four types of error into a true network: node deletion, node addition, edge deletion, and edge addition. Errors are introduced randomly at a selected rate (0.0, 0.01, 0.05, 0.1, 0.25, 0.5). Network attacks are mod-

[1] Random node removals are sometimes viewed as network "errors" or "failures", but since our experiments involve observation errors, we will call these random attacks to avoid confusion.

eled by selecting an attack type (random or targeted), a node removal rate (0.0, 0.01, 0.02, 0.03, 0.04, 0.05, 0.1, 0.3, 0.5, 0.7, 0.9) that determines the attack intensity, and, for targeted attacks, a centrality measure (degree, betweenness, closeness, or eigenvector) (Borgatti, et al., 2006) for identifying which nodes to remove.

3 Attack Vulnerability Experiments

For each combination of network topology, network size, network density, error type, error level, attack type, and node removal rate we generate 1,000 pairs of synthetic networks. For each network pair, the observed network is used to plan an attack and the targeted nodes are removed from the true network. The effect of the attack on the true network is assessed by measuring the relative size of the largest connected component (Motter & Lai, 2002). This performance metric is normalized with respect to values in the true network before the attack, so it begins with value 1.0 and decreases as the attack effectiveness increases. The remainder of this section discusses the results of these experiments.

3.1 Attacks on Error-Free Networks

A look at the error-free case, where the data about network nodes and edges is accurate, shows responses to random attacks and targeted attacks that are consistent with previously published results. **Fig. 1** compares attack effectiveness on networks having 100 nodes and 5% density, using betweenness centrality to identify which nodes to target for deletion. Targeted attacks are much more effective than random attacks for fragmenting graphs of all types, though when the removal rates are small there is little difference between the two attack strategies. As expected, the largest advantage for targeted attacks is seen in scale-free networks. The difference in performance between random and targeted attacks typically becomes significant when the attack removes at least 10% of the nodes. The exception is small world networks, where the high clustering seems to make these networks more resistant to both kinds of attacks. We observed similar results when using degree centrality and closeness centrality to identify nodes to target for deletion. Eigenvector centrality tends to provide a much smaller advantage for targeted attacks over random attacks.

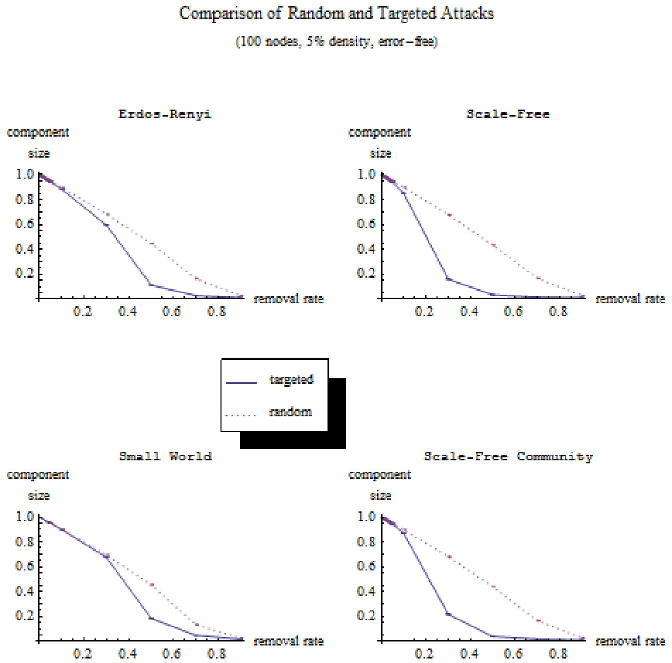

Fig. 1. Comparison of random and targeted attacks, using betweenness to select targeted nodes. This graph, like all subsequent ones, includes error bars which are almost always smaller than the plot symbols

3.2 Attacks on Networks with Observation Errors

The experiments studying the impact of observation errors on the effectiveness of network attacks also looked at graphs of size 100 with 5% density and used betweenness centrality to select targeted nodes. The expected outcome was that observation errors would reduce the effectiveness of attacks, but the impact would resemble a noise effect that increased as the amount of error increased. It turns out that the type of error makes a significant difference in the impact on attack effectiveness, and the impact does not always vary with the amount of error in the manner expected.

For random attacks, the connectivity in the graph has no bearing on which nodes are selected for deletion. Consequently, edge addition errors and edge deletion errors have no impact on the effectiveness of random attacks. Node addition errors have no impact on random attacks either, but for a different reason. Adding nodes leads to a larger observed network, which makes key nodes in the true network less likely to be selected at random for deletion. However, since attacks are defined in terms of the fraction of nodes removed, a larger observed network also means more nodes are selected for deletion and therefore the chance of selecting any particular true node is increased. The net result of these two opposing selection pressures is that node addition errors do not noticeably impact the effectiveness of random attacks.

Node deletion errors, on the other hand, have a very noticeable impact on effectiveness (see **Fig. 2**). The reason for this impact is that if a key node does not appear in the observed network, the attack has no chance to delete it. The higher the likelihood that such an event occurs, the less effective the attack can be.

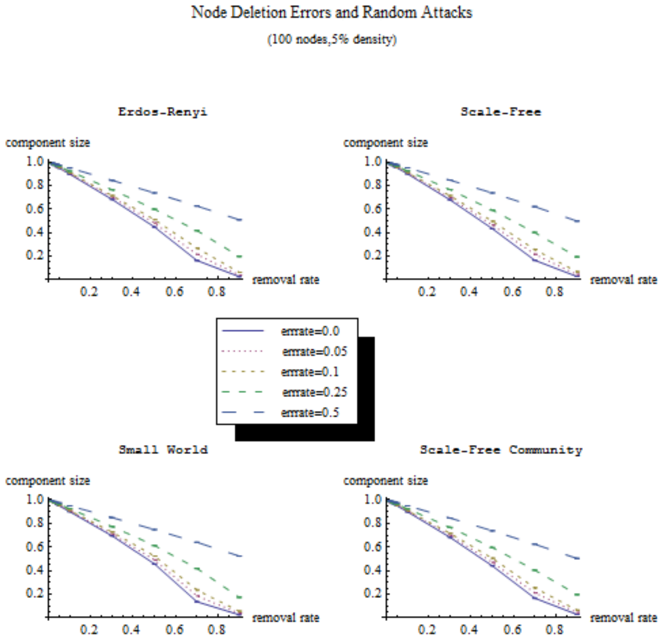

Fig. 2. The effect of node deletion errors on random attacks

For targeted attacks, connectivity has a big influence on the centrality values used to select which nodes to remove. This means that edge addition errors and edge deletion errors have a noticeable impact on targeted attacks (see **Fig. 3**). As the amount of edge error increases, there is a corresponding decrease in the effectiveness of targeted attacks. When the amount of edge error is sufficiently large, the centrality values become so inaccurate that the targeted attacks become indistinguishable from random attacks. As the fraction of nodes removed by the attack gets large and the key nodes are more certain to be targeted, the impact on attack effectiveness begins to look the same regardless of the amount of edge error.

The effect of node errors on targeted attacks is much different. As the amount of node deletion error increases, there is a steady decrease in attack effectiveness as shown in **Fig. 4**. This degradation continues even as the fraction of nodes removed by the attack gets large since, as noted earlier, if a key node does not appear in the observed network then the attack cannot remove it[2].

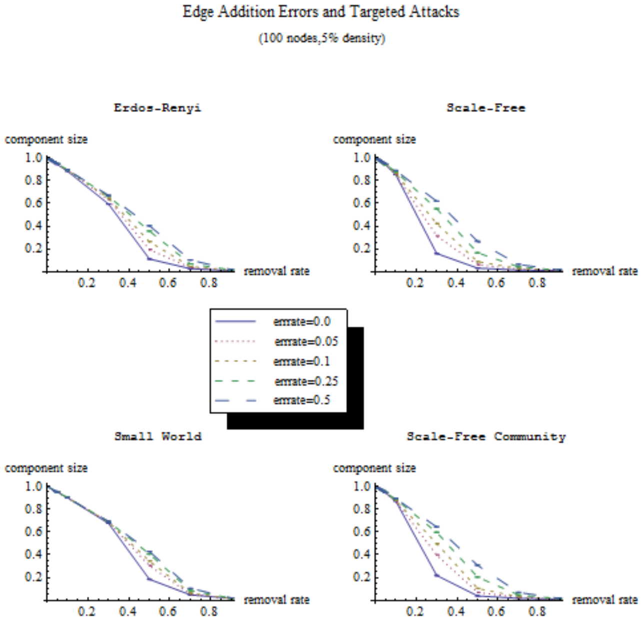

Fig. 3. The effect of edge addition errors on targeted attacks

The effect of node addition errors is more complicated (see **Fig. 5**). Node addition error has the smallest impact on the accuracy of centrality values (Borgatti, et al., 2006), so there is only a small change in attack effectiveness as the amount of error increases. The nature of that impact appears to involve another interaction between two tendencies as noted for node addition errors and random attacks, but with a different dynamic. As the fraction of nodes removed by the attack gets larger, the key nodes in the true network are more likely to be included in the top ranked observed nodes. This will increase attack effectiveness early on, since the number of nodes targeted for deletion will be greater than the number chosen in the error-free case, effectively increasing the likelihood that key nodes will be targeted. Eventually, though, the larger number of targeted observed nodes will include a larger proportion of erroneous nodes. Targeting these nodes is a waste of effort since they are not elements of the true network. This diversion of the attack focus from true nodes tends to decrease attack effectiveness.

[2] Note that the behavior does not converge to something resembling a random attack, in stark contrast to the behavior observed for edge errors. The effectiveness is eventually worse than a random attack.

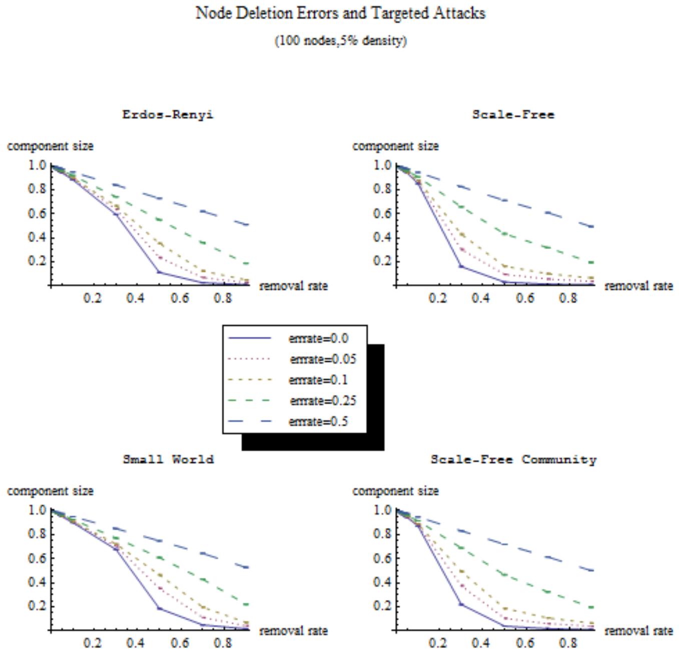

Fig. 4. The effect of node deletion errors on targeted attacks

4 Summary and Conclusions

Different kinds of observation errors have differing impact on the effectiveness of random and targeted attacks. In most cases, when the fraction of nodes targeted for removal is small, the differences in the impact of the various error types is minimal. The one exception is node deletion errors. These errors can have a significant impact even in small scale attacks. The unavoidable fact is that when a key node in the true network is never observed, an attack cannot remove it. Consequently, node deletion is the source of error that has the most disruptive impact on the effectiveness of network attacks.

Acknowledgements The author gratefully acknowledges support for this research from the Joint IED Defeat Organization.

References

1. Albert, R., Jeong, H. & Barabási, A.-L., 2000. Error and attack tolerance of complex networks. *Nature*, 27 July, Volume 409, pp. 378-381.
2. Barabási, A.-L. & Albert, R., 1999. Emergence of scaling in random networks. *Science*, 286(5439), pp. 509-512.

3. Borgatti, S. P., 2003. The key player problem. In: R. Breiger, K. Carley & P. Pattison, eds. *Dynamic Social Network Modeling and Analysis: Workshop Summary and Papers.* s.l.:National Academies Press, pp. 241-252.
4. Borgatti, S. P., Carley, K. M. & Krackhardt, D., 2006. On the robustness of centrality measures under conditions of imperfect data. *Social Networks,* Volume 28, pp. 124-136.
5. Crucitti, P., Latora, V., Marchiori, M. & Rapisarda, A., 2003. Efficiency of Scale-Free Networks: Error and Attack Tolerance. *Physica A,* Volume 320, pp. 622-642.
6. Erdos, P. & Renyi, A., 1959. On Random Graphs, I. *Publicationes Mathematicae (Debrecen),* Volume 6, pp. 290-297.
7. Gallos, L. K. et al., 2005. Stability and topology of scale-free networks under attack and defense strategies. *Physical Review Letters,* May, 94(18), p. 188701.
8. Holme, P., Kim, B. J., Yoon, C. N. & Han, S. K., 2002. Attack vulnerability of complex networks. *Physical Review E, Statistical Nonlinear and Soft Matter Physics,* 7 May, 65(5 pt. 2), p. 056109.
9. Kossinets, G., 2003. Effects of missing data in social networks, s.l.: e-print, arXiv.
10. Lancichinetti, A. & Fortunato, S., 2009. Benchmarks for testing community detection algorithms on directed and weighted graphs with overlapping communities. *Physical Review E,* 80(1), p. 016118.
11. Motter, A. E. & Lai, Y.-C., 2002. Cascade-based attacks on complex networks. *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics,* 66(6), p. 065102(R).
12. Watts, D. J. & Strogatz, S. H., 1998. Collective dynamics of 'small-world' networks. *Nature,* 4 June, Volume 393, pp. 440-442.
13. Wu, J., Deng, H.-Z., Tan, Y.-J. & Zhu, D. Z., 2007. Vulnerability of complex networks under intentional attack with incomplete information. *Journal of Physics A: Mathematical and Theoretical,* Volume 40, pp. 2665-2671.

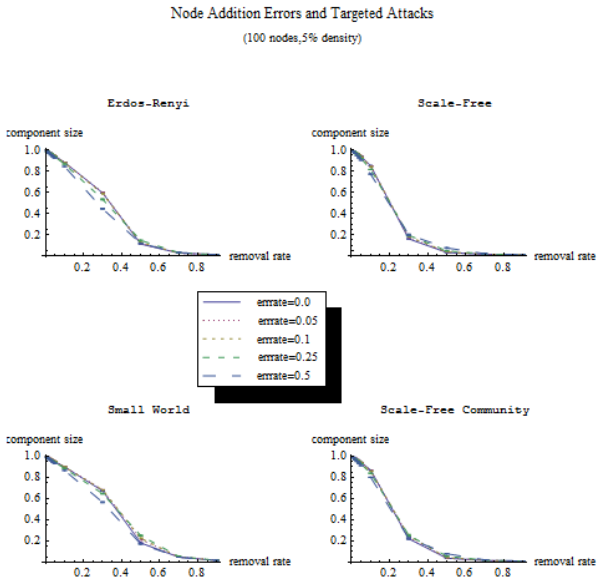

Fig. 5. The effect of node addition errors on targeted attacks