Telecommunication Systems manuscript No.
(will be inserted by the editor)

On the Resistance of Overlay Networks against Bandwidth Exhaustion Attacks

Franz Girlich · Michael Rossberg · Guenter Schaefer

Received: 01/08/13 / Accepted: 07/12/13

Abstract In order to perform private communication over public networks, such as the Internet, several different kinds of virtual overlay networks emerged. Examples are the well known Virtual Private Networks (VPN), Darknets, and anonymizing networks like Tor. All of these networks are designed to provide data delivery that is confidential, authentic and integrity protected. Nonetheless, for a secure operation also the availability must be taken into account, especially as these structures turn into vital targets for Denial-of-Service (DoS) attacks.

Within this article we present metrics to rate different network topologies with regard to their resistance against botnets, whose available attack bandwidth is not a limiting factor. The presented metrics consider random, greedy, and optimally operating attackers, and are used to derive several properties that very resilient overlay topologies must have. In particular high girth and a low constant node degree were identified and validated by simulations.

Franz Girlich
Ilmenau University of Technology
Tel.: +49-3677-694157
Fax: +49-3677-694540
E-mail: franz.girlich@tu-ilmenau.de

Michael Rossberg
Ilmenau University of Technology
E-mail: michael.rossberg@tu-ilmenau.de

Guenter Schaefer
Ilmenau University of Technology
E-mail: guenter.schaefer@tu-ilmenau.de

SPONSORED BY THE
Federal Ministry of Education and Research

Keywords Denial-of-Service · Resilience · Overlay Networks · Virtual Private Networks

1 Introduction

With the appearance of packet-oriented networks and in particular the Internet, cheap and easy means of global communication became available for a general public. In order to exploit these possibilities also for their private data exchange, governments, companies, and non-government organizations operate Virtual Private Networks (VPNs). But also private individuals use the Internet to create overlay networks. Depending on the purpose that may be friend-to-friend networks [26], or at smaller scale so-called darknets. The latter are usually private networks between trusted peers, and designed to provide information exchange without legal or governmental control.

Most of these networks offer certain security services. Usually at least the confidentiality of the transmitted data is guaranteed, and modern cryptographic protocols make it simple to also verify authenticity and integrity. However, the availability of the different overlay networks is very often neither discussed nor actively improved [38]. While this may not yet pose a problem for networks of private individuals, as they are too small and unimportant, it may cause difficulties in professional environments that depend on a working network infrastructure. Examples of such vital VPNs include the European Energy Exchange (EEX) [14], the European Network Exchange (ENX) as well as the Automotive Network Exchange (ANX).

Over the last decade, attackers tried to impose the availability of vital network services by more and more sophistical Denial-of-Service (DoS) attack tools [12,15]

and large botnets [7,31]. These attacks are hard to prevent, yet easy to execute. For example, bandwidth-exhausting attacks can only be coped with by filtering malicious traffic already at the Internet Service Provider (ISP). However, [16] documents that even this intricate approach does not completely mitigate effects on the offered network services. Due to the encryption of overlay networks it becomes even less effective, as the ISP can no longer distinguish between legitimate and malicious traffic.

In the future, these attacks will continue to be a threat to overlay networks, as no effective countermeasures exist. Furthermore, as society becomes more reliant on VPNs, the main motivations for such attacks become stronger, i.e., it may be getting more attractive to extort companies for not attacking their networks [28]. Also the attractiveness to attack for political reasons may increase, in order to harm competitors, or even a whole foreign country, as it happened in Estonia in 2007 [25] or in the Georgia conflict in 2008 [22].

In order to effectively perform an attack, the attacker needs to know which IP addresses to flood, however, depending on the structure of the overlay network, determining worthwhile targets may become more or less complicated. Currently VPNs are very often organized in so-called Hub-and-Spoke topologies. Here all participants connect to a central coordinator. The other predominant structure is a full mesh between all participating entities. In both cases observing the traffic of any VPN member is sufficient to determine targets, whose elimination leads to a failure of the whole VPN. Previously, approaches like Secure Overlay Services (SOS) [23] tried to address this problem with the introduction of dedicated relay nodes that segregate communicating entities and may handle bandwidth exhaustion attacks at the same time. Even though the dedicated nodes may form exposed points for an attacker, the approach shows that hiding the IP addresses of potential DoS targets may increase the resilience of an overlay significantly.

Another interesting observation can be made: fully meshed topologies are considered extremely robust when non-malicious errors occur, i.e., random and/or correlated ones, but they are not resilient against an intelligently operating attacker. This differentiates the problem from the usually considered creation of robust networks (e.g., [40,39]).

Considering such malicious failures, this article extents the work in [37] and presents:

- a formal model for random, greedy, and optimal DoS attacks,
- metrics to determine the resilience of overlay topologies against such attacks,
- and heuristics to create overlay networks that are optimally protected, when nodes are compromised with equally probability.

The article is structured as follows: In the next section we will discuss attackers of three different strengths' and derive formal models to quantify the resistance of overlays with regard to each attacker type. These models are then used to deduct properties of resilient overlay topologies in section 3. The insights are also used to evaluate the found topologies in comparison with well-known overlay structures in section 4. Other related work will be covered in section 5. The closing section 6 concludes the article.

2 Modeling DoS-Attacks on Overlay Networks

As already indicated, DoS attacks are usually based on a very asymmetric balance of power: While it is easy to generate large amounts of illegitimate traffic, it is rather complicated to filter that traffic already within the ISPs networks. In contrast to previous models, we assume in the following that this asymmetric relation cannot be directly shifted. Thus, we do not consider the generated traffic bandwidth a limiting factor for the attacks. This is due to the following reasons:

- There is a black market for renting botnet capacities, making it affordable for nearly everyone to launch effective attacks against arbitrary targets. Prices start at 200 US-$ for generating between 10-100 $Gbit\,s^{-1}$ for 24 hours [31].
- It might be possible for an attacker to launch so-called *amplification attacks*, i.e., sending small requests to different servers by spoofing the address of the victim and letting the servers send large answers to the victim. Such techniques are known to work with the Domain Name System (DNS), for example, and may increase the attack traffic by factors of up to 50.
- Advanced DoS techniques, such as pulsed attacks, allow to impair many nodes without using more bandwidth resources than an attack on a single target [24].

The assumption to not limit the attack bandwidth makes the determination of good targets the most difficult task for an attacker. After the identification of suited targets, the actual attack takes place like shown in figure 1. When considering the overlay network graph $G = (V, E)$ consisting of the nodes V and overlay connections E, the attacker observes the traffic of a node set $X \subseteq V$, e.g., by sniffing at a hotspot or by compromising one or more nodes of the overlay itself. After

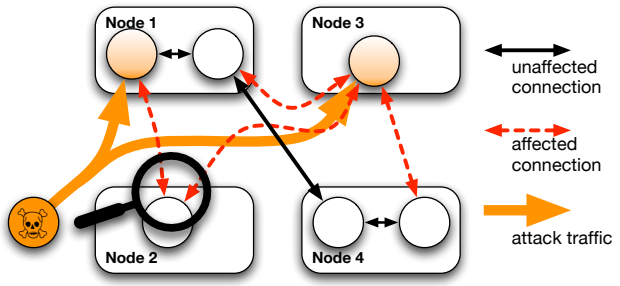

Fig. 1: DoS attack with an attacker observing node 2, leading to a failure of node 3 and one of the interfaces of node 1

an observation period the attacker starts a DoS attack against all neighboring nodes, whose addresses he identified. In the case that the nodes have a single direct connection to the Internet, this attack will lead to a failure of the entire 1-hop neighborhood of all nodes in X.

If multiple interfaces shall be considered, e.g., the nodes 1 and 4 in the example depicted in figure 1, only the uncloaked interfaces will be affected. However, it is easily possible to express also such a scenario with the graph model, by generating a fully connected subgraph for each interface within each device, as described in [5]. Nonetheless, we will concentrate on overlay nodes with single interfaces for simplicity in the following.

Our failure model that considers all traffic of a node to be affected by attacks is underpinned by the results in [18]. The congestion is most likely to occur on the link to the user, not only because of its lower bandwidth capabilities, but also as the wide-spread nature of botnets will distribute the attack traffic well in the core network. Thus, side effects of the attack traffic can be disregarded in a mathematical model without causing oversimplification.

The resilience of the common VPN topologies [33] is rather low when considering this attacker model. For fully-meshed scenarios only a single node needs to be observed to reveal all other nodes and thus allow for arbitrary attacks. In Hub-and-Spoke networks an attacker will usually only discover the central coordinator, and even though these hubs usually have large bandwidth capacities, they still represent a single-point-of-failure. Thus, the observation of a single node can lead to the failure of the whole VPN in the predominant scenarios.

It turns out that not only the overlay topology G determines the effect of attackers, but also the actually observed node set X. In the following, we will concentrate on different strategies that attackers could pursue in order to observe a set X that leads to significant failures. In particular, we consider random, greedy, and optimal observations.

2.1 Random Node Observations

When considering attackers that are not planning to observe specific nodes or behave opportunistically, i.e., observe one or more nodes by chance and then blackmail the corresponding company, the observed set X will be random. That is, each node $v \in V$ has a certain chance p_v to be successfully monitored by an attacker. For different nodes this probability may be very different as some nodes may be physically secured while others roam, for example. Without further knowledge of the network, we will assume the observations to be independent for the following calculations, which is legitimate if the overlay is operating in a global scale. Nonetheless, similar calculations are possible if certain relationships are known, e.g., because a network plan indicates that two devices are in the same building.

Let $D_G(X) \in [0,1]$ be the damage that an attacker can cause after observing a certain node set X. Then the average damage performed by an attacker can be simply derived by calculating the expected damage:

$$E[X] = \sum_{X \in \mathcal{P}(V)} \prod_{x \in X} p_x \prod_{y \in V \setminus X} (1 - p_y)\, D_G(X)$$

Thus, by enumerating the powerset $\mathcal{P}(V)$ of all nodes V, and calculating the probability of every attack times the damage. The calculation of $E[X]$ can also be done for large graphs by utilizing Monte Carlo methods [36], for example.

2.2 Intelligent Node Subset Selection

The introduced random attacker may give a first impression of the networks vulnerability against DoS attacks. However, when calculating the expected value it was assumed that the attacker tried to observe all nodes and succeeded with only a few. This may be unrealistic as an attacker might try to observe only a smaller subset and therefore increase his chances in total. Thus, actually it has to be considered that an attacker only tried to observe the node set that he actually succeeded observing. The probability that a certain node set X was successfully observed is then simply

$$P(X) = \prod_{x \in X} p_x.$$

In accordance the overall damage for random attackers for a graph G is defined by:

$$\mathcal{A}_{rand}(G) = \sum_{X \in \mathcal{P}(V)} \prod_{x \in X} p_x\, D_G(X)$$

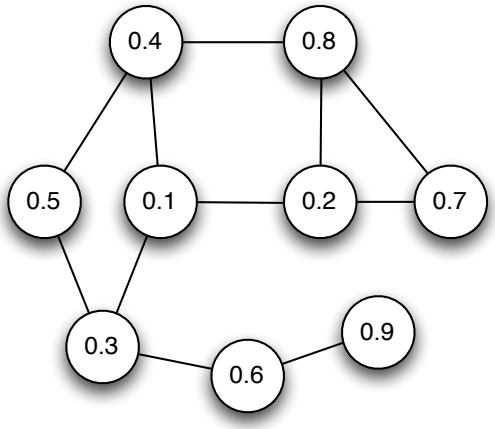

Fig. 2: Example overlay network with observation probabilities between 0.1 and 0.9

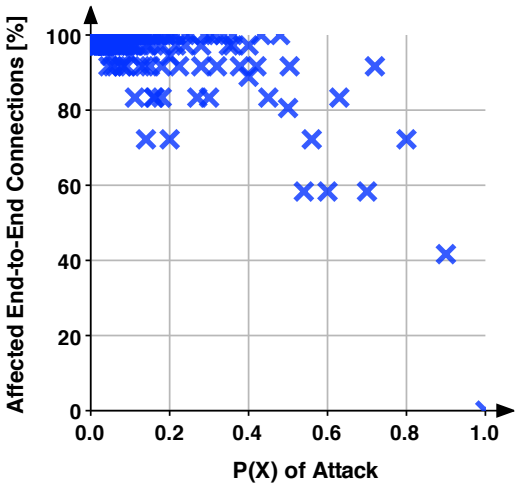

Fig. 3: Damage $D_G(X)$ for all possible observation sets X of the example network

Note that $A_{rand}(G)$ is not an expected value but rather a performance metric with the purpose of comparing topologies.

In order to calculate $\mathcal{A}_{rand}(G)$, an actual damage function $D_G(X)$ must be defined. Instead of counting the number of affected nodes, for example, we suggest using the normalized number of no longer routable end-to-end-connections. This metric better reflects possible partitions of the overlay network. Consider the example overlay in figure 2 with different observation probabilities, the attack following the observation of the node with probability 0.5 does much more damage than one that concentrated on the node probability 0.7. In both cases only 2 nodes are attacked, but in the later case no partitioning occurs. The damage is therefore much smaller.

Figure 3 shows the different number of dropped end-to-end-connections for all possible observation sets (and the corresponding attacks) in the example overlay from figure 2. The x-axis shows the probability $P(X)$, i.e., the probability that an attacker succeeds in observing the set X. The introduced metric $\mathcal{A}_{rand}(G)$ can be interpreted as the weighted average over all points. In the presented example overlay graph $\mathcal{A}_{rand}(G)$ is 29.7, which can be compared with other graphs of equal node count and observation probabilities. Nonetheless, figure 3 also show that there are some node sets that are easier to observe than others, while at the same time creating a larger damage. Thus, a more intelligent attacker might try to plan observations better. This is reflected by the following, more intelligent attacker model.

2.3 Greedy Selection of Observed Node Sets

In contrast to randomly selecting node sets for observation, an attacker might try to choose nodes in a more sophisticated way. For example, he might try to observe nodes with a higher degree, e.g., supernodes or VPN access concentrators. It might also be possible to observe nodes in different geographic regions in order to have the one and two hop neighborhoods of the observed nodes to not overlap.

Thus, the greedy attacker would plan his attack by sequentially adding nodes to the observation set X, depending on the estimated impact and the "budget" P_{min}. That is the required lower bound for the probability that the chosen observation will succeed in a given time period. Please note: While costs for greedy attackers are usually considered to grow linearly with the size of X, the proposed model uses the multiplicative cost model from random attacks. This is done as it is believed to reflect real world effects in a better way, i.e., chances for an observation to fail increases more than linear with the set size.

Thus, for a given target damage D_{min} the greedy attacker tries to find a suited node set X that maximizes the following function:

$$\max_{\substack{X \subseteq V \\ D_G(X) \geq D_{min}}} \prod_{v \in V} (1 - (1 - p_v)x_v)\,.$$

Whereas, x_v represents a binary variable that indicates whether a node v is observed (1) or not observed (0). The multiplicative effect makes it difficult to perform the actual optimizations. However, as x_v is binary and $p_v \in [0, 1]$, it is possible to reformulate the function as

follows:

$$\begin{aligned}\max \log \prod_{v\in V} (1-(1-p_v)x_v) \\ &= \max \sum_{v\in V} \log\left(1-(1-p_v)x_v\right) \\ &= \max \sum_{v\in V} \begin{cases} \log p_v & \text{if } x_v = 1 \\ 0 & \text{if } x_v = 0 \end{cases} \\ &= \max \sum_{v\in V} x_v \log p_v\end{aligned}$$

Using that formula, a greedy attacker may be implemented as outlined in algorithm 1.

Algorithm 1 Greedy attacker $D_{greedy}(G, P_{min})$

$G \leftarrow (V, E)$, $X \leftarrow \emptyset$, $P \leftarrow 1.0$
while $\{v \in V |\ P \cdot p_v \geq P_{min}\} \neq \emptyset$ **do**
  $v \in \text{argmax}\{D_G(X \cup \{v\}) \mid v \in V \setminus X,\ P \cdot p_v \geq P_{min}\}$
  $X \leftarrow X \cup \{v\}$
  $P \leftarrow P \cdot p_v$
  $G \leftarrow$ induced subgraph of G graph without v and its adjacent nodes
end while
return $D_G(X)$

Interpreting the result, the greedy attacker selects certain observable node sets from the grand total, i.e., from figure 3. Nonetheless, the actually achieved damage for a specific overlay network may vary, depending on the budget P_{min}. This makes it even more difficult to compare different graphs of equal node count. In contrast to the previous attacker model, there is only one attack set for each P_{min}, however. Thus, it is possible to obtain an overall performance index by calculating the area under the function $D_{greedy}(G, P_{min})$. Accordingly, we define $\mathcal{A}_{greedy}(G)$ to be:

$$\mathcal{A}_{greedy}(G) = \int_0^1 D_{greedy}(G, P_{min})\, \mathrm{d}P_{min}$$

For the example overlay graph from figure 2, the area under $D_{greedy}(G, P_{min})$ is plotted in figure 4.

2.4 Optimal Selection of Observed Node Sets

Usually in the context of IT security, not only the average resilience or the resilience against simple (greedy) attackers is of interest, but also the worst case. In our given scenario, this corresponds to attackers that have full topology knowledge and plan attacks with that information optimally. This knowledge might either be available when a deterministic construction algorithm is used or when necessary data can be extracted from

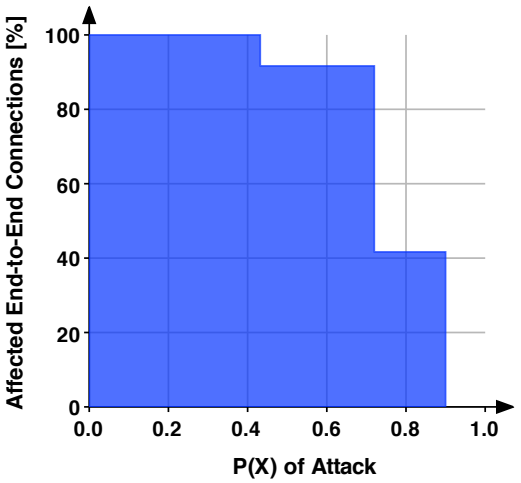

Fig. 4: Impact of the greedy attacker on the example network in figure 1

the routing information. We again assume that observing the traffic of a certain node v imposes a difficulty on the attacker, hence making an attack of required P_{min} less likely to occur. Therefore, we denote the damage caused by an attacker as:

$$D_{opt}(G, P_{min}) = max\left\{ D_G(X) \;\middle|\; X \subseteq V, \sum_{x\in X} \log p_x \geq \log P_{min} \right\}$$

In analogy to the greedy attacker, we define the overall DoS vulnerability performance index to be:

$$\mathcal{A}_{opt}(G) = \int_0^1 D_{opt}(G, P_{min})\, \mathrm{d}P_{min}$$

Given that $D_G(X) \in [0, 1]$, we will refer to resilience as $1 - \mathcal{A}_{opt}(G)$ in the following.

However, while a greedy attack can be computed in $\mathcal{O}(n^4)$, the complexity of finding an optimal attack grows exponentially with the size of G. To prove this claim, we will show that the respective decision problem $D^{\mathcal{D}}_{opt}$ is NP-complete. For a given graph G with observation probabilities p, a budget P_{min} and damage D, we define $(G, p, P_{min}, D) \in D^{\mathcal{D}}_{opt}$ if and only if there is an attack X on G, so that $\sum_{x\in X} \log p_x \geq \log P_{min}$ and $D_G(X) \geq D$.

Lemma 1 *The decision problem $D^{\mathcal{D}}_{opt}$ is NP-complete.*

Proof Since the calculation of damages of given attacks is in FP and all possible attacks can easily be enumerated, $D^{\mathcal{D}}_{opt}$ must be in NP. In order to show the NP-hardness, we reduce the well-known NP-complete problem *vertex cover* VC [21] to $D^{\mathcal{D}}_{opt}$. Using a polynomial-time computable function f, VC inputs $x = (G, k)$

are mapped to $D_{opt}^{\mathcal{D}}$ inputs $f(x) = (G', p, P_{min}, D)$ so that $x \in VC \Leftrightarrow f(x) \in D_{opt}^{\mathcal{D}}$. The transformed graph $G' = (V', E')$ is generated from the original $G = (V, E)$ by augmenting each edge by four additional nodes as sketched in figure 5. The graph G' is thus defined by:

$$
\begin{aligned}
V' &= V \cup N \\
N &= \{u_{e,i}, v_{e,i} \mid e = \{u, v\} \in E,\ i \in \{0, 1\}\} \\
E' &= \{\{x, y_{e,i}\} \mid e = \{u, v\} \in E,\ x, y \in e,\ i \in \{0, 1\}\} \\
&\quad \cup \{\{x_{e,0}, x_{e,1}\} \mid e \in E,\ x \in e\}
\end{aligned}
$$

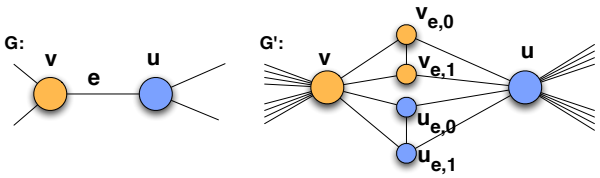

Fig. 5: Transformation of the edge $e \in E$ to map a VC instance x to an instance $f(x)$ of $D_{opt}^{\mathcal{D}}$

The observation probabilities p_v can be set to an arbitrarily chosen constant $\in (0, 1)$ for all nodes in V'. Furthermore, the lower bound for the damage D that the attacker must achieve is set to the maximum value 1. With the budget P_{min} set to p^k, a feasible attack X has to disrupt any possible communication path by observing at most k nodes.

We will now show that with our construction, any vertex cover problem solution can be derived by solving the corresponding problem in $D_{opt}^{\mathcal{D}}$.

$x \in VC \Rightarrow f(x) \in D_{opt}^{\mathcal{D}}$: If there is a vertex cover on G with at most k node, an attack X can be constructed by observing the same k nodes in G'. Since each edge in G is incident to a node in X, all nodes in N are adjacent to a node in X as well. The removal of all nodes in X and all of their neighbors leaves the nodes in $V \setminus X$ isolated. Therefore, no communication is possible and the damage is maximal.

$f(x) \in D_{opt}^{\mathcal{D}} \Rightarrow x \in VC$: If there is an attack observing at most $|X| = k$ nodes so that $f(x) \in D_{opt}^{\mathcal{D}}$, a vertex cover X' for G can be constructed as

$$X' = (X \cap V) \cup \{u \mid u_{e,i} \in X, e \in E, i \in \{0, 1\}\}$$

The observed nodes in $X \cap V$ are also a part of the vertex cover solution. If there are any nodes in $N \cap X$, then a corresponding adjacent node in V is used instead:

When considering an edge $e = u, v \in E$ that has been transformed in a subgraph according to figure 5, either of two cases can occur. In one case at least v or u are observed to disable the subnet, hence e is covered. If neither v nor u is in X, at least one $v_{e,i}$ and one $u_{e,i}$ have to be observed to keep them from communicating. In this case, either v or u may be selected for the vertex cover. Therefore, each edge in G is incident to a node in X.

While finding optimal attacks on a given topology is a hard problem for the attacker to solve, this is good news for operators only at first sight. It also makes finding optimal topologies for network sizes of practical relevance an infeasible task, as discussed in the next section. Furthermore, attacks may still be approximated or optimally determined in exponential time by solving the following Binary Linear Program (BLP):

$$D_G(X) = \frac{1}{2} \max \sum_{s,t \in V, s \neq t} 1 - f_{s,t} \equiv \frac{1}{2} \min \sum_{s,t \in V, s \neq t} f_{s,t}$$

$$\forall v \in V: \quad a_v \leq x_v + \sum_{n \in \mathcal{N}(v)} x_n \tag{1}$$

$$\forall v \in V: \quad f_{v,v} = 1 - a_v \tag{2}$$

$$\forall s \neq t \in V, n \in \mathcal{N}(t): \quad f_{s,t} + a_t \geq f_{s,n} \tag{3}$$

$$\sum_{v \in V} x_v \log p_v \geq \log P_{min} \tag{4}$$

$$\forall u, v \in V: \quad 0 \leq f_{u,v} \leq 1$$

$$\forall v \in V: \quad 0 \leq a_v \leq 1$$

$$\forall v \in V: \quad x_v \in \{0, 1\}$$

The BLP determines an attack that minimizes the number of end-to-end connections. Constraint (1) ensures that $a_v = 1$ for each node v that is either observed ($x_v = 1$) or neighbor of an observed node, with $\mathcal{N}(v)$ denoting the neighborhood of v. This marks the node being a victim of the attack. Hence, the constraints (2) and (3) force the reachability $f_{s,t}$ from s to t to become 1, if and only if it passes no attacked nodes. Finally, constraint (4) guarantees that the budget of the attacker is met.

For efficient calculation, the program is designed to use as few as possible explicitly marked binary variables. Nonetheless, the program forces the variables $f_{u,v}$ and a_v to be either 0 or 1 even if they are treated to be continuous. This eases actually solving instances of the problem. However, the decision, whether a node is observed is still binary, hence $|V|$ binary variables (x_v) are required.

Since the solved problem is NP-hard, the required effort to solve the BLP must also grow exponentially. Nonetheless, depending on the graphs' structure and the observation probabilities, optimal attacks for graphs

up to a hundred nodes can be calculated on common hardware.

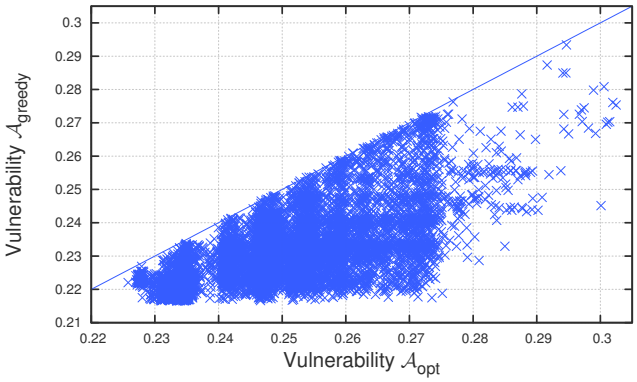

Fig. 6: Evaluation of several 3-regular graphs with optimal and greedy attacks. For graphs that are far from the line, the greedy attack fails in finding an effective attack.

Since solving the optimization problem is infeasible for large graphs, the idea might come up to use the greedy attacker to evaluate vulnerability and resilience instead. However, depending on the graph's structure, the approximation $\mathcal{A}_{greedy}$ can differ greatly from $\mathcal{A}_{opt}$. Figure 6 shows this fact by comparing the vulnerability metrics for a variety of 3-regular graphs, which are considered to be rather resilient. Details on the selection will be given in section 3.2.1. The better the graph's resilience is approximated by the greedy attacker, the closer the points are to the diagonal line. As there are several graphs with a large distance to the ideal line, the greedy attacker cannot be considered a reliable indicator for $\mathcal{A}_{opt}$. By manual analysis, we found out that graphs with a high greedy resilience are often partitioned by the optimal attacker by observing two or more nodes. Since such a partition is often only generated by the optimal attack, the induced damage is magnitudes higher.

However, modeling and analyzing attacks on given graphs is only a first step. The remainder of this article focuses on finding topologies that minimize the damage induced by optimal attacks.

3 Construction of DoS-resilient Overlays

While the resilience evaluation of a given graph is an NP-hard problem, it is still feasible for smaller networks. Thus, DoS-resilient overlay networks should consider this attacker type. However, finding a topology that provides optimal resilience against such optimal attacks is an even harder problem. In the following, the general problem of finding optimal topologies is defined. Afterwards, several properties are discussed that indicate high DoS-resilience given all nodes have equal observation probabilities.

3.1 Finding Optimal Topologies

The challenge of generating optimal DoS-resilient network topologies is a so-called *Stackelberg competition* [10]. The network operator takes the role of the *leader*, making his first move by creating the overlay topology, which is observable to the attacker. Afterwards, the attacker — in the role of a *follower* — can react based on the network given by the leader.

Therefore, when optimizing his network, the leader must not only consider minimizing the cost, e.g., delay, bandwidth, or monetary costs, but also minimizing the damage of a potential attacker. A trade-off between these possibly contradictory goals may be achieved by weighting them in the objective function. Given the designed network, the attacker calculates the attack set by solving the previously defined optimization problem $\mathcal{A}_{opt}$. This leads the following problem definition:

$$\begin{aligned}\text{operator: } & \operatorname*{argmin}_{G}\{attackDamage(G,X) + c \cdot costs(G),\\ & \quad \text{for feasible topologies } G \text{ and sets } X\}\\ \text{attacker: } & \operatorname*{argmax}_{X}\{attackDamage(G,X),\\ & \quad \text{for feasible observed node sets } X\}\end{aligned}$$

Due to the structure of this problem instance, i.e., that the follower problem is also part of the problem formulation of the leader, the optimization problem for the operator has at least one optimal solution. Finding this solution, however, is even harder than finding an optimal attack. This is because the NP-hard attacker objective function is also part of the operator's objective. Given the current state of the art in bi-level programming [1,10], finding optimal topologies is currently infeasible for graph sizes of practical relevance, e.g., hundreds of nodes.

To still approach the problem, heuristics have to be found and simplifications have to be made. First, we will focus on rather small graphs in the following, i.e., of 30 and 64 nodes, in order to derive construction principles for resilient graphs. It is not only faster to determine these graphs' resilience, but there are also fewer relevant graphs in number. Second, we will assume that the observation probability is equal for each node. Again, this shrinks the search space and speeds up the resilience calculation considerably. Furthermore, since real data on practical observation probabilities is

scarce, this produces more comparable results. Third, by deducting properties of resilient graphs, i.e., a low degree, the search space can be further reduced.

3.2 Properties of DoS-resilient Topologies

Even the simplified construction problem with the same observation probability for all nodes, makes finding optimal attacks still NP-hard. This is because observation probabilities are not essential for the proof of lemma 1. However, there are several metrics that are easier to calculate and correlate with the resilience as will be shown in the following.

3.2.1 Node degree in DoS-resilient Graphs

When considering the node degree, two main effects may be observed. On the one hand, connectivity can be increased by increasing the node degree. This results in structures that are more resilient against partitioning. On the other hand, a high node degree implies that more nodes are disclosed by observed nodes. Since these effects are contradictory, there is a break even point, where the vulnerability against partitioning and the vulnerability due to high node degree are balanced.

Furthermore, it can be argued that resilient graphs must be regular. In order to increase its gain, the optimal and greedy attackers will prefer observing nodes with a high degree. Therefore, graphs that contain only nodes of equal degree are potentially more resilient since they do not contain exposed nodes.

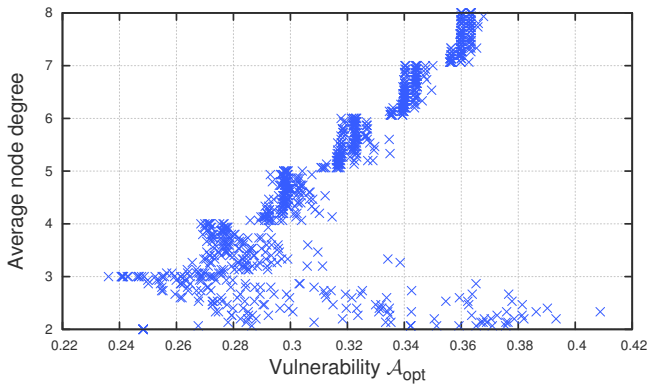

Fig. 7: Vulnerability against optimal attacks: random graphs consisting of 30 nodes with varying node degree; within each graph the node degree differs by at most one. Sample correlation $\rho = 0.714166$ ($\rho' = 0.963795$ for average degree ≥ 3).

In order to evaluate these arguments and find the optimal node degree, we performed an experiment in which 1,000 graphs with 30 nodes were optimally attacked. By construction, each of these graphs was either regular or the difference between the lowest and highest node degree was one. The vulnerabilities were obtained by solving the optimization problem from section 2.4 using the Gurobi 5.0 [20] optimizer. The results depicted in figure 7 indicate an optimal node degree of 3. For graphs of higher degree, the vulnerability increases in a step function. Each step occurs with graphs that are almost regular, i.e., have only one node with one more neighbor than the others. This node is exposed to the attacker and significantly decreases the resilience of the whole graph.

The resilience of graphs with a degree below 3 varies greatly. This is because some of them can be partitioned by observing a single node. Surprisingly the only connected 2-regular graph (a ring) has a relatively high resilience. Even though its connectivity is low, it takes at least two observations to partition it. Furthermore, each observation unveils at most two other nodes. However, as already outlined, several random 3-regular graphs show a higher resilience in this experiment.

3.2.2 The Influence of Short Cycles

The previous experiment indicates that a higher resilience can be achieved by a low and homogeneous node degree. However, figure 7 also suggests that this criterion is not sufficient for resilient graphs.

Even 3-regular graphs are vulnerable, if they contain small cycles. Because of the low node degree, nodes within such a cycle are not well connected to the remaining graph, even though they are strongly connected with each other. An example for this effect is depicted in figure 8: the cycle of length 3 is connected to at most 3 different other nodes. If these 3 other nodes are affected by an attack, the cycle becomes partitioned from the remainder of the graph and the damage increases significantly. This gives also an explanation for the relatively high resilience of the 2-regular graph in figure 7. It contains only one large cycle. By augmenting any additional edge the length of the smallest cycle is reduced and the graph can be partitioned by observing a single node.

Therefore, we conjecture that the longer the shortest cycle in a graph, the so-called girth, the higher its resilience. In order to evaluate this hypothesis, 10,000 3-regular graphs with different girth were generated [29] and assessed by calculating an optimal attack. Figure 9 shows the vulnerability in correspondence to the average length of the smallest non-trivial cycle through

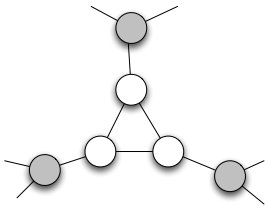

Fig. 8: Segment of a 3-regular graph containing a small cycle that becomes partitioned when the gray nodes are taken down by an attack

each node, where each cycle size 3 or higher is considered non-trivial. A visible result is that there is correlation between vulnerability and girth, and it becomes stronger for more resilient graphs. In particular, all evaluated graphs of girth 7 and 8 have a rather high resilience.

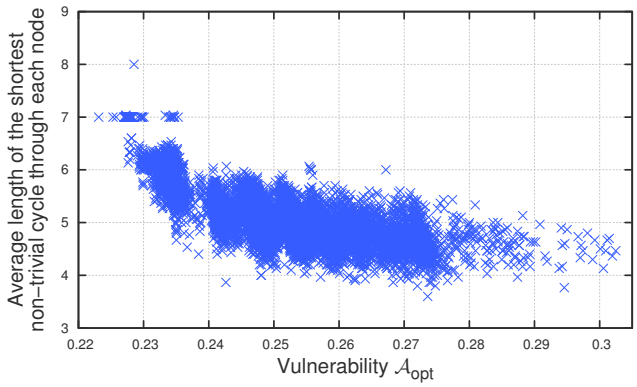

Fig. 9: Vulnerability against an optimal attacks: 30 node3-regular random graphs with varying girths; Evaluation of average length of the smallest cycle through a node. Sample correlation $\rho = -0.739117$.

According to this correlation it could be assumed that graphs with minimal size at given girth, so-called *cage graphs*, are the most resilient. Surprisingly, the only cage with 30 nodes, a so-called *Levi graph* [9], has a girth of 8 and is not the most resilient. Possible reasons for this are inconsistencies due to the small size of the examined graphs or the even cycle length which causes the graph to be bipartite and therefore potentially easier to partition.

3.2.3 Homogeneous Graph Structures

With an increase in the girth, the graph simultaneously becomes more homogeneous. This means in extension to a homogeneous node degree, i.e., the size of the one hop neighborhood, also the sizes of neighborhoods of two or more nodes become more equal among the nodes.

Lemma 2 *Given a d-regular graph ($d \geq 2$) with girth $g \geq 3$, the size of a node's $\lceil g-2/2 \rceil$ hop neighborhood $\mathcal{N}$ consists of exactly $\sum_{i=1}^{\lceil g-2/2 \rceil} d(d-1)^{i-1}$ distinct other nodes.*

Proof 1. $|\mathcal{N}| \leq \sum_{i=1}^{\lceil g-2/2 \rceil} d(d-1)^{i-1}$
The maximum distinct k-hop neighborhood in a d-regular graph can be achieved by constructing a $d-1$-ary tree of depth k where only the root has d neighbors. Each node has degree of at most d and the root has a distinct k-hop neighborhood of $\sum_{i=1}^{k} d(d-1)^{i-1}$.
2. $|\mathcal{N}| \geq \sum_{i=1}^{\lceil g-2/2 \rceil} d(d-1)^{i-1}$
Assuming there are less nodes in the neighborhood $\mathcal{N}$, then at least two nodes from the previously constructed tree must be identical. This closes a cycle of at most $2\lceil g-2/2 \rceil < g$ nodes.

Therefore, larger neighborhood sets become of equal size as the girth increases. The increasing homogeneity within the graph can be measured by comparing the results of node metrics throughout the graph. One well known node metric is the *closeness*, i.e., the average length of the shortest path to all other nodes. When calculating the closeness of all nodes, the sample variance of these values can be used to measure homogeneity. We conjecture that the more equal nodes are, regarding their closeness, the more equal they are as potential targets to the attacker.

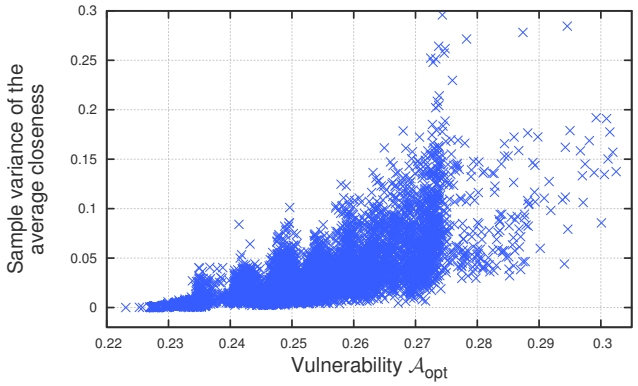

Fig. 10: Vulnerability against optimal attacks: 30 node 3-regular random graphs with varying girths; Evaluation of the sample variance of the average closeness per node. Sample correlation $\rho = 0.805688$.

Figure 10 shows the relation between vulnerability and sample variance of the average closeness for the

previously introduced set of 10,000 graphs. A clear correlation becomes visible, and the most resilient graphs from the sample do not show any variance at all.

Therefore, we conjecture that given any node-specific metric, its sample variance throughout the graph approaches zero as the graphs become more resilient. To further substantiate this hypothesis a more strict node metric was designed.

Let $r_x = (r_{x,1}, \ldots, r_{x,l})$ be the reachability vector, so that $r_{x,i}$ is the number of nodes x that can be reached over exactly i hops using the shortest path. When the average number of reachable nodes over exactly i hops is $\bar{r}_i = \frac{1}{|V|} \sum_{x \in V} r_{x,i}$, the inhomogeneity $I(G)$ of a graph G can be defined to be the sum of the standard sample deviations of $r_{x,i}$:

$$I(G) = \sum_{i \in \{1,\ldots,l\}} \sqrt{\frac{1}{|V|-1} \sum_{x \in V} (r_{x,i} - \bar{r}_i)^2}$$

While being similar to the closeness, $I(G)$ compares the size of k-hop neighborhoods for each k instead of their weighted average. Therefore, the influence of possible inhomogeneities should be measured with more significance.

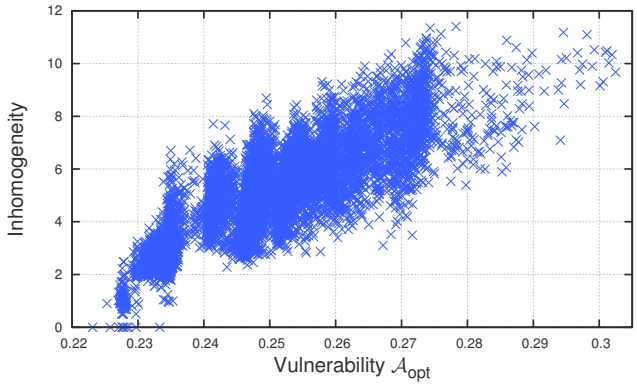

Fig. 11: Vulnerability against optimal attacks: 30 node 3-regular random graphs with varying girths; Evaluation of inhomogeneity $I(G)$. Sample correlation $\rho = 0.850833$.

Figure 11 depicts the correlation of $I(G)$ and the vulnerability $\mathcal{A}_{opt}(G)$. This plot strengthens the previous assumptions that more homogeneously structured are more resilient. The most resilient graphs are even found to be perfectly homogeneous regarding this metric, i.e., for each k, each node can reach the same amount of nodes via k hops.

3.2.4 The Influence of High Connectivity

Given a constant degree d, graphs that are highly connected are suspected to be more resilient since they are harder to partition. One metric to intuitively express connectivity in regular graphs of equal size is the diameter.

The previously introduced lemma 2 implies that in a graph of girth g each node has a k-neighborhood of maximum size for $d \leq \lceil g-2/2 \rceil$. By increasing the number of distinct nodes reachable over paths of length $\leq k$, a high girth therefore decreases the diameter.

Thus, the diameter is a simple indicator for a high girth and by extension an indicator for a high resilience. Another reason for low-diameter graphs having a high resilience is the fact that short paths support the resilience directly, since it becomes harder for an attacker to disrupt communication at intermediate nodes.

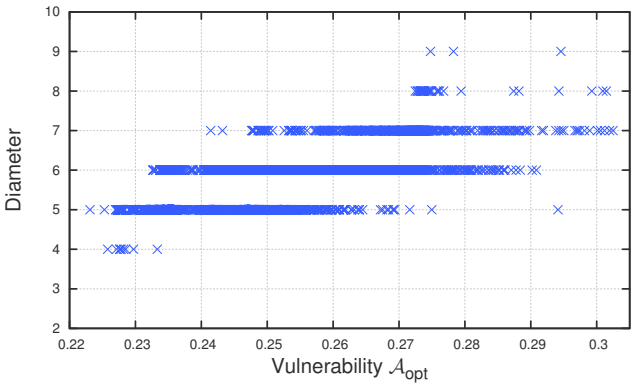

Fig. 12: Vulnerability against optimal attacks: 30 node 3-regular random graphs with varying girths; Evaluation of diameter. Sample correlation $\rho = 0.756561$.

Figure 12 shows the examined graphs' diameters in relation to their vulnerability. As expected, a significant correlation was found. However, this correlation is not as strong as with the previously examined metrics. This is likely caused by the fact that the diameter is a maximum metric. Even a few longer paths increase it, while they may have only little influence on the resilience.

By calculating the average instead of the maximum path length these effects can be reduced. The correlation indicated by figure 13 is much stronger and clearly shows the impact of high connectivity on the vulnerability.

A different connectivity metric is the so-called *algebraic connectivity*, which is defined as the smallest non-zero eigenvalue of the Laplacian matrix of the graph. In this special case of regular graphs, the algebraic con-

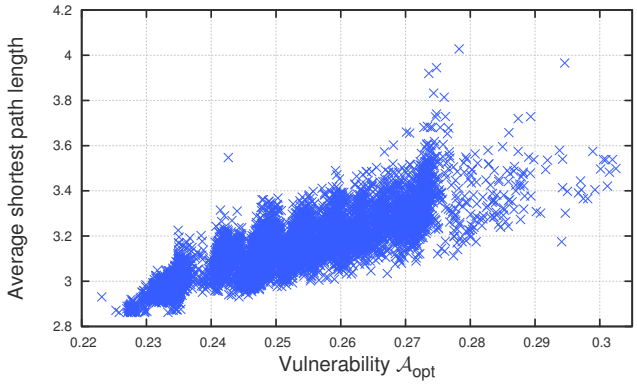

Fig. 13: Vulnerability against optimal attacks: 30 node 3-regular random graphs with varying girths; Evaluation of average shortest path length. Sample correlation $\rho = 0.852539$.

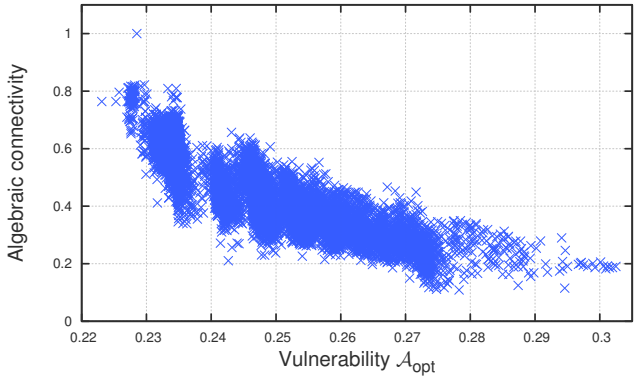

Fig. 14: Vulnerability against optimal attacks: 30 node 3-regular random graphs with varying girths; Evaluation of Algebraic connectivity. Sample correlation $\rho = -0.823763$.

nectivity is equal to the spectral gap, which is the difference of the two largest eigenvalues of the adjacency matrix. It can be shown that the algebraic connectivity is bounded below by $4/d|V|$ with diameter d [30]. The results in figure 14 indicate that the algebraic connectivity is closely correlated with the vulnerability as well.

3.2.5 Constructing DoS-resilient Topologies

When summarizing the presented metrics, the node degree has the strongest impact on resilience. More precisely, 3-regular graphs were the most resilient. While this constraint is fairly easy to respect when designing a network, it has one significant downside. There are only 3-regular graphs of even size. Furthermore, our evaluation showed, that a single node violating the regularity usually increases the graph's vulnerability significantly. Furthermore, short paths and high connectivity also strengthen the topology against DoS attacks; properties that are normally considered anyways due to efficiency implications.

A harder to enforce property is the lack of short cycles and the implied homogeneity, especially in distributed environments. Even in a centralized planning procedure, finding graphs with optimal high girth – so called *cages* – is a hard combinatorial problem. For small instances calculation is still possible [29] and there are several approximation algorithms for larger graphs [3, 43].

4 DoS-resilience of Well-Known Graphs

The introduced graph metrics can improve the understanding of what makes some overlay networks more DoS-resilient than others. They can also assist in the vulnerability evaluation of large topologies, where solving the optimization problem is infeasible. However, there are several well-known overlay topologies designed to be efficient and resilient with regard to previously existing metrics. In this section we will compare several of these commonly used topologies regarding their properties and DoS-resilience.

4.1 Qualitative Analysis

Before taking a look at the vulnerability in numbers, we introduce the graphs and compare their relevant properties, such as node degree and diameter. Furthermore, we *strengthen* several graph types in order make a direct comparison with the novel metric more fair.

4.1.1 Hypercube

Because of their low diameter of $\log(|V|)$ and relatively low node degree of $\log(|V|)$, hypercubes have been applied in a variety of scenarios, such as parallel computing [13] and overlay systems [27]. They are highly connected and homogeneous, however, regarding the defined vulnerability metric, their non-constant node degree imposes a weakness, since attackers observing a single node can disable more nodes in larger graphs. Furthermore, as hypercubes are only defined for $|V| = 2^k$, graphs of arbitrary size can only approximate the hypercube structure. Theses imperfect structures do not necessarily maintain the high connectivity and homogeneity the ideal hypercubes provide. This weakens their practical relevance significantly, as it is desirable to construct networks of any size with a consistently high resilience.

4.1.2 Cube Connected Cycles

In order to overcome the hypercubes' potential weak spot, their growing node degree, transformations can be applied. By replacing each hypercube node of degree d by a cycle of d nodes and connecting each to a former neighbor a 3-regular graph is generated, the so-called *cube connected cycles* [32]. The reduction of node degree should improve the resilience significantly. However, the transformation also decreases homogeneity and increases the diameter. According to the previous results, this again indicates a potential loss of resilience. Furthermore, there are only cube connected cycles with $|V| = k2^k$ nodes, since each of the hypercube's $|V| = 2^k$ nodes are replaced by a k cycle.

4.1.3 Butterfly Graphs

Quite comparable to the cube connected cycles are butterfly graphs, since they also provide a low diameter and a constant node degree. They are constructed from $k + 1$ interconnected layers of 2^k nodes each. However, this scheme induces inhomogeneities in the first and last layer: Nodes in these respective layers have only three neighbors, while all other nodes have four. This inhomogeneity can be fixed by identifying the first and the last layer and transforming the graph to a *wrapped* Butterfly graph [19]. The transformed graphs have a constant and consistent node degree of four and contain $k2^k$ nodes. Compared to the cube connected cycles a higher vulnerability can be expected, due to the slightly increased degree.

4.1.4 "De Bruijn" Graphs

"De Bruijn" graphs are directed graphs of b^k nodes. They can be constructed by identifying each node with a k-adic number to the base b. Each node is then connected to all other nodes that are a left shift of its number. Therefore, nodes have a constant out-degree (and in-degree) of b. Since our analysis focuses on undirected graphs with low diameter, we chose $b = 2$ and undirected each edge. With $b = 2$, the nodes 0...0 and 1...1 have self edges, so their degree is reduced to 2 in a digraph. Furthermore, the nodes 0101... and 1010... are bidirectionally connected, thus their undirected degree becomes 3. All other nodes have a distinct incoming and outgoing neighborhood and therefore degree 4.

In order to fix these inconsistencies we connected 0...0 with 1...1, 0...0 with 1010..., and 1...1 with 0101.... This creates 4-regular *fixed* "De Bruijn" graphs with increased connectivity and homogeneity. Their properties are otherwise similar to wrapped Butterfly graphs.

4.1.5 Content-Addressable Network (CAN)

The CAN [35] topology is non-deterministically constructed by segmenting a toroidal d-dimensional hyperplane. Each node is sequentially added to the plane at a random position. If the segment it is positioned on already contains a node, the segment is split in two so that each part contains one node. The split is done either horizontally or vertically. After all nodes have been added, the plane is segmented into $|V|$ patches, each containing a single node. The graph is then derived by connecting each pair of nodes whose corresponding segments are adjacent.

In contrast to the previously discussed structures this method allows for graph generation of arbitrary size. The average node degree depends on d, but not on $|V|$. However, the degree varies within the graph, leaving some nodes more vulnerable than others to the defined attacker, thus weakening the graph. Furthermore, the diameter is proportional to $n^{1/d}$, indicating that CAN graphs are more loosely connected than the previous graphs. In order to keep the degree minimal, we chose $d = 2$, which results in planar graphs with an average degree between 4 and 5.

4.2 Resilience Evaluation

Each of the presented structures represents a class of infinite graphs. In order to compare their resilience specific instances have to be analyzed. This imposes a difficulty, since most of these structures are only defined for particular sizes such as 2^k or $k2^k$. Fortunately, with the exception of the Butterfly Graphs, each structure has a graph with 64 nodes. Furthermore, it is still feasible to determine the resilience of graphs with 64 nodes using the BLP given in section 2.4. As previously discussed, Butterfly Graphs can be strengthened by *wrapping* them, reducing their size from $(k + 1)2^k$ to $k2^k = 64$ with $k = 4$. Since random 3-regular graphs and CAN graphs are not constructed deterministically, the resilience of 100 graphs each was averaged.

To compare the well-known structures with the results from the previous section, 5 3-regular graphs with 64 nodes and a girth of 9 – the maximum with these properties – were generated [29] and optimized as well.

The results depicted in figure 15 verify the presumptions and substantiate the results from section 3. The 3-regular structures (random, cube connected cycles, and high girth) were found to be the most resilient followed by the 4-regular candidates (Wrapped Butterfly and Fixed "De Bruijn" Graph), the highly homogeneous Hypercube, the inhomogeneous "De Bruijn" Graph and CAN structures.

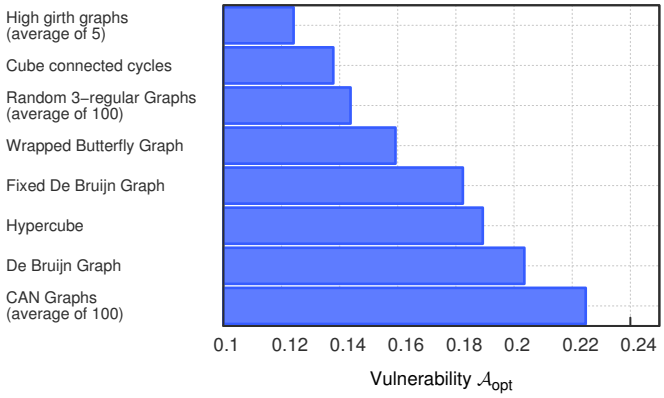

Fig. 15: Vulnerability comparison of different 64 node graphs.

CAN graphs suffer not only from their inhomogeneous node degree, but also from the fact that they are planar and therefore potentially easy to separate into larger partitions. The "De Bruijn" graph has an inhomogeneous node degree as well. Fixing this issue by increasing the connectivity reduced the vulnerability of this structure significantly. Among the regular graphs, the Hypercube was found to be the least resilient. While it is highly homogeneous and connected, its dynamically increasing node degree – 6 with 64 nodes – cannot be compensated, the 4-regular Wrapped Butterfly Graph is much more resilient. Cube connected cycles attempt to combine the properties of Hypercubes with a low and constant node degree. This shows that there are graphs that are more resilient than 3-regular random graphs. However, by increasing the size of the smallest non-trivial cycle – the so-called *girth* – the remaining vulnerability can be further reduced.

5 Related Work

Constructing overlay networks to improve the DoS-resilience has been in the scientific focus for almost a decade. One of the first approaches was SOS [23], where the IP addresses of centralized servers are hidden from the attacker by routing all requests via three layers of a trusted overlay network. However, the security of the system depends on the premise that overlay nodes are trusted and cannot be compromised. In extension to this idea, Fu et al. [17] proposed a capability-based filtering at the edges and migration of services in order to evade node failures. Other means to achieve resilience are smart placement of backup resources and selection of backup paths [8,34]. Furthermore, scalable VPN can provide a resilient overlay network [5]. For special applications such as Application Layer Multicast (ALM) where no any-to-any communication is required, properties of optimal resilient overlay topologies can be shown analytically [6,4].

Strongly related to this article is the research of Wang et al. [41]. Similar to our analysis in section 4, the authors assessed several well known graphs regarding their resilience. Since they assumed an uncoordinated epidemic attacker model, their findings are quite contrary to our results. When the attack is spreading from an initial node zero over the network, the stronger connected the network, the more it deteriorates. Therefore, the CAN overlay topologies were found to be most resilient. This result was verified independently by simulations [2]. However, our results show clearly that these are quite vulnerable, if the attacker operate more intelligent.

According to our attacker model, highly connected graphs with a low constant degree and high girth, such as cages [42] are the most resilient. Furthermore, cage based network structures – named *entangled graphs* – were found to be also useful to achieve other objectives, such as low diameter or resistance against random failures [11].

6 Conclusion & Future Work

As the importance of overlay networks for VPN or peer-to-peer-based applications in the private and commercial sector grown over the last decade, they also became more and more attractive to potential attackers. With the increase in attack bandwidth provided by botnets, DoS protection is a pressing issue.

In this article we presented an attacker model that is more suitable to reflect todays threads on overlay networks than single node attacks. We proved that finding DoS attacks that induce the maximum possible damage with a given budget is an NP-hard problem and presented and BLP to calculate such attacks for small graphs. Furthermore, we were able to identify several easier to compute criteria, that correlate with the resilience against optimal attacks. In particular, a low degree of 3, high girth, low average path length, low variation in closeness, and high algebraic connectivity were found to decrease a topology's vulnerability. When regarding these metrics designing network structures, highly resilient topologies can be created.

However, these results were found assuming homogeneous observation probabilities. Therefore, our future work will focus on finding resilient topologies that contain nodes of varying observation probabilities. It is expected that this modification leads to less homogeneous graphs and even non-regular graphs. In particular, re-

lations between the distribution of observation probabilities and node degree distributions are interesting. There is also uncertainty of whether or not more resilient nodes form clusters.

Another aspect of future research will circle around scalable distributed self-organizing resilient network construction. This introduces more requirements as nodes should not have knowledge of the full topology and must not gain increasing knowledge of other's IP addresses over time. A particular challenge will be maintaining efficiency and resilience under dynamics, i.e., insertion and removal of nodes.

References

1. Bard, J.F.: Practical Bilevel Optimization: Algorithms and Applications. Kluwer Academic Publishers (1999)
2. Beitollahi, H., Deconinck, G.: Comparing Chord, CAN, and Pastry overlay networks for resistance to DoS attacks. In: CRiSIS '08, pp. 261–266 (2008)
3. Biggs, N.: Constructions for Cubic Graphs with Large Girth. Journal of Combinatorics **5** (1998)
4. Brinkmeier, M., Fischer, M., Grau, S., Schaefer, G.: Towards the design of unexploitable construction mechanisms for multiple-tree based p2p streaming systems pp. 193–204 (2009)
5. Brinkmeier, M., Rossberg, M., Schaefer, G.: Towards a Denial-of-Service Resilient Design of Complex IPsec Overlays. In: Proceedings of the IEEE International Conference on Communications (ICC) (2009)
6. Brinkmeier, M., Schafer, G., Strufe, T.: Optimally DoS Resistant P2P Topologies for Live Multimedia Streaming. IEEE Transactions on Parallel and Distributed Systems (TPDS) **20**(6), 831–844 (2009)
7. Broersma, M.: Botnet price for hourly hire on par with cost of two pints. ZDNet Article (2010)
8. Cinkler, T.: Some more aspects of resilience. Telecommunication Systems pp. 1–22 (2011)
9. Coxeter, H.S.M.: Self-dual configurations and regular graphs. Bulletin of the American Mathematical Society **56**, 413–455 (1950)
10. Dempe, S.: Foundations of Bilevel Programming. Kluwer Academic Publishers (2002)
11. Donetti, L., Hurtado, P.I., Muñoz, M.A.: Entangled networks, synchronization, and optimal network topology. Physical Review Letters **95**, 188,701 (2005)
12. Douligeris, C., Mitrokotsa, A.: DDoS attacks and defense mechanisms: classification and state-of-the-art. Computer Networks **44**(5), 643–666 (2004)
13. Dunigan, T.: Performance of the Intel iPSC/860 and Ncube 6400 hypercubes. Parallel Computing **17**(10–11), 1285–1302 (1991)
14. European Energy Exchange AG: EEX - Technical Connection Alternatives. White Paper (2011)
15. Fisher, D.: Attack Tool Released to Exploit SSL DoS Issue. Internet Publication (2011)
16. Forrester Consulting: The Trends And Changing Landscape Of DDoS Threats And Protection (2009)
17. Fu, X., Crowcroft, J.: GONE: an Infrastructure Overlay for Resilient, DoS-Limiting Networking. In: Proceedings of ACM Network and Operating System Support for Digital Audio and Video (NOSSDAV) (2006)
18. Gamer, T., Mayer, C.P.: Large-scale Evaluation of Distributed Attack Detection. In: Proceedings of the 2nd International Conference on Simulation Tools and Techniques, Proceedigs of Simutools '09 (2009)
19. Gruska, J.: Foundations of Computing. ITCP Computer Science Series. Thomson International Computer Press (1997)
20. Gurobi Optimization, I.: Gurobi optimizer reference manual (2012). URL `http://www.gurobi.com`
21. Karp, R.M.: Reducibility Among Combinatorial Problems. In: R.E. Miller, J.W. Thatcher (eds.) Complexity of Computer Computations, pp. 85–103. Plenum Press (1972)
22. Kastenberg, J.E., Korns, S.W.: Georgia's cyber left hook. Parameters **38**(4), 60–76 (2008)
23. Keromytis, A.D., Misra, V., Rubenstein, D.: SOS: An Architecture for Mitigating DDoS Attacks. IEEE Journal on Selected Areas in Communications (JSAC) **22**, 176–188 (2004)
24. Kuzmanovic, A., Knightly, E.W.: Low-rate tcp-targeted denial of service attacks: the shrew vs. the mice and elephants. In: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications, SIGCOMM '03, pp. 75–86. ACM (2003)
25. Lesk, M.: The New Front Line: Estonia under Cyberassault. IEEE Security & Privacy **5**(4), 76–79 (2007)
26. Li, J., Dabek, F.: F2F: reliable storage in open networks. In: Proceedings of 5th International Workshop on Peer-to-Peer Systems (IPTPS) (2006)
27. Maymounkov, P., Mazieres, D.: Kademlia: A peer-to-peer information system based on the xor metric. In: Proceedings of 1st International Workshop on Peer-to-Peer Systems (IPTPS) (2002)
28. McAfee: In the Dark: Crucial Industries Confront Cyberattacks. White Paper (2011)
29. Meringer, M.: Fast generation of regular graphs and construction of cages. J. Graph Theory **30**, 137–146 (1999)
30. Mohar, B.: The laplacian spectrum of graphs **2**, 871–898 (1991)
31. Ollmann, G.: Want to rent an 80-120k DDoS Botnet? (2009). URL `http://blog.damballa.com/?p=330`
32. Preparata, F.P., Vuillemin, J.: The cube-connected cycles: a versatile network for parallel computation. Communications of the ACM **24**(5), 300–309 (1981)
33. Raghunath, S., Ramakrishnan, K.K., Kalyanaraman, S., Chase, C.: Measurement based characterization and provisioning of IP VPNs. In: ACM SIGCOMM, pp. 342–355 (2004)
34. Rak, J., Walkowiak, K.: Reliable anycast and unicast routing: protection against attacks. Telecommunication Systems pp. 1–18 (2011)
35. Ratnasamy, S., Francis, P., Handley, M., Karp, R., Shenker, S.: A Scalable Content-Addressable Network. In: Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM), pp. 161–172 (2001)
36. Ross, S.M.: Probability Models for Computer Science. Academic Press (2001)
37. Rossberg, M., Girlich, F., Schaefer, G.: Analyzing and Improving the Resistance of Overlay-Networks against Bandwidth Exhaustion. In: Proceedings of the 4th International Workshop on Reliable Networks Design and Modeling (RNDM) (2012)
38. Rossberg, M., Schaefer, G.: A Survey on Automatic Configuration of Virtual Private Networks. Computer Networks **55**, 1684–1699 (2011)

39. Sterbenz, J.P., Çetinkaya, E.K., Hameed, M.A., Jabbar, A., Shi, Q., Rohrer, J.P.: Evaluation of Network Resilience, Survivability, and Disruption Tolerance: Analysis, Topology Generation, Simulation, and Experimentation. Telecommunication Systems (2011)
40. Vajanapoom, K., Tipper, D., Akavipat, S.: Risk Based Resilient Network Design. Telecommunication Systems (2011)
41. Wang, J., Lu, L., Chien, A.A.: Tolerating Denial-of-Service Attacks Using Overlay Networks - Impact of Topology. In: ACM workshop on Survivable and self-regenerative systems (SSRS), pp. 43–52 (2003)
42. Wong, P.K.: Cages — A Survey. Journal of Graph Theory **6**(1), 1–22 (1982)
43. Wormald, N.C.: Models of Random Regular Graphs. Surveys in Combinatorics **267**, 239–298 (1999)

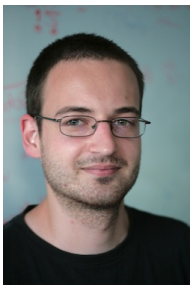

Franz Girlich received his master degree in computer science in 2011 at the Technische Universität Ilmenau, Germany, where he subsequently joined the Telematics/Computer Networks Research Group. His interests lay in the area of network robustness and protection of communication infrastructures.

Michael Rossberg obtained his computer science diploma and Ph.D. at Technische Universität Ilmenau in 2007 and 2011 respectively. He is part of the Telematics/Computer Networks Research Group and researching in the fields of network security especially the DoS-resistance of communication infrastructures.

Guenter Schaefer holds a diploma (1994) and Ph.D. (1998) in Computer Science from the University of Karlsruhe, Germany. He worked as a postdoctoral researcher at the Ecole Nationale Superieure des Telecommunications in Paris, France, between February 1999 and July 2000, focusing on network security and access network performance of third-generation mobile communication networks. From August 2000 to March 2005, he worked at the Technische Universität Berlin, Germany in the areas of network security and advanced mobile communication architectures and services. In April 2005, he became a full professor of telecommunications/computer networking at the Technische Universität Ilmenau, Germany. His main research interests are communication protocols and architectures, network security and protection of communication infrastructures.