Telling faults from cyber-attacks in a multi-modal logistic system with complex network analysis

Dario Guidotti, Giuseppe Cicala, Tommaso Gili, Armando Tacchella

KEYWORDS

Cyber-Security and Critical Infrastructure Protection, Complex Networks, Discrete Event Simulation.

ABSTRACT

We investigate the properties of systems of systems in a cybersecurity context by using complex network methodologies. We are interested in *resilience* and *attribution*. The first relates to the system's behavior in case of faults/attacks, namely to its capacity to recover full or partial functionality after a fault/attack. The second corresponds to the capability to tell faults from attacks, namely to trace the cause of an observed malfunction back to its originating cause(s). We present experiments to witness the effectiveness of our methodology considering a discrete event simulation of a multimodal logistic network featuring 40 nodes distributed across Italy and daily traffic roughly corresponding to the number of containers shipped through in Italian ports yearly averaged daily.

INTRODUCTION

Complex networks are significantly present in many science disciplines and have recently received much attention [Bar]. Many studies have been devoted to measuring networks' robustness against attacks or random degradation failures causing deletion of nodes or connections. Such measures are used to increase the security of complex systems and possibly to improve their robustness [AJB00], [KG14]. Edges of a network usually play the role of transmitting information or load and maintaining network connectivity. The load model of cascading failures can be used to investigate small-world network performance subject to deliberate attacks on node and edge. Results show that edge attacks produce more significant cascading failures than node attacks. On the other hand, in real-world networks, the nodes vulnerable to attack are often well protected, while edges are a relatively easy target for attackers [NGZL15].

Systems of systems, e.g., water treatment plants, electric grids (power plants and associated distribution networks), industrial plants, transportation networks, and smart homes, are the ideal field of application for complex network theory to obtain useful insights about the behavior of the systems under scrutiny. In such systems, wireless communication among components and external network access for supervisory control and data acquisition (SCADA) make them an ideal target for cyber-attacks. It has been demonstrated that malicious users can gain control of such systems and/or disrupt their functionality severely [FR11]. This is also true for systems that are part of critical national infrastructure (CI). As such, intentional or accidental incidents that alter their normal behavior can have dramatic effects on the safety of citizens [WFD10].

Among other security-related issues, resilience is recognized as one of the keys to understanding how much damage can be brought to a system and its surrounding environment in case of a successful cyber-attack [DRKS08]. The concept of resilience — defined as "*the quality of being able to return quickly to a previous good condition after problems*" — emerges as an additional target, complementary to protection from external threats, but not subordinate to it. More recently, the term *cyber-resilience* has been coined to identify specifically "*the ability to continuously deliver the intended outcome despite adverse cyber events*" [BHSZ15], and this is the interpretation we consider in this paper, where we are interested in applying complex networks analysis to obtain a measure of resilience.

Our research goal is to discriminate whether a random fault or a cyber-attack causes the performance degradation a system incurs into and the amount of such degradation. We call this *attribution* and we hypothesize that it relates strongly to the ability to trace the cause of an observed malfunction back to its cause(s). We need to stress that we are not interested in the specific originating event but rather differentiating system-related events from cyber-related events. We find that a clear answer to this question may be the solid basis for any attribution process targeted to spot attackers.

We implemented complex network metrics on a realistic multi-modal logistic system. We embedded a hypothetical network into the Italian railway system, providing coverage of the entire national territory using 40 terminals (nodes) so that each one serves an area of approximately 150Km in radius. Simulation parameters – e.g., number of trains with their schedules and routes, number of containers with their origins and final destinations – are chosen according to stochastic models. Events generated by the simulator are stored (i) in a database, and basic KPIs can be computed out of these data. The user can inject both faults and attacks in the simulation, so that their effects can be observed in the results. Simulations tested our methodology's effectiveness considering daily traffic scenarios approximately corresponding to the number of containers shipped through Italian ports yearly.

Results clearly show that complex network analysis enables the assessment of cyber-resilience and gives us indications to understand whether a system's performance degra-

Dario Guidotti, Giuseppe Cicala and Armando Tacchella are with "Dipartimento di Informatica, Bioingegneria, Robotica e Ingegneria dei Sistemi" (DIBRIS), University of Genoa, Viale Causa 13, 16145 Genoa, Italy. E-mail: dario.guidotti@edu.unige.it, giuseppe.cicala@unige.it, armando.tacchella@unige.it. Tommaso Gili is with IMT Lucca ... E-mail: tommaso.gili@imtlucca.it. The authors wish to thank ... The corresponding author is Armando Tacchella.

dation is due to a fault or some malicious activity.

BACKGROUND

In our methodology, we consider different topological measures from undirected graphs.

Definition 1 (Graph) A Graph is a pair $G = (V, E)$ where V is a set whose elements are called vertices and E is a set of paired *vertices*, whose elements are called *edges*. We briefly present our measures of interest in the following. The first measure we consider is Laplacian Energy, i.e., the sum of the absolute values of the eigenvalues of the Laplacian matrix of the graph. This quantity is often studied in the context of spectral graph theory and chemistry studies.

Definition 2 (Laplacian Energy) [GZ06] Let G be a graph with n vertices and no loops or parallel edges. Let L be the Laplacian matrix of G and μ_i, $i = 1, ..., n$ the eigenvalues of L. Then the Laplacian energy of the graph is defined as:

$$E(G) = \sum_{i=1}^{n} \mu_i^2 \tag{1}$$

Another measure of interest for graphs, in general, is the centrality of vertices. In this work, we have chosen to consider Laplacian Centrality [QFW+12] and Betweenness Centrality [Fre77]. We use these measures to understand the relevance of the nodes and edges of the graph.

Definition 3 (Betweenness Centrality) Let G be a graph with n vertices and no loops or parallel edges. The Betweenness centrality of the vertex i is defined by the number of shortest paths that pass through i. Precisely, let L_{hj} be the total number of shortest paths from a vertex h to another vertex j and $L_{hj}(i)$ be the number of shortest paths that pass through the vertex i. The Betweenness centrality of vertex i can be defined as

$$\frac{2}{(n-1)(n-2)} \sum_{h \neq i} \sum_{j \neq i, j \neq h} \frac{L_{hj}(i)}{L_{hj}} \tag{2}$$

Definition 4 (Laplacian Centrality) Let G be a graph with n vertices and no loops or parallel edges. Let $E_L(G)$ be the laplacian energy of G and $E_L(G_i)$ the laplacian energy of G after the vertex i has been removed. The laplacian centrality of vertex i is defined as

$$\frac{E_L(G) - E_L(G_i)}{E_L(G)} \tag{3}$$

Definition 5 (Community Structure and Communities) A graph is said to have a community structure if the graph's nodes can be easily grouped into sets of nodes such that each set of nodes is densely connected internally: each set of nodes is a community. One of the reasons for the importance of communities is that they often present very different properties than the average properties of the corresponding network, therefore concentrating only on the average property usually misses important and interesting features of the network. In this work, we consider the communities generated using the Louvain algorithm [BGLL08].

Definition 6 (Giant Component) It is the largest connected component of a given graph that contains a finite fraction of the vertices. We partitioned the graph into several connected components by removing the least important edges according to a percolation approach [LLL+21].

Fig. 1: Graphical representation of ONTOMIL network.

Edges removal stops when the difference of the two largest connected components' size is less or equal to a specific value.

ONTOMIL SIMULATOR

The simulator models an *Intermodal Logistics System* (ILS) which support receiving, storing and shipping goods packaged in *Intermodal Transport Units* (ITUs, also known as "containers"). A detailed description of the context is provided in [CCT13]. Here we restrict our attention to ILSs wherein rail transportation is supported by a network of terminals equipped with systems for fast ITU handling. The overall network is "covered" by relatively frequent short-distance trains with a fixed composition and a predefined daily schedule. ITUs enter the network at some terminal and travel to their destination according to a predefined route, usually boarding more than one train along the way. While this solution enables efficient utilization of resources, information technology is vital to operate it effectively.

Operation of the simulated ILS involves several customers forwarding their goods through the system and an handling agent, i.e., the business responsible for managing the entire network. Given its role, the handling agent is also the main stake-holder, and the one who is thought to collect *key performance indicators* (KPIs) to be computed on data about the system. Transportation across the network is organized by having customers emit *requests for work* which contain ITUs to be sent from a given terminal to other destinations on the network. The handling agent associates to each request for work a number of *transport orders*, one for each ITU listed in the request for work. The transport order contains all the data related to the shipping, like ITU route through the network and expected time of delivery. Once the ITU corresponding to a given transport order is collected at a terminal, it is boarded on the first outgoing train whose

destination is compatible with its route. Since trains travel across relatively short distances, it is possible to dispatch ITUs more than once during a 24 hours time-span.

The main activity of the simulated ILS is to satisfy the supplied demand of transportation in a timely way in spite of events potentially disrupting the service like, either due to natural causes, e.g., network and rolling stock failures, or due to malicious activity, e.g., cyber-attacks targeting single terminals or the whole network. We describe how such events are injected in the simulator later on as part of our methodological approach, but here we observe that monitoring ITUs from the departure terminal to their final destination is a key enabler for every kind of analysis on the network. In particular, the data obtained through monitoring enables the computation of KPIs which summarize the overall status of the system and its ability to handle a given workload over time. In particular, ONTOMIL provides the following "standard" indicators:

1. Late transport orders on a daily basis, i.e., the number of transport orders issued on a given date whose ITUs did not reach the final destination on the same date.
2. Cumulative number of ITUs handled in terminals.
3. Average number of ITUs unloaded per hour in the network terminals.
4. Late trains, i.e., trains suffering one or more delays with respect to their schedule on a specific route.
5. Recent sink/source terminals, i.e., terminals wherein the only operations were loadings or unloadings over the past hour.
6. Number of customers whose request for work contains transport orders backlogged for more than two days – calculation done on a daily basis.
7. Number of loading and unloading operations for each terminal on an hourly basis.
8. High-activity customers on a daily basis, i.e., those customers shipping more ITUs than a given daily threshold.
9. Average train utilization on an hourly basis, i.e., number of ITUs vs. number of trains travelling across the network.
10. Route utilization, i.e., cumulative number of ITUs which traveled along a given route.

The above KPIs can be grouped in different categories. For instance, KPI 2 and 9 are considered *critical success factors*, in the sense that they highlight potential flaws in network organization which should be corrected to maintain efficiency, e.g., wrong train scheduling and routing. Most of the remaining KPIs are so called *dashboard indicators*, in the sense that they provide useful information to quantify the overall health status of the network, and can support tactical decision making. Change in some of the dashboard indicators impacts on the ability of the whole system to generate revenues for the handling agent.

The actual simulated model consists of a hypothetical logistic network covering the entire Italian territory using 40 intermodal terminals connected between them through railroads. A representation of such network is given in Figure 1. Albeit ONTOMIL simulates an ideal system, the railroad connection correspond to the actual freight lines along which goods are forwarded by train. As shown in Figure 2, ONTOMIL enables the customization of different parameters, including the number of simulation days, the minimum

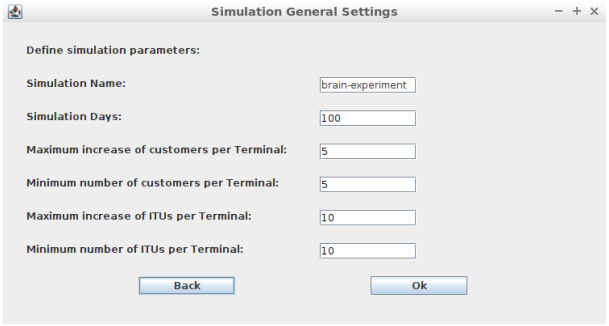

Fig. 2: Graphical Interface for the selection of the parameters of the ONTOMIL simulation.

number of customers insisting on any terminal on each single day and others that affect the number of ITUs in the system. The actual values are chosen randomly for each terminal before the beginning of each day of simulation respecting the bounds defined by the user. In our simulations we considered a number of customers and ITUs that results in a number of ITUs circulating in the network which is comparable to the number of ITUs handled daily by Italian ports and forwarded on railway trains.

METHODOLOGY

Analyzing the ILS

To analyze the ILS we abstract it as two different kinds of graph and we analyze how they change in different temporal intervals. The terminals correspond to the vertexes of the graphs whereas the connections between the terminals correspond to the edges of the graphs. The first graph we considered is the Flux Graph (FG) whose weights are the number of ITUs present on the corresponding connection during the chosen interval of time. The second graph is the Difference Flux Graph (DFG) whose weights are the difference in absolute value between the number of ITUs present on the corresponding connection during the chosen interval of time and the number of ITUs present on it during the previous interval of time. In this work we have chosen a single day as the interval of time of interest given the characteristic of the simulation. The idea behind the FG is to provide a snapshot of the performance of the ILS during a particular day, whereas the idea behind the DFG is to provide a snapshot of the evolution of the ILS between a particular day and the next.

Fault and attack injection

In order to test the proposed methodology to analyze the ILS, we added on top ofthe ONTOMIL simulator the capability of injecting faults and attacks. We think of the former as naturally occurring events, e.g., delay along a line due to a locomotive malfunction, and the latter as the result of a malicious activity, e.g., an hacker infiltrating the shipping network and altering the transport orders. The capability to inject faults and attacks, alone or combined, is crucial to demonstrate the effectiveness of the methodology proposed. As shown in Figure 3, ONTOMIL currently supports the injection of two kinds of anomalies:

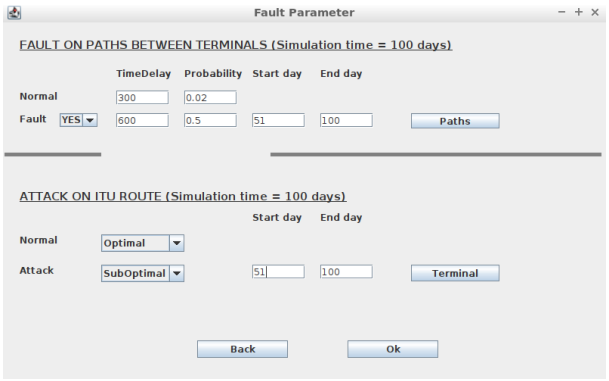

Fig. 3: Graphical Interface for the control of faults and attacks in the ONTOMIL simulation.

- Fault on paths between terminals, where the user can tune the "physiological" delay that affects trains traveling across all routes with a given probability ("Normal" time delay and related probability), and introduce a "pathological" delay which is larger in magnitude and probability of occurrence and it is meant to affect specific paths — including the possibility to target all the paths.
- Attach on ITU route, where the user can decide to change the normal routing of ITUs through the network, based on the minimum number of terminals hops, to some anomalous routing that could be just suboptimal or plain wrong, e.g., an ITU that should go from terminal A to terminal B is sent to a terminal C where no connection to B exists; the attack can affect all the terminals or just specific one, again based on user's choice.

It is important to observe that these anomalies do not cover all potential incidents due to natural causes or malicious attacks that can happen in a network like the one simulated by ONTOMIL. However, train delays represent a frequent event and alteration of transport orders is the easiest way that an hacker has to alter the normal behavior of the network and cause disruption in service. Also, they represent two fundamentally different ways to cause such a disruption, one that relates to the physical nature of the process, the other that relates to the control of the same. In principle, further anomalies can be injected into the ONTOMIL simulator, but these will be the subject of future work.

EXPERIMENTAL EVALUATION

The results of our experiments can be seen in Tables I and II. For each measure of interest, we have computed the p-value with the Wilcoxon Signed-Rank test and the Cohen's d coefficient, comparing the samples collected during normal operation of the system and the ones collected during the fault/attack. The goal is to understand whether the distribution of a specific measure differs, in a statistically significant way, when considering normal operation and fault/attack injection. Given the results, it is clear that some of the measures of interest present distributions which have this characteristic, because they are significantly different during the normal operations and the fault/attack and/or they are different based on the event injected (fault or attack). In the tables, we have highlighted in green all the p-values smaller than 10^{-4} and all the Cohen's d coefficients greater than 1.

As we would expect, the simulations without any attack or fault do not present any statistically significant difference between the various measures. However, when the network is under attack, the Louvain Energy measures computed over the Flux Graphs present a significant difference. Moreover, when the attack is applied to the high importance terminals also the Giant Energy measure present a comparable difference. The same phenomenon occurs for the high importance fault simulations regarding the Louvain Energy measure computed over the Difference Flux Graph. Regarding the simulation in which both the attack and faults occurs on the high relevance terminals and routes, all the measures except the Giant Energy computed over the DFGs present a significant difference.

Given the observations above we can define a set of rules (represented as a decision tree in Figure 4) based on the statistical significance of the difference of the measures computed on the data sampled during different time-windows of the simulations. In particular, it appears clear that an attack can be identified using the Louvain Energy or Giant Energy measures computed on the FGs, a fault can be identified using the Louvain Energy measure computed on the DFGs and the presence of both fault and attack can be identified by the significance of both Louvain Energy computed on FGs and DFGs at the same time. In general, variations of the Louvain Energy computed on the FGs pinpoint attacks both on low and high importance Terminals, whereas variations of the Louvain Energy computed on the DFGs pinpoint faults only on the high importance routes.

To understand the motivations behind the general inability to identify faults on low importance routes we must refer to Images 5, 6 and 7 in which we show the trend of two KPIs of interest during a simulation in which the system was under attack, one in which it was experiencing a fault and one in which it was experiencing both. As it can be seen, the performances of the system clearly deviates from normal conditions either under attack and under both attack and fault, and this happens both when their intensity is low and when it is high. On the other hand, when only the low intensity fault is applied, the trend of the KPIs is almost identical to the baseline one. This shows clearly that we are unable to identify low intensity faults because their effects on the system are negligible.

CONCLUSIONS

We have shown that, considering a hypothetical but realistic case study, complex network analysis is capable of quantifying the decrease of resilience in a system under fault/attack and to tell the difference between the two. As a future work, we plan to consolidate our methodology by further integrating by formalizing the theoretical connections between the observed measures and the dynamics of the underlying system. On the engineering side, we wish to extend our analysis to cover other systems of systems, and further validate our methodology by extending it to evaluate resilience of other critical-infrastructure facilities, with a focus on energy production plants and distribution networks.

Experiment	Simulation	Louvain Energy		Giant Energy	
		FG	*DFG*	*FG*	*DFG*
EXP 1	Standard	0.341	0.701	0.233	0.142
	Attack (H)	$1.302 * 10^{-9}$	0.039	$3.091 * 10^{-8}$	0.716
	Attack (L)	$1.339 * 10^{-8}$	0.001	$1.150 * 10^{-4}$	0.059
	Fault (H)	0.009	$7.774 * 10^{-6}$	0.717	0.411
	Fault (L)	0.437	0.134	0.060	0.124
	Both (H)	$1.383 * 10^{-9}$	$2.789 * 10^{-9}$	$1.585 * 10^{-8}$	0.126
	Both (L)	$1.983 * 10^{-8}$	$6.569 * 10^{-4}$	0.260	0.043
EXP 2	Standard	0.653	0.020	0.919	0.289
	Attack (H)	$1.302 * 10^{-9}$	0.527	$1.417 * 10^{-8}$	0.838
	Attack (L)	$4.790 * 10^{-8}$	0.026	0.012	0.043
	Fault (H)	0.454	$4.458 * 10^{-7}$	0.081	0.694
	Fault (L)	0.143	0.988	0.114	0.452
	Both (H)	$2.661 * 10^{-9}$	$1.309 * 10^{-8}$	$4.790 * 10^{-8}$	0.005
	Both (L)	$2.478 * 10^{-8}$	0.069	0.105	0.924
EXP 3	Standard	0.151	0.489	0.813	0.716
	Attack (H)	$7.159 * 10^{-9}$	0.002	$4.012 * 10^{-9}$	0.389
	Attack (L)	$1.549 * 10^{-7}$	0.005	$1.150 * 10^{-4}$	0.754
	Fault (H)	0.015	$1.544 * 10^{-7}$	0.382	0.988
	Fault (L)	0.881	0.608	0.739	0.650
	Both (H)	$7.556 * 10^{-10}$	$9.773 * 10^{-9}$	$4.067 * 10^{-8}$	0.208
	Both (L)	$1.468 * 10^{-9}$	0.001	0.017	0.020

TABLE I: Table of p-values for our experiments. The **Experiment** column indicates the experiments of interest. The **Simulation** column identifies the specific simulation inside a specific experiment, L indicates that the low importance terminals/routes were chosen for the attack/fault whereas H indicates that the high importance ones were chosen. **Louvain Energy** and **Giant Energy** indicate the measure considered and they represent the Laplacian Energies of the Louvain Communities graph and the Giant Component graph respectively. *FG* and *DFG* represent the Flux Graph and the Difference Flux Graph respectively.

Experiment	Simulation	Louvain Energy		Giant Energy	
		FG	*DFG*	*FG*	*DFG*
EXP1	Standard	0.135	0.003	0.288	0.310
	Attack (H)	2.341	0.290	1.623	0.167
	Attack (L)	1.682	0.576	0.802	0.409
	Fault (H)	0.497	1.005	0.102	0.171
	Fault (L)	0.208	0.191	0.376	0.344
	Both (H)	2.033	1.683	1.567	0.272
	Both (L)	1.460	0.702	0.338	0.261
EXP2	Standard	0.116	0.430	0.004	0.245
	Attack (H)	2.233	0.104	1.749	0.017
	Attack (L)	1.617	0.466	0.548	0.467
	Fault (H)	0.242	1.285	0.321	0.200
	Fault (L)	0.296	0.060	0.369	0.157
	Both (H)	1.827	1.613	1.609	0.485
	Both (L)	1.591	0.416	0.431	0.027
EXP3	Standard	0.260	0.031	0.040	0.017
	Attack (H)	1.698	0.585	1.969	0.200
	Attack (L)	1.295	0.598	0.940	0.118
	Fault (H)	0.499	1.535	0.203	0.044
	Fault (L)	0.052	0.188	0.060	0.194
	Both (H)	2.012	1.681	1.446	0.339
	Both (L)	2.043	0.672	0.481	0.488

TABLE II: Table of the Cohen's d coefficients for our experiments. The **Experiment** column indicates the experiment of interest. The **Simulation** column identifies the specific simulation inside a specific experiment, L indicates that the low importance terminals/routes were chosen for the attack/fault whereas H indicates that the high importance ones were chosen. **Louvain Energy** and **Giant Energy** indicates the measure considered and they represent the Laplacian Energies of the Louvain Communities graph and the Giant Component graph respectively. *FG* and *DFG* represent the Flux Graph and the Difference Flux Graph respectively.

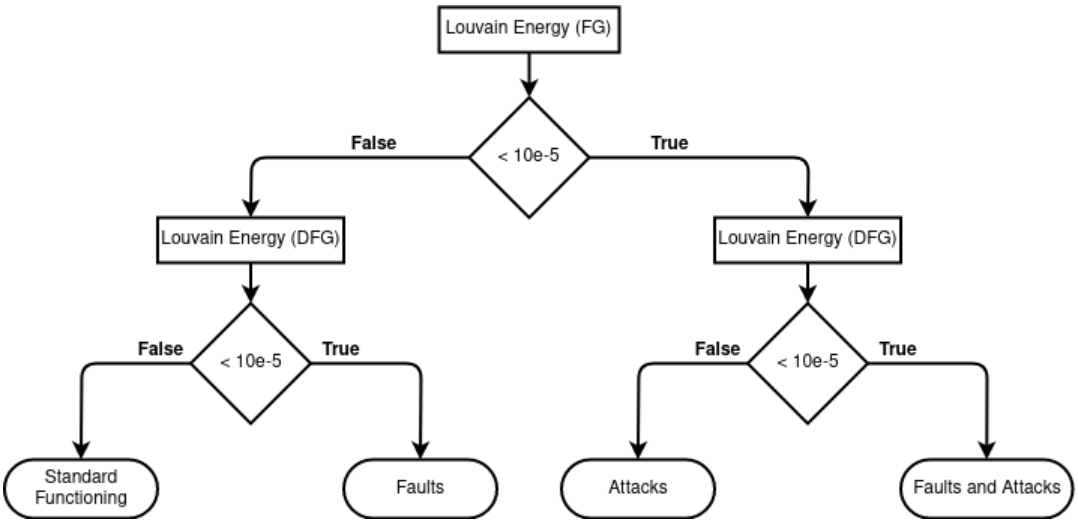

Fig. 4: Graphical representation of the rules extracted by our analysis on the ILS.

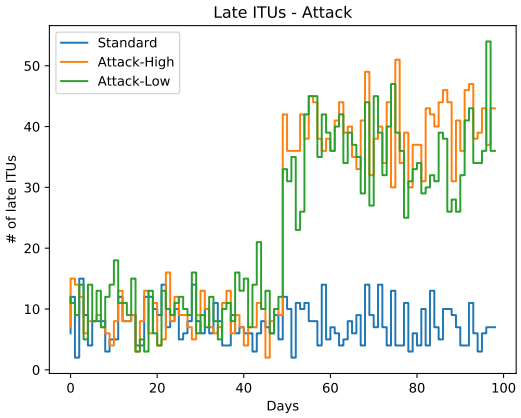

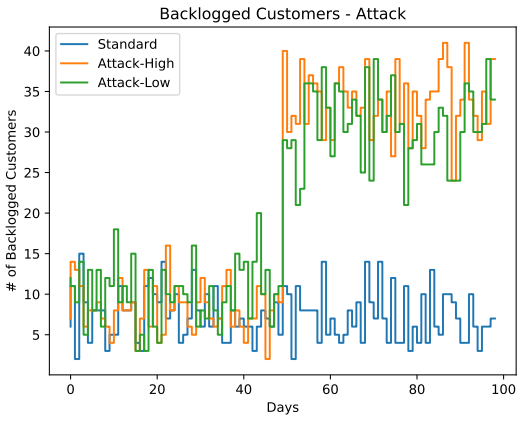

Fig. 5: Trends of the KPIs Late ITUs and Number of Backlogged Customers during a simulation subject to an attack.

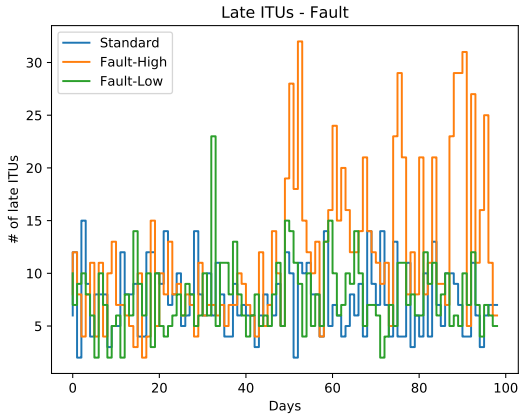

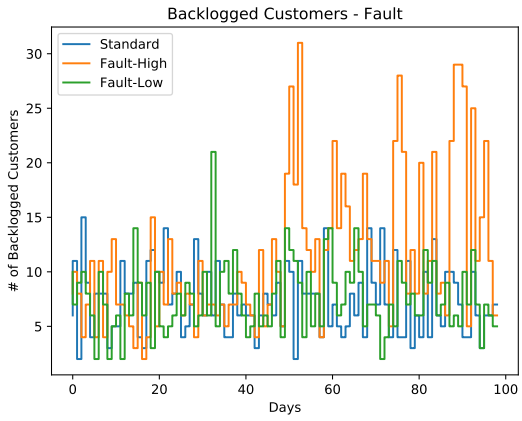

Fig. 6: Trends of the KPIs Late ITUs and Number of Backlogged Customers during a simulation subject to a fault.

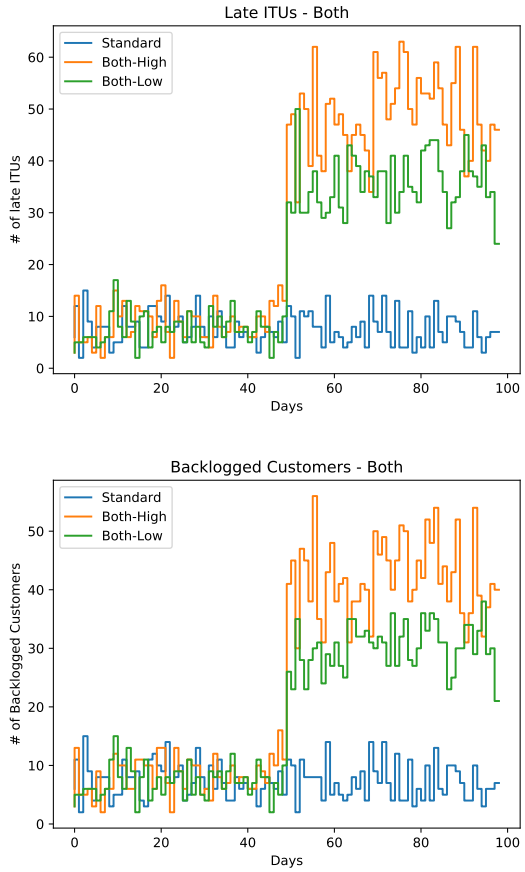

Fig. 7: Trends of the KPIs Late ITUs and Number of Backlogged Customers during a simulation subject both to a fault and an attack.

REFERENCES

[AJB00] Reka Albert, Hawoong Jeong, and Albert-Laszlo Barabasi. Error and attack tolerance of complex networks. *Nature*, 406:378382, 2000.

[Bar] Albert-Laszlo Barabasi. *Network Science*. Cambridge University Press.

[BGLL08] Vincent D Blondel, Jean-Loup Guillaume, Renaud Lambiotte, and Etienne Lefebvre. Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008, 2008.

[BHSZ15] Fredrik Björck, Martin Henkel, Janis Stirna, and Jelena Zdravkovic. Cyber resilience-fundamentals for a definition. In *WorldCIST (1)*, pages 311–316, 2015.

[CCT13] Matteo Casu, Giuseppe Cicala, and Armando Tacchella. Ontology-based data access: An application to intermodal logistics. *Inf. Syst. Frontiers*, 15(5):849–871, 2013.

[DRKS08] Salvatore DAntonio, Luigi Romano, Abdelmajid Khelil, and Neeraj Suri. Increasing security and protection through infrastructure resilience: the inspire project. In *International Workshop on Critical Information Infrastructures Security*, pages 109–118. Springer, 2008.

[FR11] James P Farwell and Rafal Rohozinski. Stuxnet and the future of cyber war. *Survival*, 53(1):23–40, 2011.

[Fre77] Linton C Freeman. A set of measures of centrality based on betweenness. *Sociometry*, pages 35–41, 1977.

[GZ06] Ivan Gutman and Bo Zhou. Laplacian energy of a graph. *Linear Algebra and its applications*, 414(1):29–37, 2006.

[KG14] Marek Korytar and Darja Gabriska. Integrated security levels and analysis of their implications to the maintenance. *Journal of Applied Mathematics, Statistics and Informatics*, 10:3342, 2014.

[LLL+21] Ming Li, Run-Ran Liu, Linyuan Lu, Mao-Bin Hu, Shuqi Xu, and Yi-Cheng Zhang. Percolation on complex networks: Theory and application. *Physics Reports*, in press, 2021.

[NGZL15] Tingyuan Nie, Zheng Guo, Kun Zhao, and Zhe-Minhg Lu. New attack strategies for complex networks. *Physica A: Statistical Mechanics and its Applications*, 424:248–253, 2015.

[QFW+12] Xingqin Qi, Eddie Fuller, Qin Wu, Yezhou Wu, and Cun-Quan Zhang. Laplacian centrality: A new centrality measure for weighted networks. *Information Sciences*, 194:240–253, 2012.

[WFD10] Chunlei Wang, Lan Fang, and Yiqi Dai. A simulation environment for scada security analysis and assessment. In *Measuring Technology and Mechatronics Automation (ICMTMA), 2010 International Conference on*, volume 1, pages 342–347. IEEE, 2010.