Delft University of Technology

A Moment of Weakness

Protecting Against Targeted Attacks Following a Natural Disaster

Oostenbrink, Jorik; Kuipers, Fernando

DOI

[10.1145/3397776.3397780](https://doi.org/10.1145/3397776.3397780)

Publication date

2020

Document Version

Final published version

Published in

ACM SIGMETRICS Performance Evaluation Review

Citation (APA)

Oostenbrink, J., & Kuipers, F. (2020). A Moment of Weakness: Protecting Against Targeted Attacks Following a Natural Disaster. *ACM SIGMETRICS Performance Evaluation Review*, 47(4), 12-15.
<https://doi.org/10.1145/3397776.3397780>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

A Moment of Weakness: Protecting Against Targeted Attacks Following a Natural Disaster

Jorik Oostenbrink
Delft University of Technology
J.Oostenbrink@tudelft.nl

Fernando A. Kuipers
Delft University of Technology
F.A.Kuipers@tudelft.nl

ABSTRACT

By targeting communication and power networks, malicious actors can significantly disrupt our society. As networks are more vulnerable after a natural disaster, this moment of weakness may be exploited to disrupt the network even further. However, the potential impact and mitigation of such a follow-up attack has yet to be studied.

In this paper, we propose a framework to analyze the impact of a combination of a natural disaster followed by a targeted single node failure. We apply this framework on empirical disaster data and two network topologies. Our experiments show that even small targeted attacks can significantly augment the already grave network disruption caused by a natural disaster. We further show that this effect can be mitigated by adopting a calculated repair strategy.

1. INTRODUCTION

Communication and power networks are critical to our society. This makes them a prime target for malicious actors trying to destabilize or terrorize a country. In fact, Admiral Mike Rogers, the former director of the NSA, warned, “It’s only a matter of the when, not the if, you are going to see a nation state, a group or an actor engage in destructive behavior against critical infrastructure of the United States” [1].

For many of these actors, the exact timing of their attack may not be essential. Their focus is to deal as much damage as possible, preferably using only a small amount of resources. A strategy they might adopt is to delay their attack until the network is most vulnerable, such as after a natural disaster. By attacking the network at its weakest moment, a bad actor can multiply the damage he, or the disaster by itself, could otherwise inflict.

As it takes time to prepare and execute an attack, a network operator has a limited window to try to reduce the impact of any potential attack. However, to the best of our knowledge, while there is a large body of research on the resilience of networks to natural disasters (see [4]) and targeted attacks (e.g. [2, 5, 8]), the potential combination of a disaster followed by a targeted attack has yet to be studied.

In this paper, we propose a framework to analyze the impact of a combination of a natural disaster and targeted single node failure¹. Our main contributions are as follows:

¹Be it through physical means or cyber attacks.

- We extend our successive disaster framework from [7] to incorporate *targeted* attacks.
- We apply our framework to empirical disaster data and show that a small follow-up attack can significantly increase the impact of a natural disaster. In addition, we study the effect of changing the repair strategy to prepare for potential follow-up attacks.

2. FRAMEWORK

In [7], we have introduced a model and framework for assessing the resilience of networks to successive disasters, taking into account network recovery. In this paper, we extend this framework to include the risk of targeted attacks.

We model the network as a directed multigraph $G = (V, E, \psi)$ with nodes $v \in V$ connected by links $e \in E$, where $\psi : E \rightarrow V \times V$ and $e \in E$ connects v_1 to v_2 if and only if $\psi(e) = (v_1, v_2)$. We define the *network state* s of this network by its failures: network component $c \in V \cup E$ is functioning if and only if $c \notin s$.

Now, to be able to assess the resilience of this network to follow-up attacks, we first need to model the impact of a disaster itself. We assume disaster occurrences are Poissonian, and we are given a multiset D^* of disaster processes $d = (a_d, \lambda_d)$, where $a_d \subseteq V \cup E$ are the components affected by d and λ_d is the rate of d . Thus, if disaster $d \in D^*$ occurs at time t , when the network state is s , the new network state at time t will be $s \cup a_d$. As the combination of multiple Poisson processes is itself Poissonian, the disaster processes in D^* can be combined as follows:

$$D = \{(a_d, \lambda_d) | a_d \neq \emptyset \wedge \lambda_d = \sum_{(a_d, \lambda_d^*) \in D^*} \lambda_d^* > 0\} \quad (1)$$

In this paper, we only consider attacks after a single disaster. In other words, we assume a single disaster occurs and is then followed by an attack on the network. However, by applying the techniques of [7], our model can be easily generalized to capture an attack after an arbitrary number of successive disasters or any other mix of natural disasters and attacks. We fix the time of the initial disaster (D_1) at $T_1 = 0$. Now, we can compute the distribution of the network state at T_1 , S_1 , by

$$P(S_1 = s) = \sum_{d \in D | a_d = s} P(D_1 = d) = \sum_{d \in D | a_d = s} \frac{\lambda_d}{\lambda_D} \quad (2)$$

where λ_D is $\sum_{d \in D} \lambda_d$.

A follow-up attack after a disaster can be pre-planned or opportunistic. In either case, it will take some time to react

to the disaster and execute the attack. We consider two different attack models: (1) the attack occurs after a fixed amount of time t_{attack} , and (2) the time between the disaster and attack is exponentially distributed with rate λ_{attack} . In both cases, if the network has been fully repaired before the attack has been executed, we assume the attack will be canceled and the network will not suffer any further damage.

Let T_{attack} be the time of the attack. We assume the attacker has perfect knowledge of the network at all times, and will always take down the node that maximizes the number of disconnected node pairs at T_{attack} . In other words, an attack is modeled as a worst-case node failure.

The target and impact of this attack greatly depend on the progress of network repair at T_{attack} . We consider a deterministic repair model. That is, we assume that, given a certain starting state, the recovery of the network is fixed (until the attack occurs). For each possible starting state s , we define a repair function $r_s : \mathbb{R}^+ \rightarrow V \cup E$. $r(t)_s \in V \cup E$ is the state of the network at time $t \leq T_{\text{attack}}$, given that the state of the network after being struck by the initial disaster was $S_1 = s$. Thus, the state of the network just before the attack is $r_{S_1}(T_{\text{attack}})$.

Let $M(s)$ be the number of connected node pairs in network state s and let $R_s := \min\{t \geq 0 | r(t)_s = \emptyset\}$ be the time it takes to fully repair the network (assuming no attack occurred beforehand). Given M , we can consider the follow-up attack as a function $\text{att} : V \cup E \rightarrow V \cup E$ from the network state just before the attack to the network state just after the attack:

$$\text{att}(s) = \begin{cases} \emptyset & \text{if } s = \emptyset \\ s \cup \underset{v \in V}{\text{argmin}} M(s \cup v) & \text{otherwise} \end{cases} \quad (3)$$

By combining our disaster, repair, and attack models, we can now directly compute the distribution of the state S_{attack} of the network just after the attack. In the fixed attack time case, the distribution of S_{attack} is given by

$$P(S_{\text{attack}} = s) = \sum_{d \in D | s = \text{att}(r_{a_d}(t_{\text{attack}}))} \frac{\lambda_d}{\lambda_D} \quad (4)$$

while in the random attack time case the distribution of S_{attack} is given by

$$\begin{aligned} P(S_{\text{attack}} = s) &= \sum_{d \in D} \frac{\lambda_d}{\lambda_D} P(S_{\text{attack}} = s | D_1 = d) \\ &= \sum_{d \in D} \frac{\lambda_d}{\lambda_D} (\exp(-\lambda_{\text{attack}} \min(\mathcal{M}_{a_d,s}, R_{a_d})) \\ &\quad - \exp(-\lambda_{\text{attack}} \min(\mathcal{S}_{a_d,s}, R_{a_d}))) \end{aligned} \quad (5)$$

where $[\mathcal{M}_{a_d,s}, \mathcal{S}_{a_d,s})$ is the period of time during which an attack would result in network state² s .

Given the distribution of S_{attack} , we can directly compute the distribution of any performance metric after the follow-up attack, such as the number of remaining connections, $M(S_{\text{attack}})$. In addition, our framework allows network operators to assess the impact of different repair strategies or network configurations by simply exchanging repair functions or modifying the initial network.

² $\mathcal{M}_{a_d,s}$ is the first time t at which $\text{att}(r_{a_d}(t)) = s$ (or ∞ if no such time exists), and $\mathcal{S}_{a_d,s}$ is the first time t after $\mathcal{M}_{a_d,s}$ at which $\text{att}(r_{a_d}(t)) \neq s$ (or ∞).

3. EXPERIMENTS

In this section, we apply our framework to two slightly modified³ versions of undirected networks from the topology zoo [6]: Sinet and Deltacom. Sinet is a Japanese network of 47 nodes connected by 49 links, and Deltacom is a US network of 99 nodes connected by 151 links.

We make use of the same disaster set as was used in [7]: a set of earthquake scenarios and historical tropical cyclones. For Sinet, we consider both types of disasters, while for Deltacom, we only consider tropical cyclones. We assume only network nodes are affected by these disasters and all network links remain functioning. This gives us a yearly disaster rate λ_D of 1.597 for Sinet and 1.342 for Deltacom.

For ease of reading, we make the assumption that one node is repaired every day. However, by scaling both the attack and repair time, our results can easily be transformed to any other repair time.

We use the Average Two-Terminal Reliability (ATTR) as an impact measure. The ATTR is defined as the number of connected node pairs divided by the total number of node pairs in the network. That is, $\text{ATTR} = \frac{M}{|V|^2}$.

3.1 Impact

We first compare the impact of follow-up attacks to those of a disaster or attack by itself. We assume both networks use a greedy repair function that continuously chooses the node with the largest impact on ATTR to repair. To make a fair comparison, we modify the attack function att by continuing the follow-up attack even if the network has been fully repaired.

Figure 1 shows the expected ATTR after a targeted attack, disaster, or disaster and follow-up attack. Sinet is clearly more vulnerable to both targeted attacks and disasters than Deltacom. However, for both networks, a follow-up attack can significantly increase the impact of a disaster. For Sinet, the combination of disaster and follow-up attack disconnects more than half of all node pairs on average.

We have assumed that a malicious actor strikes after the *first* natural disaster that hits the network. On average, such an opportunity occurs more than once per year. However, he could also decide to wait for a larger disaster, which would allow him to inflict even more damage to the network. We consider an attacker that waits for a disaster that damages at least 5 (10) of Deltacom's (Sinet's) nodes⁴. Figure 2 shows the impact of this more patient follow-up attack. Waiting for larger events allows the attacker to inflict much more damage. In the case of Sinet, the expected impact of the follow-up attack is the disconnection of around half of the *remaining* node pairs.

3.2 Repair Strategies

After a disaster, the network operator will typically try to restore as much functionality as quickly as possible. The impact of the follow-up attack greatly depends on the progress of these repair operations at the time of the attack. Although speeding up repair would have the largest effect, the network operator can also change the order in which components are repaired to try to minimize the impact of any

³We have removed all nodes without a geographical location or with degree 0

⁴An opportunity that occurs around once every 2 years on average

(a) Deltacom

(b) Sinet

Figure 1: Impact of a follow-up attack. The disaster occurs at $t=0$, and is followed by a targeted attack after 0 to 16 days (even if the network is already repaired). One node is repaired every day.

(a) Deltacom

(b) Sinet

Figure 2: Impact of a follow-up attack on Deltacom (Sinet) if the attacker waits for a disaster that damages at least 5 (10) nodes. The disaster occurs at $t=0$, and is followed by a targeted attack after 0 to 16 days (even if the network is already repaired). One node is repaired every day.

(a) Deltacom

(b) Sinet

Figure 3: Cumulative Distribution function of the ATTR after the follow-up attack for different repair strategies. The disaster occurs at $t=0$, and is followed by a targeted attack after 3 days. If the network is repaired before the attack, the attack is canceled (and $ATTR = 1$).

attacks. While this might lower the speed at which network functionality is restored, it could be a worthy trade-off if the network is under threat.

In this section, we consider the effect of changing the node repair order on the impact of the follow-up attack. We compare 5 different repair strategies:

- *Betweenness Repair*: Repair nodes in the order of their betweenness centrality [3].

- *Greedy Repair*: Every day, repair the node with the highest impact on the ATTR.
- *Optimal Unprepared Repair Strategy*: Maximizes the ATTR at the time of attack.
- *Greedy Prepared Repair*: Every day, repair the node that would increase the ATTR the most if the network would be attacked immediately afterwards.

Figure 4: Cumulative Distribution function of the ATTR 3 days after the initial disaster (without follow-up attack) for different repair strategies.

- *Optimal Prepared Repair Strategy*: Maximizes the ATTR immediately after the attack.

Figure 3 shows the Cumulative Distribution Function (CDF) of the ATTR after a follow-up attack with a fixed t_{attack} of 3 days. The large spikes at $\text{ATTR} = 1$ show the probability of completely repairing the network within 3 days (0.340 for Sinet and 0.588 for Deltacom). In these cases, the attack is canceled and the repair strategy has no impact. However, in most other cases the repair strategy does significantly impact the ATTR after a follow-up attack. In particular, there is a large gap between the strategies that try to maximize the ATTR by itself compared to those that try to maximize the ATTR after the attack. Experiments on randomly delayed follow-up attacks show very similar results to those of fixed-time attacks (figures not included due to space limitations).

Knowing that changing the order of repair can reduce the impact of follow-up attacks, one might wonder what the impact of these prepared repair strategies is on the performance of the network during repair without an attack. Or in other words, what does preparing for a follow-up attack cost us if no such attack occurs? Figure 4 shows the CDF of the ATTR 3 days after the initial disaster (without any follow-up attack). For the considered networks, the difference between the different repair strategies is extremely small and even the prepared strategies perform close to optimally.

4. CONCLUSION

Critical infrastructure networks are prime targets for malicious actors trying to destabilize or terrorize a country. As part of their attack strategy, they might wait for points in time when critical infrastructure is significantly more vulnerable, for example right after a natural disaster has struck. However, current disaster vulnerability frameworks do not consider the potential risk of a network to these potential follow-up attacks.

We have proposed a framework for assessing the impact of follow-up attacks. Our framework can take into account a variety of natural disasters and two kinds of attacks: a worst-case node failure after (1) a fixed amount of time or (2) exponentially distributed random delay after an initial disaster.

In our experiments, we have shown that small targeted attacks can significantly augment the impact that a natural disaster has on the network. Fortunately, our results also reveal that the right choice of repair strategy allows network operators to reduce the threat of follow-up attacks at almost no cost to network performance compared to other repair strategies. Our framework aids in determining the efficacy of repair strategies.

5. REFERENCES

- [1] NSA chief worries about cyber attack on US infrastructure. <https://www.securityweek.com/nsa-chief-worries-about-cyber-attack-us-infrastructure>, March 2016. Accessed: 18-04-2019.
- [2] R. Cohen, K. Erez, D. Ben-Avraham, and S. Havlin. Breakdown of the internet under intentional attack. *Physical review letters*, 86(16):3682, 2001.
- [3] L. C. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, pages 35–41, 1977.
- [4] T. Gomes, J. Tapolcai, C. Esposito, D. Hutchison, F. Kuipers, J. Rak, A. De Sousa, A. Iossifides, R. Travanca, J. André, et al. A survey of strategies for communication networks to protect against large-scale natural disasters. In *2016 8th international workshop on resilient networks design and modeling (RNDM)*, pages 11–22. IEEE, 2016.
- [5] S. Iyer, T. Killingback, B. Sundaram, and Z. Wang. Attack robustness and centrality of complex networks. *PloS one*, 8(4):e59613, 2013.
- [6] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan. The internet topology zoo. *Selected Areas in Communications, IEEE Journal on*, 29(9):1765–1775, october 2011.
- [7] J. Oostenbrink and F. A. Kuipers. The risk of successive disasters: A blow-by-blow network vulnerability analysis. In *2019 IFIP Networking Conference (to appear)*. IEEE, 2019.
- [8] C. M. Schneider, A. A. Moreira, J. S. Andrade, S. Havlin, and H. J. Herrmann. Mitigation of malicious attacks on networks. *Proceedings of the National Academy of Sciences*, 108(10):3838–3841, 2011.