Scopus

## Documents

Arnold, H., Masad, D., Pagani, G.A., Schmidt, J., Stepanova, E.

**Network disruption and recovery: Co-evolution of defender and attacker in a dynamic game**

(2014) *Studies in Computational Intelligence*, 549, pp. 141-155.

**DOI:** 10.1007/978-3-319-05401-8_14

[a] University of Oregon, Eugene, OR, United States
[b] George Mason University, Fairfax, VA, United States
[c] University of Groningen, Groningen, Netherlands
[d] University of Natural Resources and Life Sciences, Vienna, Austria
[e] Santa Anna School of Advanced Studies, Pisa, Italy

**Abstract**
The evolution of interactions between individuals or organizations are a central theme of complexity research. We aim at modeling a dynamic game on a network where an attacker and a defender compete in disrupting and reconnecting a network. The choices of how to attack and defend the network are governed by a Genetic Algorithm (GA) which is used to dynamically choose among a set of available strategies. Our analysis shows that the choice of strategy is particularly important if the resources available to the defender are slightly higher than the attackers'. The best strategies found through GAs by the attackers and defenders are based on betweenness centrality. Our results agree with previous literature assessing strategies for network attack and defense in a static context. However, our paper is one of the first ones to show how a GA approach can be applied in a dynamic game on a network. This research provides a starting-point to further explore strategies as we currently apply a limited set of strategies only. © 2014 Springer International Publishing Switzerland.

**Publisher:** Springer Verlag

**ISSN:** 1860949X
**ISBN:** 9783319054001
**Language of Original Document:** English
**Abbreviated Source Title:** Stud. Comput. Intell.
**Document Type:** Conference Paper
**Source:** Scopus

# Network Disruption and Recovery: Co-Evolution of Defender and Attacker in a Dynamic Game*

Holly Arnold[1], David Masad[2], Giuliano Andrea Pagani[3], Johannes Schmidt[4], and Elena Stepanova[5]

[1] University of Oregon, Eugene, Oregon, USA
email:arnold3@uoregon.edu
[2] George Mason University, Fairfax, Virginia, USA
email: dmasad@gmu.edu
[3] University of Groningen, Groningen, The Netherlands
email: g.a.pagani@rug.nl
[4] University of Natural Resources and Life Sciences, Vienna, Austria
email: johannes.schmidt@boku.ac.at
[5] Santa Anna School of Advanced Studies, Pisa, Italy
email: e.stepanova@sssup.it

**Abstract.** The evolution of interactions between individuals or organizations are a central theme of complexity research. We aim at modeling a dynamic game on a network where an attacker and a defender compete in disrupting and reconnecting a network. The choices of how to attack and defend the network are governed by a Genetic Algorithm (GA) which is used to dynamically choose among a set of available strategies. Our analysis shows that the choice of strategy is particularly important if the resources available to the defender are slightly higher than the attackers'. The best strategies found through GAs by the attackers and defenders are based on betweenness centrality. Our results agree with previous literature assessing strategies for network attack and defense in a static context. However, our paper is one of the first ones to show how a GA approach can be applied in a dynamic game on a network. This research provides a starting-point to further explore strategies as we currently apply a limited set of strategies only.

## 1 Introduction

Networks have been used to elegantly model systems with many interacting elements in many different disciplines [16] including biology [10], linguistics and social sciences [18], epidemics [4], infrastructures [17], and banking [3]. A central question in network science is to understand the robustness of a network if nodes or edges fail or come under attack [9,2]. The study of network robustness has many different applications, such as assessing the vulnerability of power grids [1], subway networks [13], and airline transportation networks [7]. Additionally, social networks of interest are covert networks such as criminal or terrorist

* The authors have evenly contributed to the work presented in the paper.

organizations [12]. For example, targeting one individual over another by police force might have more effect on the communication capability of the network depending on network topology. Analogously, technical networks of interest are computer networks, where the maintainers of computer networks might attempt to identify the best strategy to defend against cyber attacks or random failures.

Network topology plays a large role in how effective an attack is, and how the network is able to defend itself. Albert *et al.* [2] demonstrated that scale-free networks, unlike random networks, are very robust to random failure but vulnerable to targeted attacks. This is due to the fact that most nodes in their scale-free model had few connections, so the probability of randomly targeting a highly connected and central node was low. The targeted attack, however, was able to remove the small percentage of highly connected nodes rapidly, thereby crippling the network connectivity much faster than random attacks. Several researchers have addressed the issue of network robustness using iterative attack and defense games on networks where attackers and the defending network employ static attack and defense strategies against one another [14,8,5]. Holme *et al.* [8] considered static attack strategies on edges as opposed to nodes, and suggested edge betweenness as a more effective target of an attacker than attacks on high degree nodes. Nagaraja and Anderson [14] extend Holme's approach by considering both static attack and defense strategies. The network is allowed to defend, or rewire its connections to become less vulnerable to attack via a set of predefined defense strategies. Likewise attacks on the network are performed with a predefined strategy, where attacks based on node centrality were found to perform best on disconnecting the network. Like Nagaraja and Anderson, Domingo-Ferrer *et al.* [5] allow for iterated attack and defense rounds, and show that the attacker's knowledge of the network is also an important factor in the effectiveness of an attack.

But while previous literature on iterated attack and defense has considered many different attack and defense strategies, to date, no research has been done to allow the attacker (or defender) to dynamically change strategies during the course of the game. We extend previous approaches by allowing the attacker and defender to operate with a set of strategies in each time step and to make decisions based on mixing strategies. This allows not only for the possibility that a single strategy could go to fixation, but also cyclical pattern of attack and defense strategies to emerge. A second possibility is that it could simply be advantageous to attack (or defend) based on mixing strategies during attack and defense rounds. Or, it could be that attack and defense strategies simply reach an equilibrium, where no further improvement of attack (or defense) strategy is found by the participants. We examine attacker strategies which identify network nodes to maximize the damage to the network defender. Contrary, network defenders identify the best way to rewire the network following the attack. The choice of strategies is dynamically determined by a genetic algorithm (GA) for both attackers and defenders, and thus representing coevolution between attacker and defender, or a coevolutionary 'game'.

## 2 The Model

In our model we have three fundamental entities that we deal with:

1. A **network** composed by a set of $n$ nodes and $m$ edges.
2. An **attacker** attempting to disrupt the network.
3. A **defender** attempting to repair the network after an attack to guarantee its continuing functionality.

An attacker disrupts the network by removing a node and all its associated edges. The defender, on the other hand, is allowed to reintroduce a node that has been previously disconnected as a consequence of an attack by re-connecting it to the network. The defender also adds edges to the network if he has enough resources to spend. In fact, the attacker and defender each have an assigned set of resources that they can use in their attack or defense process. The resources for the attacker correspond to the number of nodes that he can remove, whereas defender resources correspond to the number of edges that can be added to the network following an attack. We assume that attackers and defenders have complete knowledge of the network topology and that they perform their actions one after the other beginning with an attack followed by a defense.

A particular simulation starts by generating an initial (first generation) population of an equal number of attackers and defenders. Their genes are initialized randomly, and attackers and defenders are randomly paired up. Each attacker-defender pair is assigned a network of $n$ vertices and no edges. Based on the rules defined by their genomes (which are explained in detail in section 3), each defender adds new edges to the network, up to a total number of $m$ edges. So we start with a set of disconnected nodes and start to build the network from scratch, not fixing any specific network topology at the start. However, fixing the defender and attacker rules will create networks that are similar in topology.

After the network is initially built, the attacker removes $k$ nodes in the network, $k$ being the amount of resources assigned to the attacker, which are the same for all attackers. The choice of the nodes to remove depends on the attacker genome. Once the attack phase is completed, the defender is allowed to add a total of $w$ edges to the network, $w$ being the amount of resources assigned to the defender, which are the same for all defenders. First, the nodes removed by the attacker in this round are re-connected to the network. The nodes to which they will be connected depends on the defender genome. If defender resources allow additional edges to be inserted into the network, those edges are added to the network by the following rule: the starting point for the edges is a random node from the list of nodes which lost edges in the previous attack. The end point is determined by the genetic algorithm. If there are still resources left after reconnecting each of the nodes that have lost an edge in the previous attack, random nodes in the network are picked as starting points. Again, the end points of the new edges are determined by the genetic algorithm.

This process of attack and defense on the network is repeated for $r$ rounds. In summary, a round is an execution of the game with iterative attacks each based on the $k$ resources for the attacker and a (re-)wiring process consisting of

$w$ resources for the defender. In our simulations $r$ is equal to 20, i.e. a total of 20 attack-defense rounds is played in each generation of the genetic algorithm.

After each round, the fitness (see Section 3 for the thorough fitness description) of the attackers and defenders is calculated and a final average fitness after $r$ rounds is computed for each individual in the population. Recombination of individuals and mutations which are necessary to generate a new generation of attackers and defenders are discussed in the next section. We are interested to track over generations the evolution of the fitness function for both, attacker and defender as a measurement of their performance in the game. We track over generations the change in genomes as well, because we are interested to identify prevailing strategies.

## 3 Genetic Algorithm

The GA is used to evolve the strategies applied by the attackers and defenders and thus, allows for a dynamic development of the strategies that are applied by the two groups. A strategy is a mechanism for both the attacker and the defender to decide which node to attack or edge to create/rewire based on some rules, measures or indicators on the network. First, we define the fitness function, then we discuss the genomes of attackers and defenders, and finally we present recombination and mutation strategies.

### 3.1 The fitness function

We define the fitness of the defender to be the number of nodes of the Largest Connected Component ($LCC$) divided by the total number of nodes in the initial network $n$, i.e.

$$f^{def} = \frac{LCC}{n} \tag{1}$$

The attacker's fitness is the opposite, i.e.

$$f^{att} = 1 - f^{def} \tag{2}$$

The size of the LCC is a good proxy of the resilience of the network, its ability to keep its structure connected and thus allow interaction between the nodes. The same metric has been used in previous studies [11,15], allowing our results to be compared to previously-published ones. However, depending on the application of our model, different fitness functions may be appropriate. In section 6 we discuss this aspect in more detail.

### 3.2 Attacker genome

A set of strategies is available to the attacker indexed by $j = \{1, 2, 3\}$ – these strategies have been developed previously in the literature [11,15,5]:

1. High-degree removal: nodes are prioritized for removal in decreasing order with respect to their degree.
2. High-centrality removal: nodes are prioritized for removal in decreasing order with respect to their betweenness centrality, which is known to be more related to connectivity than other centrality measures.
3. Random removal: nodes are prioritized randomly.

Each gene $G_j$ corresponds to a weight on one of the strategies, and its value varies from 0 to 100. Each strategy calculates a specific network metric (e.g. degree or betweenness centrality) for every node $i$. The metric is normalized to the interval $[0, 1]$. Thus, to each node $i$ in the network, a value $N_{ij}$ in the interval $[0, 1]$ is assigned by each strategy. In combination with the importance of the strategy as defined by the genome, this represents the removal ranking of a node $i$. For each node in the network, the attacker's genome assigns a number

$$TotalN_i = \sum_j G_j N_{ij} \tag{3}$$

which is a linear combination of all available strategies weighted by the attacker genome. The probability of a node $i$ to be attacked $Pr_i$ is $TotalN_i$ divided by the sum over $TotalN_i$ for all network nodes, i.e.

$$Pr_i = \frac{TotalN_i}{\sum_i TotalN_i} \tag{4}$$

A node is removed from the network based on its probability $Pr_i$.

### 3.3 Defender Genome

The strategies of the defender are similar to the attacker strategies as they are based on the same weighting algorithm. The starting point of an edge that is added to the network is not determined by this weighting algorithm, but by a sequence of rules as outlined in the previous section. Only the endpoint of the new edge is determined by the defender's genome.

The following strategies are available to the defender indexed by $j = \{1, 2, 3\}$ - these strategies have been developed previously in the literature [11,15,5]:

1. Preferential replenishment: nodes are ranked in decreasing order with respect to their degree.
2. Balanced replenishment: nodes are ranked in increasing order with respect to their betweenness centrality.
3. Random replenishment: nodes are ranked randomly.

The weighting of nodes is performed similar to the attacker, i.e. the genome determines how the value of a certain metric for the nodes is weighted. See the description of the attacker genome above for details.

### 3.4 Genome reproduction process

The indexed set of genes $G_j$, $j = \{1, 2, 3\}$ representing the attacker and the defender genome are initially randomly sampled from a uniform distribution in the range $[0, 100]$. Reproduction consists of gene recombination: two attackers or defenders from the current population are randomly chosen from the current generation. The mechanism of selection follows the principle of genetic algorithms known as *roulette wheel selection* [6]: the probability of being picked is not uniform, but is proportional to the fitness of the agent. A random position in the genome is chosen for crossover. At this position, the two individuals will exchange their genetic material, taking the first part from the first parent and the second part from the second parent[6], as shown in Figure 1. The offspring replaces the previous generation (i.e., parents), thus providing the new base of the genetic material for the following evolution step.

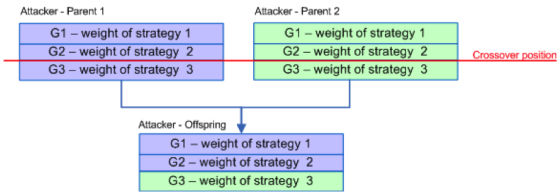


**Fig. 1.** Example gene crossover

A mutation process occurs with a fixed 5% probability. The mutation in a gene is obtained by sampling a value from a Gaussian distribution with the mean equal to the current value of the gene and a standard deviation of 5.

## 4 Scenarios

We are interested in the following research problems: first, how does an attacker applying a genetic algorithm perform against a static defender, i.e. a defender with only one, fixed defense strategy. We next look at the inverted scenario, i.e. how a static attacker performs against an evolving defender. Finally, we allow both the attacker and defender to co-evolve against each other. For the purpose of comparison, we also run each static attacker strategy against each static defender strategy. Both defender and attacker have 3 different strategies each. This implies that there are 16 different scenarios to assess in total.

[6] As we have only 3 genes in the genome, there are only two possibilities: the offspring will inherit the first gene from his first parent and second and third genes from his second parent, or he will inherit the two first genes from the first parent and the third gene from the second parent

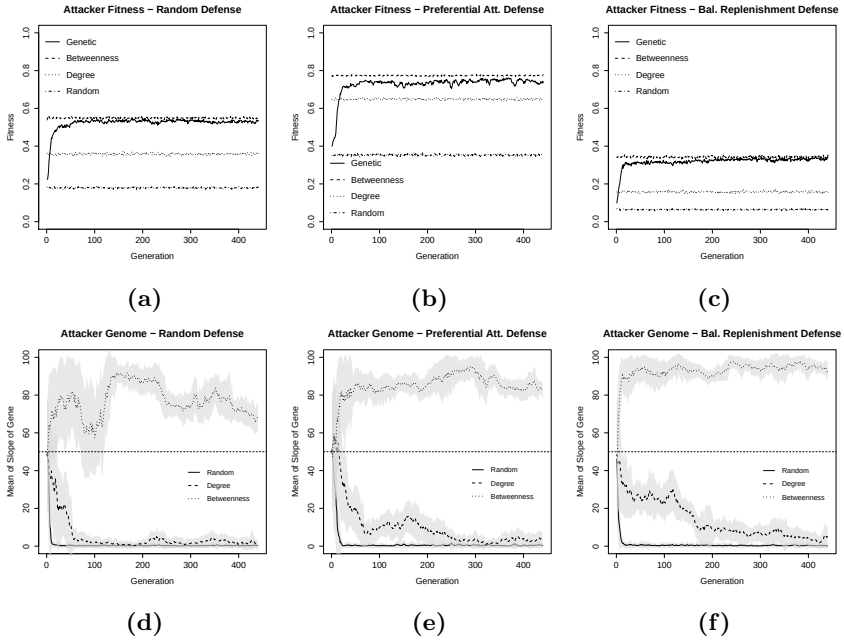


**Fig. 2.** Top: evolution of the mean of fitness in the attacker population when attackers use the genetic algorithm against 3 static strategies. Bottom: Evolution of the mean of attacker weights for different strategies in the genetic case. The transparent areas indicate the standard deviation. Left: Attacker vs. Random defender. Middle: Attacker vs. Preferential defender. Right: Attacker vs. Balanced Replenishment.

In the base run, we start with a population of 200 attackers and defenders, operating on a network of 100 nodes and 150 edges, and run the GA for 500 generations. Attackers are allowed to remove 3 nodes while defenders rewire 5 edges. In a sensitivity analysis we test different defender budgets of 3,7, or 9 edges. The whole simulation is driven by random choices of attackers and defenders and by a random (although directed) process of selection of individuals in the genetic algorithm. That implies, that a different run of the same simulation may show a different dynamical outcome. At the current moment, we did not run the simulations for several times to analyze the variance of results due to time constraints with the exception of the co-evolution case which was run 25 times. Further runs are left to be presented in future versions of this paper.

## 5 Results

### 5.1 Scenarios Results

**Static Defenders** Figure 2a shows that the dynamic attacker quickly approaches the fitness of the single best attacker strategy against a static random

defender. The genes evolve accordingly (Figure 2c) , prioritizing high weights for the betweenness strategy and much lower weights for the other two strategies. It can also be observed that the standard deviation in the genes decreases over time, indicating that the individuals in the population converge. Playing against the other two static defender strategies show similar results (Figures 2e and 2f). The worst static defense strategy is preferential attachment which can be derived from the fact that the attacker fitness is highest in that case (middle in Figure 2b). The best possible static defense strategy is balanced replenishment as indicated by the low attacker fitness (Figure 2c). In all cases, the betweenness attack strategy is selected by the attacker's GA.

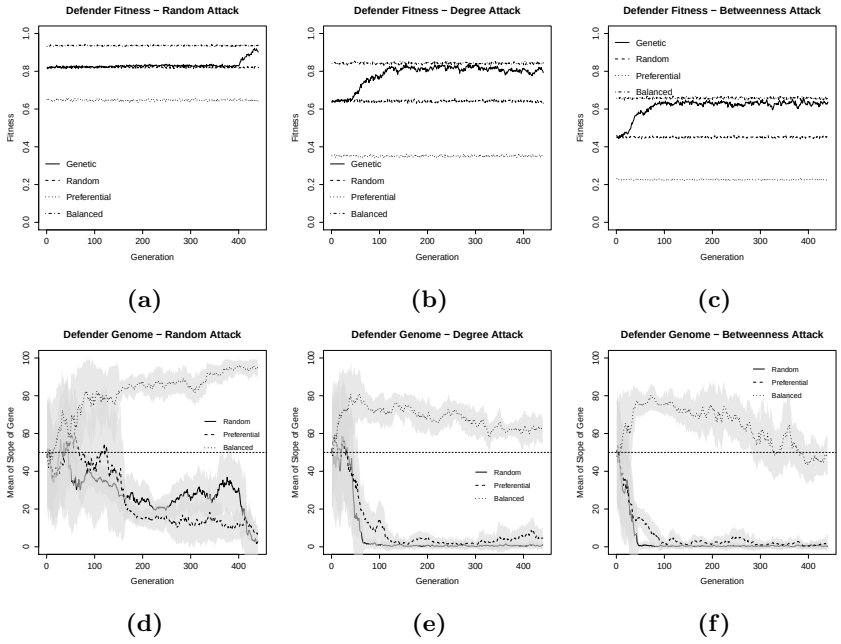


**Fig. 3.** Results of simulation runs: Defenders applying the genetic algorithm against 3 static attack strategies. Top: Mean of fitness of defender. Bottom: Mean of defender genes. The transparent areas indicate the standard deviation. Left: Random attack vs. Defender. Middle: Degree attack vs. Defender. Right: Betweenness attack vs. Defender.

**Static Attackers** Also the defender has a preferred strategy, independent of the static attacker strategy. It is balanced replenishment. However, the GA takes more time to find the dominating strategy in comparison to the attacker's GA in some cases. Defending against a random attacker (Figure 3a) shows that the defender's fitness approaches the fitness of the best possible solution only after 400 generations - even though the balanced replenishment strategy is selected

earlier as can be observed by the graph in Figure 3d. However, as long as the random strategy has a rather high weight, the fitness of the defender is not significantly increased. Only after ruling out the random defense, the fitness increases rapidly. That indicates that even a small amount of mixing of strategies may cause a rather bad performance of the defender. This is not the case for the second and third comparison in Figures 3b,3c, 3e, 3f - if the attacker applies the degree attack and betweenness strategy respectively, the defender evolves rapidly in using the balanced replenishment strategy only. The fitness, accordingly, increases quickly in both cases. The defender can deal best with the random attack strategy, as indicated by the comparativley high overall fitness in Figure 3a, while the best strategy for the attacker seems to be betweenness attacks, as also confirmed by the results in the previous section.

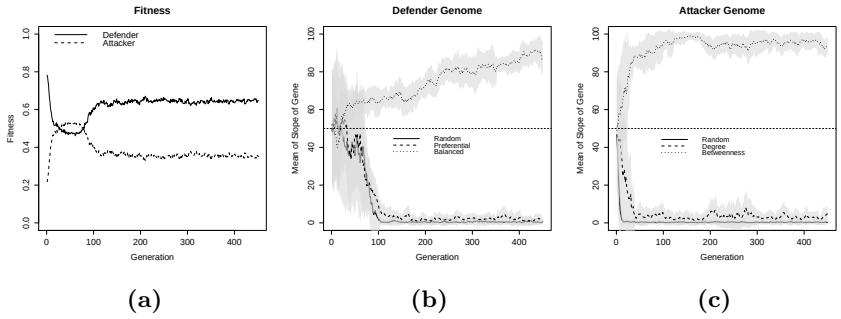


**Fig. 4.** Results of simulation runs. Left: evolution of the mean of fitness in the defender and attacker population in the co-evolution case. Middle: Evolution of the mean of defender weights for different strategies in the GA case. The transparent areas indicate the standard deviation. Right: Evolution of the mean of attacker weights for different strategies in the GA case. The transparent areas indicate the standard deviation.

**Co-Evolution** In the case of co-evolution, i.e. both, defenders and attackers employ a genetic algorithm to select their strategy, attackers evolve quicker towards the more efficient strategy, causing a decline in the fitness of the defender (see Figure 4). However, after about 50 generations, there is a turn-around and the defender starts selecting the best defense strategy, causing an increase in the defender's fitness. After defenders and attackers have evolved into applying the balanced replenishment and betweenness attack strategies respectively, the fitness function stabilizes and no further major fluctuations are observed – an equilibrium is reached. This co-evolutionary process was tested for 25 different instances (while the cases described in the previous section was only tested for 1 instance) and the variance in the overall observed outcome of the gene weights and the fitness of defender and attacker was very low. The pattern shown in

Figure 4 for one instance could, in a similar way, be observed in all instances of the problem.

### 5.2 Sensitivity analysis

The sensitivity analysis assesses the effect of different defender budgets, i.e. the number of edges that are rewired after an attack, on the overall outcome. A high defender budget plus an efficient defense strategy (i.e. balanced replenishment) almost completely reduce the possibility of the attacker to increase her fitness (see Table 1, row Attacker GA vs. Balanced Replenishment and budget of 9). On the other hand, a low budget decreases the fitness improvements over time for the defender (see Table 1, budget of 3). This indicates that a meaningful game can only be played if the available budgets are in a certain, rather limited interval - too high of a budget for one of the two sides will make any response strategy inefficient. In the co-evolution case, the defender shows a lower fitness at the end of the evolution process than in the beginning if the budget is smaller or equal to 5 edges, while it is the other way round for a budget above that level.

| **Defender Budget** | **3** | | **5** | | **7** | | **9** | |
|---|---|---|---|---|---|---|---|---|
| Attacker GA vs. | FAS | FAE | FAS | FAE | FAS | FAE | FAS | FAE |
| Random Defense | 0.38 | 0.63 | 0.22 | 0.52 | 0.12 | 0.37 | 0.08 | 0.18 |
| Preferential Defense | 0.48 | 0.76 | 0.40 | 0.76 | 0.36 | 0.64 | 0.32 | 0.62 |
| Balanced Replenishment | 0.37 | 0.54 | 0.10 | 0.34 | 0.01 | 0.02 | 0.01 | 0.01 |
| Defender GA vs. | FDS | FDE | FDS | FDE | FDS | FDE | FDS | FDE |
| Random Attack | 0.67 | 0.68 | 0.81 | 0.92 | 0.90 | 0.97 | 0.94 | 0.98 |
| Degree Attack | 0.49 | 0.54 | 0.63 | 0.81 | 0.81 | 0.98 | 0.90 | 0.98 |
| Betweenness Attack | 0.36 | 0.42 | 0.45 | 0.63 | 0.61 | 0.95 | 0.82 | 0.97 |
| Co-Evolution | FDS | FDE | FDS | FDE | FDS | FDE | FDS | FDE |
| GA vs. GA | 0.62 | 0.38 | 0.78 | 0.66 | 0.88 | 0.95 | 0.92 | 0.98 |

**Table 1.** Fitness of attackers and defenders with varying budgets. FAS and FAE indicate the average fitness of the attacker at the start and the end of the simulation (i.e. generation 1 and generation 500), respectively. FDS and FDE indicate the average fitness of the defender at the start and at the end of the simulation, respectively.

## 6 Related Work

Several researchers have assessed the robustness of networks in case of attacks on nodes or edges. Here we look more in detail to studies where the concepts of evolution of a network, in terms of its topology, is tied to the behavior of an attacker of the network. In a seminal paper by Albert *et al.* [2], the authors demonstrate that scale-free networks are vulnerable to targeted attacks of nodes of high degree, while fairly robust to random attacks. Holme *et al.* [8] consider

attacks on edges as opposed to nodes, and suggest edge centrality as an effective target of an attacker.

As already mentioned in Section 1, the work of Nagaraja and Anderson [15] is relevant to our paper since it considers an evolutionary game theory approach that takes place on a network. In a way similar to our interpretation of the evolutionary game, their game is organized in rounds and each round consists of an attack followed by a recovery. The attack consists of targeting a number of nodes to be removed, depending on the attacker budget. However, the recovery is different than the one we propose in this paper, and consists in two stages, namely replenishment and adaptation. The first stage deals with inserting new nodes into the network and establishing new connections based on the defender's budget, while the second deals with rewiring existing links. The objective for the attacker is to split the network in separate components. The authors also consider betweenness as a type of attack and the effects are more disrupting against all types of defense. Our approach is more flexible giving the possibility to the attacker and defender to adapt or change their strategies (i.e., type of attack/defense) during the game, while in [15] the strategies are chosen and kept fixed through the game. Our model allows to identify the strategies for attackers and defenders that provide the maximum fitness out of a potentially broad set of strategies. In [15] the test performed takes into account scale free networks as initial topologies, whereas our approach starts with an initial topology that is already optimized by the defender under the assumption that the defender initially generates the network. One aspect that we prove through the evolution of the genome is the superiority in attack of the balanced replenishment strategy that is highlighted also in [15]. Nagaraja and Anderson's work is not without limitations, however. The cost of implementing an edge is essentially zero since the network is allowed to rewire with an arbitrary amount of newly added edges.

Kim and Anderson [11] expand upon the work of Nagaraja and Anderson. Kim and Anderson give each attacker and defender a fixed budget, or cost to add nodes and edges after an attack, and analyze the effect of attacks on a variety of different network topologies. They find a strategy of connecting low centrality nodes is the best defense strategy. However, as the edge to node ratio increases, the network becomes more robust, and even adding edges randomly is effective against targeted attacks. They find that there is a threshold value for the proportion of edges to nodes at which point the effectiveness of attacks decreases drastically.

The work of Domingo-Ferrer and Gonzalez-Nicolas [5] is based on the ideas and findings of previous work by Nagaraja and Anderson [15] and Kim and Anderson [11] and adds further properties to the networks and the experiment set. In the paper the authors analyze the evolution of the *order* and average path length of scale-free networks (weighted and unweighted) under attack and defense. The only strategy of attack considers betweenness centrality as the measure to identify the most critical node; whereas defense is achieved following two types of strategies: delegation and node replenishment. The results show basically that an important factor is the visibility that an attacker has of the network,

while there is basically no difference in the disruption behavior of weighted and unweighted networks. Our approach is more flexible considering the possibilities of different strategies of attack and defense and networks that are not fixed a priori, but built by the defender that is usually the organization that has to defend from the attacks.

## 7 Conclusions and future work

We have shown that our approach to model interactions between attackers and defenders can be successfully modeled using genetic algorithms. Our results confirm what has been found in previous papers which compared various static strategies. In addition, our work shows that strategies for link placement can also be applied to generate networks from scratch, as we do in generating the networks, achieving already an initial strength against some types of attacks (in contrast to other papers, which only used them to rewire networks after they have been attacked)[7]. Obviously, the success of a defense and attack depends on the available resources. The choice of the strategy matters primarily when the defender's resources are slightly larger than the attacker's resources. In any other case, the results of the game are going to be biased towards the side with the resource advantage. If the defender resources are slightly higher than the attacker's and if the defender's goal is to maintain or increase the LCC and the attacker aims for the opposite, there are clear winning strategies among the ones tested in this study: the balanced replenishment and betweenness attack strategy, respectively, can be considered to be the most efficient ones, independent of which strategy is applied by the opponent. An equilibrium situation arises if the two opponents apply these strategies, although the defender appears to evolve slower than the attacker.

This result may be applied to social networks, computer networks, or any other kind of network. From an empirical perspective, it would be interesting if similar strategies are observed in real networks (i.e. where they have evolved 'naturally'). From a normative point of view, the results of this paper and related work can be used to design strategies to defend against attacks or to target attacks against certain nodes in networks.

Future work will include the development and testing of new defender and attacker strategies - currently, only three strategies are included. A larger number of strategies may make the game dynamics more complex than the current version, which allows for a stable equilibrium in the co-evolution case. Additionally, the current fitness function emphasizes connectedness of the network, but does not assess the efficiency of the network in providing transportation or communication services. Different fitness functions which may include a combination of the largest connected component with some measure of efficiency as, for example, the diameter or effective diameter of the network, therefore might be considered interesting options for future research.

[7] However, this difference is somehow minor if we consider that many attack-defense rounds applying the same defense strategies will cause the network topology to resemble a network that was built from scratch using the very same defense strategy.

**Acknowledgements:** This paper is the outcome of a group project started at the Complex Systems Summer School of the Santa Fe Institute in June, 2013. We are very grateful to the organizers of the summer school, and to all the lecturers and presenters. To all of our fellow summer school students: thanks for the good times – and the greenhouse sessions. Thanks to Mauricio Cantor and Bruno Pace for very valuable discussion on the subject of the paper. Very special thanks go to Tom Carter for all of his efforts and in particular for his enlightening performance of the "power-law blues". The simulations for this paper were partly executed on the Vienna Scientific Cluster (VSC) – thanks a lot to the VSC team for their support during the whole process.

## References

1. R. Albert, I. Albert, and G. Nakarado. Structural vulnerability of the north american power grid. *Physical Review E*, 69, 2004.
2. R. Albert, H. Jeong, and A. L. Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, July 2000.
3. M. Boss, H. Elsinger, M. Summer, and S. Thurner. The network topology of the interbank market. *Quantitative Finance*, 4:677–684, 2004.
4. V. Colizza, A. Barrat, M. Barthélemy, and A. Vespignani. Predictability and epidemic pathways in global outbreaks of infectious diseases: the sars case study. *BMC Med*, 5:34, 2007.
5. J. Domingo-Ferrer and rsula Gonzlez-Nicols. Decapitation of networks with and without weights and direction: The economics of iterated attack an d defense. *Computer Networks*, 55(1):119 – 130, 2011.
6. D. E. Goldberg and K. Deb. A comparative analysis of selection schemes used in genetic algorithms. *Urbana*, 51:61801–2996.
7. R. Guimerà, S. Mossa, A. Turtschi, and L. A. N. Amaral. The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *PNAS*, 102(22):7794–9, May 2005.
8. P. Holme, B. J. Kim, C. N. Yoon, and S. K. Han. Attack vulnerability of complex networks. *Physical Review E*, 65(5):056109, 2002.
9. S. Iyer, T. Killingback, B. Sundaram, and Z. Wang. Attack robustness and centrality of complex networks. *PloS one*, 8(4):e59613, 2013.
10. H. Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A. L. Barabási. The large-scale organization of metabolic networks. *Nature*, 407(6804):651–4, Oct. 2000.
11. H. Kim and R. Anderson. An experimental evaluation of robustness of networks. *Systems Journal, IEEE*, 7(2):179–188, 2013.
12. V. E. Krebs. Mapping networks of terrorist cells. *Connections*, 24(3):43–52, 2002.
13. V. Latora and M. Marchiori. Is the boston subway a small-world network? *Physica A: Statistical Mechanics and its Applications*, 314(1-4):109 – 113, 2002.
14. S. Nagaraja. Topology of covert conflict. In *Security Protocols*, pages 329–332. Springer, 2007.
15. S. Nagaraja and R. Anderson. The topology of covert conflict. Technical Report UCAM-CL-TR-637, University of Cambridge, Computer Laboratory, July 2005.
16. M. Newman. *Networks: an introduction*. Oxford University Press, 2009.
17. G. A. Pagani and M. Aiello. The power grid as a complex network: A survey. *Physica A: Statistical Mechanics and its Applications*, 392(11):2688 – 2700, 2013.
18. J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32(4):425–443, 1969.