An Analysis of Link and Node Level Resilience on Network Resilience

Danique Lummen
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands
d.l.m.lummen@student.utwente.nl

ABSTRACT

Communication networks should be resilient to be able to offer an acceptable level of service even in the face of challenges. However, how to measure the network resilience is not straightforward. Moreover, the resilience of the network depends on the type of risk it is exposed to, e.g., targeted attacks or random failures, and the scale of the risk, e.g., small or large scale failures. Therefore, in this paper, we first overview the literature on the network resilience metrics and the potential risks a network might experience. As the resilience of a communication system depends on the resilience of the levels it relies upon, we focus on the node and link level resilience. Via simulations, we analyse the impact of the resilience of links and nodes on the network resilience. Our analysis reveals that link placement in networks has a large influence on the resilience and should therefore be considered carefully when designing resilient wired networks.

Keywords

Network Resilience, Link Level Resilience, Node Level Resilience, Risk Models, Network Challenges, Metrics

1. INTRODUCTION

The Internet, or more generally speaking: communication networks, have become an essential part of our daily lives. These networks are used for a variety of things, most notably they provide access to information and a means of communication with others. Numerous instances, such as governments, depend on the functioning of the Internet for their daily operation and disaster response. Therefore, the Internet may be classified as a critical infrastructure: an asset that is essential for the functioning of a society and economy. An example on how the Internet is a critical infrastructure is the dependency of the electrical grid on the Internet and vice versa. The Internet relies on the electrical grid for power, whilst the electrical grid depends on the Internet for SCADA (supervisory control and data acquisition) [5].

As the dependence on the Internet increases, there is also an increased vulnerability of communication systems to various problems. If an incident occurs where part of the Internet goes down, it creates large problems for other systems. Therefore, measures should be taken to ensure that even when part of the system fails, this should not have a significant impact on the functioning of the network, which we call *resilience*. A resilient system can be described as "one that continues to offer an acceptable level of service even in the face of challenges, whatever the nature of the challenge that it faces" [26]. As the goal is to minimise the system failures, it is necessary to ensure that communication networks are as resilient as possible when faced with a number of challenges. An example of such a challenge is a natural disaster or human error.

Significant research has already been done on methods and frameworks to ensure resilience in communication networks [19], but significantly more work needs to be done to understand and define resilience metrics [9, 6] as well as on how to quantify the resilience of a network [8]. As a communication network consists of multiple levels, the resilience of each level builds on the resilience of the levels it depends upon [9]. Not much is known yet on the specific impact each level has on the total resilience. Therefore, the purpose of this paper is to get a deeper understanding of the resilience specifically on the link and node level and to quantify the impact of this resilience on the overall resilience. More specifically, we will address the following three questions: (i) What metrics can be used to quantify link and node level resilience in a communication network?, (ii) Which risks and challenges might a network experience that test the resilience of a the network?, and (iii) How does the link and node level resilience compare for a communication netowkr with different links.

The rest of the paper is organized as follows. First, Section 2 provides background on the network resilience, whereas Section 3 elaborates on the metrics currently being used for resilience quantification. Section 4 provides an overview of challenges that networks might experience and in Section 5 the methodology for the simulations is explained. Finally, Section 6 discusses the results and concludes with a list of future research directions.

2. BACKGROUND

As mentioned before, resilience is the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation. In telecommunication, this acceptable level of service is usually defined in a Service Level Agreement (SLA) between the customer and the network service provider [6]. This agreement specifies the service levels that are considered acceptable to the customer, as well as the service levels where the service is impaired or unacceptable. Faults and challenges that the network faces, such as for example a natural disaster, impact the level of operation for the network, which in turn can cause the level of service

33rd Twente Student Conference on IT Jul. 3rd, 2020, Enschede, The Netherlands.

to degrade to an impaired or unacceptable state.

In order to evaluate the resilience of a network a resilience state space model was created by Hutchison et al. [9], which has also been used in a variety of other studies [12, 19]. This state space model is created in a three step process. First, the operational condition of the network is represented using metrics, which are called *operational metrics* as they explain the operational state of the network. Second, the level of service that is being provided by the network is quantified using *service parameters*. As the third and final step these operational metrics and service parameters are aggregated into *network states*, which represent the network. A representation of the state space of resilience in which the difference between a resilient service and a non-resilient service can be seen is Figure 1.

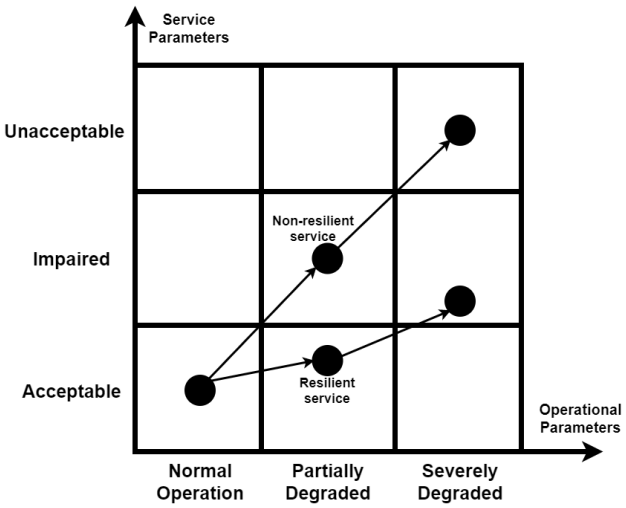

Figure 1: Resilience State Space, adapted from [9, 6].

As the network is exposed to challenges, the network state transforms from one state to another based on how the service parameters and operational metrics are impacted by the challenge. Therefore, we evaluate the resilience of a given network based on their transitions in the network state when exposed to challenges.

As already stated above, operational metrics explain the operational state of a given network. Different properties can be used to derive operational metrics, all dependent on what type of network is being used and at which level of the network the resilience is being evaluated. An example of an operational metric on the physical level of a network is propagation loss.

The operational state of a network can be represented by a set of operational metrics. The operational state space in which this operational state is represented can be divided into three regions: *normal*, *degraded* and *severely degraded*, which specify the level of operation of a network. State boundaries for the operational metrics are defined to determine when a state transition occurs.

When an event happens that degrades the operational state of the network, due to impacting one of the operational metrics, two things can happen: a state transition or a sub-state transition. When the event results in a state transition, the level of operation of the system transitions from one state to another. An example of this is a state transition from *normal* operation to *partially degraded* operation. When a sub-state transition occurs, one or more of the operational metrics are impacted, but not heavily enough to cause transition to another state. For example, the state boundary between *normal* and *partially degraded* service for the operational metric 'delay' is 200ms. The value of this state boundary depends on the application, as 200ms is not long for a data download, but is too long for mission-critical applications. A network with a delay of 150ms is challenged and the delay is increased. If the delay stays below 200ms, only a sub-state transition will occur, but if the delay exceeds 200ms a state transition will occur and the operational state of the network will shift from *normal* to *partially degraded*.

Service parameters specify the level of service that is being provided by the network, an example of a service parameter is latency. When the latency of a VoIP network increases due to challenges, the service level might become unacceptable.

When a resilient service and a non-resilient service face the same level of degradation of the operational parameters due to a fault or disturbance of the normal operation, the resilient service will have less degradation of the service parameters than the non-resilient service [19]. In other words, we can define the resilience of a network as the slope between two states of a network: the initial state and the state when a challenge occurs. The lower the slope, the higher the resilience.

In resilience research, networks are abstracted as graphs. We assume a network of N nodes and E edges is defined as graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$. The set of nodes in the network is denoted by $\mathcal{N} = \{n_1, ..., N\}$. Each node is defined by a set of properties, e.g., number of outgoing links, number of neighbours, CPU, storage capacity and failure probability. We denote an edge between node i and node j by e_{ij}. The set of edges in the network is denoted by $\mathcal{E} \subseteq \{e_{ij} | (n_i, n_j) \in \mathcal{N}^2 \wedge n_i \neq n_j\}$. Each edge is also defined by a set of properties, e.g. bandwidth, capacity, centrality and failure probability.

3. RESILIENCE METRICS

As stated in the previous section, in order to evaluate the resilience of a network, operational metrics and service parameters are needed. Resilience can be evaluated at multiple levels of the system and the resilience of the higher levels of the system depends on the resilience of the lower levels [9]. Also, in multilevel resilience analysis, the service parameters of one level become the operational metrics of the level above. In other words, the service provided by a given level becomes the operational state of the level above [12]. For example, when looking at the 7 layers of the Open System Interconnection (OSI) model, the service parameters of the first level, which is the physical level, in turn are the operational metrics for resilience on the second data link level. For our research on node and link level resilience we will be focusing on the third layer in the OSI model, the network layer, having to do with network topology, routing and policy [13].

The European Network and Information Security Agency (ENISA) has done a study with a group of stakeholders on resilience measurements, and on what metrics can be used to quantify resilience [6, 7]. A number of these metrics have a specific focus on the links and nodes of a network, such as operational availability, operational reliability, delay variation (jitter), packet loss and link/node failure. Link/node failure is an indicator for the robustness of a network to link and or network nodes failures and is expressed as a network performance parameter in function of the number of links, network nodes or components of the network nodes that are removed [7].

Jabbar et al. [12] also used some metrics specifically aimed at link and node level resilience in their research, e.g. the relative number of connected components in a network as well as the clustering coefficient. In a study done by Rosenkrantz et al. [25] several other node and edge connectivity metrics are mentioned, such as the average node degree, the number of components and the largest component size. Simulation analysis done by Çetinkaya et al. [3] evaluates the network performance by using the aggregate packet delivery ratio metric. Hop count and system stability are metrics used in research by Ibrahim et. al [10].

From all the previously mentioned metrics, metrics such as the clustering coefficient, average node degree, relative largest component size and number of connected components seem to be the ones most commonly used.

The average node degree is the number of edges per node in a graph. Assume graph $\mathcal{G}$ has N nodes and E edges, the average node degree is equal to $deg(G) = \frac{E}{N}$. The clustering coefficient measures the degree to which nodes in a graph tend to cluster together. Therefore, it measures how connected a node's neighbours are to one another. This clustering coefficient for node i: (C_i), can be calculated by dividing the number of edges connecting i's neighbours by the total number of possible edges between i's neighbours. The network clustering coefficient C is the average of all the local clustering coefficients: $C = \frac{1}{n}\sum_{i=1}^{n} C_i$. The relative largest component size is the size of the largest connected component compared to the total number of nodes in the network. This is calculated by dividing the number of nodes in the largest connected component by the total number of nodes in the network.

4. NETWORK CHALLENGES

The need for resilience in a communication system can be derived from the catastrophic damages resulting from a non-resilient system being faced with challenges. A *challenge* is any characteristic or condition that impacts the normal operation of a network [26]. A challenge can trigger the fault $\rightarrow$ error $\rightarrow$ failure chain, ultimately resulting in failure of the system. The challenge triggers a fault, which in turn could cause an error. If this error propagates it may lead to the failure of the network service. Therefore, in order to design a resilient network, it is important to understand how the network behaves under these challenges.

As Figure 2 shows, challenges to a communication network can be grouped into seven categories [4]: large-scale disasters, socio-political and economic challenges, dependent failures, human errors, malicious attacks, unusual but legitimate traffic and environmental challenges.

Large-scale Disasters
This challenge category includes a large number of challenges to communication systems, which can be split into two groups: disasters with natural causes and disasters with human-made causes. Large-scale natural disasters can be caused by terrestrial events such as an earthquake or fires, meteorological events such as hurricanes and by cosmological events such as solar storms. Human-made disasters are the challenges with big impact that are caused by human action, either by accident or by deliberate malicious intent, for example when early warnings in the operation of a system are ignored. The impact of large-scale disasters is often enormous; regions impacted are often big and the time needed to undo the damage done is usually long.

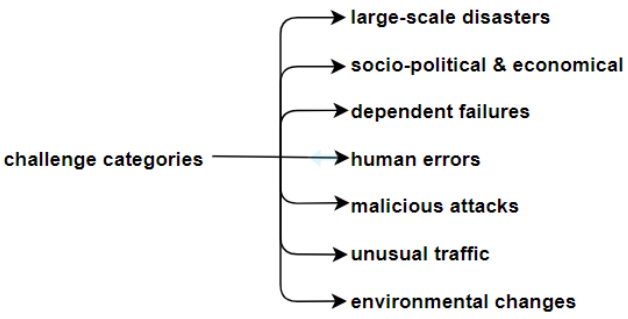

Figure 2: Taxonomy of major challenges [4, 16].

Socio-political and Economic Challenges
These challenges are specifically caused by human actions, with the intent of social, political or economic gain, such as gaining an advantage on the economical markets [2]. An example of a political challenge to a communication system is the DDoS attack against the country of Estonia by Russia. The Estonian government decided to move a statue honouring fallen WWII soldiers, angering the Russians and causing them to start DDoS attacks on all major networks of Estonia. In the end Estonia was unable to stop the attack and ultimately decided to cut the Internet connection with the outside world so that Estonian residents could continue to use the national services [14].

Dependent Failures
These challenges occur when a system on which a network is dependent fails, causing a failure in the system itself. The failure of the supportive system causes a failure in the dependent system, therefore causing a disruption. These kinds of failures have the possibility to have cascading effects, resulting in large scale damage and therefore changing to a Large-scale disaster. An example of a dependent failure is the failure of the electrical grid when there is a failure in the Internet, as the electrical grid relies no the Internet for SCADA (supervisory control and data acquisition) [5].

Human Errors
Human actions can also lead to failures of a system. These actions are usually performed without malicious intent and can either be accidental (non-deliberate) or due to incompetence (deliberate). An example of such a challenge is the "This site may harm your computer" Google accident. On January 31^{st} 2009, almost every search result from Google led to a "This site may harm your computer"-message. Normally these messages are used to signal the user that they are about to enter a website that may possibly harm their computer. This accident occurred due to a simple human error, as the list with these harmful web addresses was edited and a single '/' was mistakenly added, causing all websites to be marked as possibly malicious [17].

Malicious Attacks
These are challenges that are deliberately targeted to cause disruption to a system, with malicious intent. An example of such a challenge is the use of the Stuxnet worm as an attack on the Iranian nuclear facilities. The worm infected the programmable logic controllers used to control the centrifuges that enrich uranium, causing these centrifuges to self-destruct [20].

Unusual but Legitimate Traffic
These challenges are called flash crowds, in which a large number of users makes a request to access a service at the same time. The effects of flash crowds look similar to that

of a DDoS-attack, however unlike DDoS-attacks they do not have a malicious intent. An example of a flash crowd is the unavailability of a large number of news websites after the 9/11 terror attack on the World Trade Center. Due to the large amount of people wanting to know what happened, a flash crowd occurred, causing the website of a number of news stations to be unresponsive [23].

Environmental Challenges
The final category of challenges has to do with the network environment itself. These challenges include unpredictably long delay, weak connectivity of wireless channels and mobility of nodes.

All these challenges can also be characterised based on their time duration and the spatial region, both regarding the challenge itself and the impact afterwards. For example, the time duration of an earthquake, which only takes a few seconds, differs largely from the time duration of a hurricane, which can take hours. However, both have an impact on a large spatial region, and the time duration of this impact can take days.

Some challenges could fall into multiple categories of challenges, depending on their scale, goal and target. For example, the DDoS-attack on the Estonian government falls under both the socio-political and economic challenges category as well as the malicious attack category.

5. METHODOLOGY

In this section we will be describing the method used to address the third research question. To compare the link and node level resilience for networks with different types of links, as stated in **RQ3**, we will be creating and using a simulator. This simulator will be used to simulate a number of challenges on different networks to determine the impact on the network resilience of these challenges.

5.1 Network Topologies

We will be evaluating the performance of three separate topologies under different challenges. The first topology is the Surfnet inferred topology (shown in Figure 3a), which dataset we got from The Internet Topology Zoo from the University of Adelaide [28]. The Surfnet network is the backbone network of all institutions for higher education in the Netherlands, which is used for communication between the different institutions [27].

The second and third topology are synthetic topologies generated using a topology generation tool called KU-LocGen [11]. The topologies are generated with the same number of nodes, all at the same geographic location as the nodes in the Surfnet topology (shown in Figure 3b and 3c.) Using the KU-LocGen generator these two topologies are generated using the Waxman topology model [21], which takes into account the geographic location of the nodes when placing the links. Therefore, all three topologies have the same nodes in the same locations, but they all differ in link placement. The first synthetic topology, generated with the Waxman model with $\alpha = 0.4$ and $\beta = 0, 2$, has a lot more redundant links than the original Surfnet topology. On the other hand, the second synthetic topology, generated with the Waxman model with $\alpha = 0, 19$ and $\beta = 0, 21$, has around the same number of links as the Surfnet topology, but the distance between the connected nodes is much larger. Graph characteristics of all three topologies can be found in Table 1.

Table 1: Characteristics of Network Topologies.

	Surfnet	**Synthetic1**	**Synthetic2**
Number of Nodes	50	50	50
Number of Edges	73	118	65
Clustering Coefficient	0.0958	0.1099	0.0579
Average Node Degree	2.92	4.72	2.6
Average Hopcount	4.364	2.784	4.219
Network Diameter	11	6	10

5.2 Challenge Scenarios

As this research is specifically focusing on the effects of resilience at the link and node level of a communication network, we will be looking at challenges that may do direct damage to these levels. We will be looking at three different categories of challenges; malicious attacks, random node/link failures and large-scale disasters, based on the challenge categories from Section 4. In order to simulate the effects of these challenges on a network, we create a number of challenge scenarios. Each of these scenarios explains a challenge which will be simulated to occur to the networks in the simulator.

The scenarios will be defined based on a selection of classes from the challenge taxonomy defined by Çetinkaya et al. [4] and the fault taxonomy developed by the IFIP 10.4 working group [1]. The following template has been created based on these two taxonomies:

Challenge - Name of the challenge
Cause - The phenomenological cause of the challenge, can either be *human-made* or *natural*.
Intent - The intent of the act challenge can either be *deliberate*, or *non-deliberate*.
Scope - The challenge can impact the *nodes* or *links* within a network, or the entire *geographic area* of the network.
Simulations - Explanation of the simulation with regards to where the node and/or link failures will occur.

We will consider three categories of challenges: malicious attacks, random failures and large scale disasters. For malicious attacks we simulate two attacks on critical nodes, one targeting critical nodes based on node betweenness, the other targeting on node degree, as well as one attack on critical links. Which link or node fails is determined by the criticality of the links/nodes, where the most critical ones fail first. For random failures, we simulate both random node and random link failure within the system. Which nodes and link fail in these simulations is determined randomly. These simulations are executed 50 times for accuracy, and the 95% confidence intervals are shown in the results. Finally, for large-scale disasters we simulate large-scale failure in three separate areas of the network, two of which are on the edges of the network, whilst the third is in the critical centre of the network.

We consider the following six scenarios:

S1 Critical Node Attack using Node Betweenness: This scenario simulates a human-made, deliberate attack on the network nodes. In the simulation, 1%-50% of the nodes will fail. The node criticality is determined by the node betweenness, a higher betweenness centrality indicates a higher node criticality.

S2 Critical Node Attack using Node Degree: This

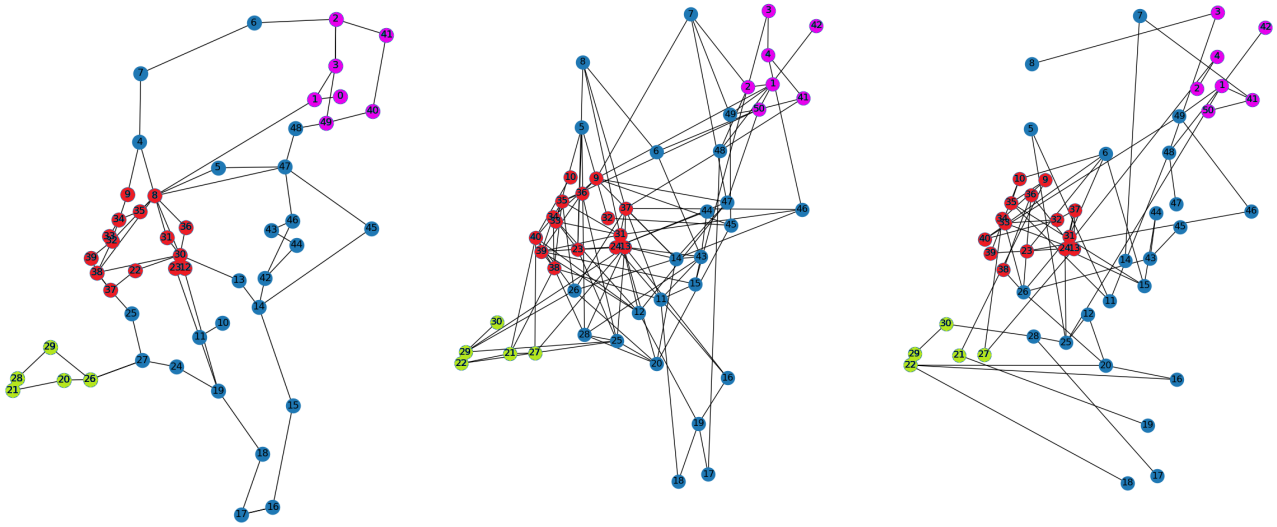

(a) Surfnet Topology (b) Synthetic Topology 1 (c) Synthetic Topology 2

Figure 3: Topologies used in simulation

scenario simulates a human-made, deliberate attack on the network nodes. In the simulation, 1%-50% of the nodes will fail. The node criticality is determined by the node degree, a higher node degree indicates higher node criticality.

S3 Link Attack using Link Betweenness: This scenario simulates a human-made, deliberate attack on the links in a network. In the simulation, 1%-50% of the links will fail. The link criticality is determined by the link betweenness, a higher betweenness centrality indicates a higher link criticality.

S4 Random Link Failure: This scenario simulates a human-made, non-deliberate failure of the links in a network. In this simulation, 1%-50% of the links in the network will fail.

S5 Random Node Failure: This scenario simulates a human-made, non-deliberate failure of the nodes in a network. In this simulation, 1%-50% of the nodes in the network will fail.

S6 Large-scale disaster This scenario simulates a natural, non-deliberate failure of nodes and links in a specific fixed geographic area. For each of the geographic areas coloured in Figure 3 all nodes and links will fail.

5.3 Operational Metrics

As the goal of this research is to determine the effect of the link and node level resilience, the operational metrics that are used for the simulations are *percentage of link failures* and *percentage of node failures*. Which one of these metrics is used depends on the failures being simulated; the percentage of link failures is used for all link failure challenges and vice versa for the percentage of node failures. The regions for the *normal* operation, *partially degraded* operation and *severely degraded* operation are defined in Table 2 and can be tuned depending on the service of interest.

5.4 Service Parameters

For our simulations, we define the service provided by the network with four parameters: the relative largest component size, clustering coefficient, average node degree and the number of connected components. The service regions for *acceptable* service, *impaired* service and *unacceptable* service are defined in Table 3 and can be changed depending on the situation and service being evaluated.

Table 2: Operational Regions

Region	Percentage of Link Failures x	Percentage of Node Failures y
Normal	$0 < x < 10$	$0 < y < 10$
Partially Degraded	$10 \leq x < 25$	$10 \leq y < 25$
Severely Degraded	$x \geq 25$	$y \geq 25$

In order to determine the final value for the service parameter which we can use to determine the resilience we aggregate and normalise all service parameter values. We define the service parameter considering these four metrics as follows:

$$SP = \frac{1}{(p_1/1) + (p_2/0.08) + (p_3/2.7) + (1/p_4)}.$$

Using the boundaries for all service parameters, the derived boundaries for the final service parameter can be found in Table 4.

Table 3: Service Parameters

	Acceptable	Impaired	Unacceptable
Rel. LC Size p_1	$p_1 = 1$	$0.90 \leq p_1 < 1$	$p_1 < 0.90$
Clustering Coefficient p_2	$p_2 \geq 0.08$	$0.04 \leq p_2 < 0.08$	$p_2 < 0.04$
Avg. Node Degree p_3	$p_3 \geq 2.7$	$2.3 \leq p_3 < 2.7$	$p_3 < 2.3$
Nr. of Conn. Components p_4	$p_4 = 1$	$1 < p_4 \leq 4$	$p_4 > 4$

5.5 Simulation Environment

The simulation environment is created in Python and uses the NetworkX [22] and Plotly [24] libraries to visualise the abstracted graphs and results from the simulations. For

Table 4: Final Service Parameter Boundaries

	Final Service Parameter SP
Acceptable	$SP < 0.25$
Impaired	$0.25 \leq SP < 0.4$
Unacceptable	$SP > 0.4$

the sake of simplicity we assume that all routing is done through Dijkstra's shortest path algorithm and that all links have the same bandwidth and transmission delay. We measure the traffic over the network by the means of stress centrality, which measures the amount of communication that passes through a link based on the number of shortest paths passing through that link [18]. This stress centrality of edge e is calculated by:

$$c_s(e) = \sum_{s \in N} \sum_{t \in N} \sigma_{st}(e).$$

Where $\sigma_{st}(e)$ denotes the number of shortest paths containing edge e. In our simulations we use this stress centrality as a failure model for the links due to stress. Should the stress centrality of a certain edge surpass a specific boundary, the link will fail due to stress.

The value of the stress centrality of a link e can lay between 0 and 1. If $c_s(e) = 0$ then out of all shortest paths, no paths use link e. If the value of $c_s(e) = 1$ then all shortest paths use link e. After evaluation of the stress centrality values of the networks under normal conditions, as well as under partial link/node failure, the value for the stress centrality boundary has been set at 0.3. The value is chosen in such a way that a few failures will not cause the boundary to be exceeded, however in case of multiple failures the boundary will be exceeded by some critical links.

After every failure of a link or node, the traffic balance over the network is recalculated and if there are any links that surpass the stress centrality boundary, these links will fail due to stress.

6. PERFORMANCE ANALYSIS

In this section, we apply our challenge scenarios to the network topologies as specified in Section 5 and analyse and discuss the findings. Next to all the results in this paper, separate figures for each of the service parameters, including the 95% confidence intervals for all the random link and node failure simulations, as well as all the other results can be found on the github repository [15].

Malicious Attacks

Figure 4c shows the results of the simulation on all three network topologies for the targeted link failure. We observe that both the Synthetic2 and Surfnet topologies transition to an *impaired* service level after only a few critical link failures. This is to be expected, as the most critical links are targeted and therefore the system is being attacked at place with the most impact. The results also show that, compared to the other two topologies, the Synthetic1 topology is much more resilient with regards to targeted link failure. The level of service for this topology only transitions to an *impaired* state when the operational level of the network is *severely degraded*.

Figures 4a and4b show the results of the simulations of the targeted node failure scenarios. We observe that the Surfnet and Synthetic2 topology rapidly transition to an *impaired* and later *unacceptable* state after a relatively low percentage of critical node failures. Interesting to see is that the Synthetic2 topology seems to be the least resilient for targeted node failures on node betweenness, whilst the Surfnet topology is the least resilient for targeted node failures on node degree. As the length of the links is the main difference between the two topologies, it can be argued that topologies with longer links are less resilient to targeted attacks on node betweenness than topologies with shorter links. The same can be said the other way around for targeted attacks on node degree.

Compared to the other two topologies, the Synthetic1 topology is again more resilient. Though it can be seen that the topology is more resilient against targeted attacks on node degree than on node betweenness. In both the targeted node failure simulations we see a big jump in the service level of Surfnet around 7-8% of node failures, we believe this is due to additional link failures because the stress centrality traffic boundary is exceeded.

When comparing the results from all three malicious attack simulations we observe that all three topologies are less resilient to targeted node attacks than to targeted link attacks. Therefore it might be more valuable for a network operator to use resources to improve the resilience against targeted node attacks than to improve resilience against targeted link attacks.

Random Failures

As Figure 5a depicts the Synthetic1 topology is more resilient than the two other topologies, as the slope of the graph for this topology is less steep. The service level of the Synthetic1 topology also never reaches the unacceptable state, whilst the service level of the other two topologies does.

Figure 5b shows the results of the simulation on random node failure. The figure shows a faster decline of the service parameters for the Surfnet and Synthetic2 topologies compared to the Synthetic1 topology. Both the random link failure and the random node failure simulation results show that the Synthetic2 topology is the least resilient towards these kinds of failures.

Table 5 shows the slope of the values of all topologies from Figures 4c to 5b. Recall that the slope corresponds to the resilience of the network, a lower slope represents higher resilience. From this data we can conclude that the Synthetic1 topology is the most resilient of all topologies, as for all simulations, its slope is the lowest out of the three topologies. We believe this to be the case due to the number of redundant links present in this topology, which cause the system to be better equipped to handle node and link failures. Even if some links fail, traccif can be routed through alternative paths helping the network maintaing its service.

Both the Surfnet and Synthetic2 topologies are roughly equally resilient to both random and targeted link failures, with Surfnet being slightly more resilient to random link failures and Synthetic2 more resilient to targeted link failures. Overall, the resilience of the Synthetic2 topology seems to be slightly worse than that of the Surfnet topology. We believe that this is mostly due to the fact that the distance between the nodes connected by a link is larger in the Synthetic2 topology than it is in the Surfnet topology.

When comparing the slopes for the targeted node failure simulations it is shown that all topologies are less resilient to attacks that target nodes based on node degree compared to node betweenness. For network operators this means that it is most important to prevent these targeted node attacks, as the systems are the least resilient to it.

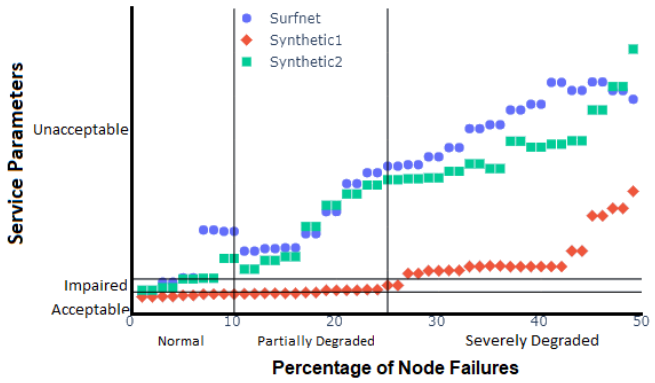

(a) Resilience after Targeted Node Failure. Criticality determined by Node Degree.

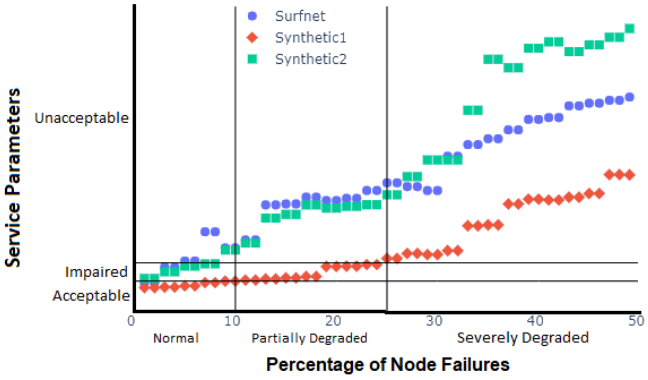

(b) Resilience after Targeted Node Failure. Criticality determined by Node Betweenness.

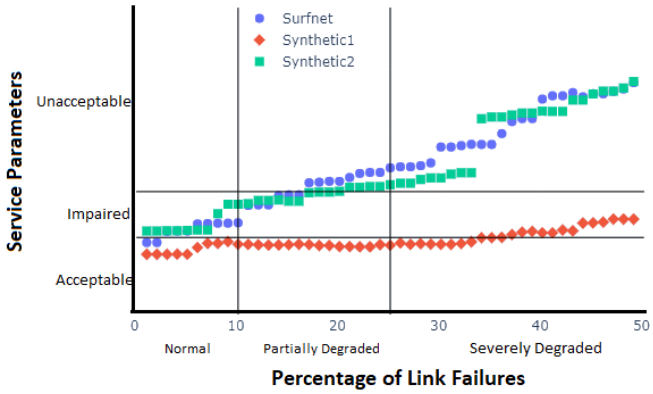

(c) Resilience after Targeted Link Failure.

Figure 4: Resilience after Malicious Attacks.

Table 5: Resilience Slope of all Simulations

	Surfnet	**Synthetic1**	**Synthetic2**
Random Link Failure	0.617	0.232	0.616
Random Node Failure	0.448	0.142	0.560
Targeted Link Failure	0.611	0.131	0.572
Targeted Node Failure - Betweenness	1.784	1.082	2.397
Targeted Node Failure - Degree	2.607	1.417	3.234

Large-scale disasters

From Table 6 we observe that for the large-scale disaster in which the green geographic area is hit (as shown in Figure 3) the service level of only the Synthetic2 topology transitions to an impaired state, whilst both the Surfnet and Synthetic2 topology remain in the acceptable service level state.

For the large-scale disaster in the red geographic area we can see that the service levels of the Surfnet and Synthetic2 topologies transition to an unacceptable state, whilst the service level of the Synthetic1 topology only barely transitions to the impaired level. The impact of this large-scale disaster specifically has a very high impact on the Surfnet topology, we believe this to be due to the fact that this attack targets the entire critical portion of the Surfnet topology, causing the remaining network to completely fall apart. The final large-scale disaster targeting the purple geographic area only causes a state transition to an impaired state in the Synthetic2 topology, both the Surfnet and Synthetic1 topology keep an acceptable level of service.

Table 6: Large-scale Disaster Final Service Parameters

	Surfnet	**Synthetic1**	**Synthetic2**
Normal	0.220	0.195	0.271
Area 1 (Green)	0.226	0.195	0.385
Area 2 (Red)	1.001	0.279	0.515
Area 3 (Purple)	0.224	0.183	0.329

In this research it was assumed that all links in the network have the same bandwidth and that all traffic goes through the shortest paths, taking an arbitrary boundary for link failure due to an overloaded link. This does not accurately represent a real network, and therefore in future research these simulations might be modelled in such a way that bandwidth and routing are taken into account. Another shortcoming of our analysis is that the topologies used in this research all have wired links and therefore no information is known about the resilience of networks that are wireless or partially wireless. There also exist some metrics that might be of interest to resilience research, but have not yet been used widely. An example of such a metric is the number of people connected to a specific network node, which can be used to determine the number of people affected should certain nodes fail. This is also another

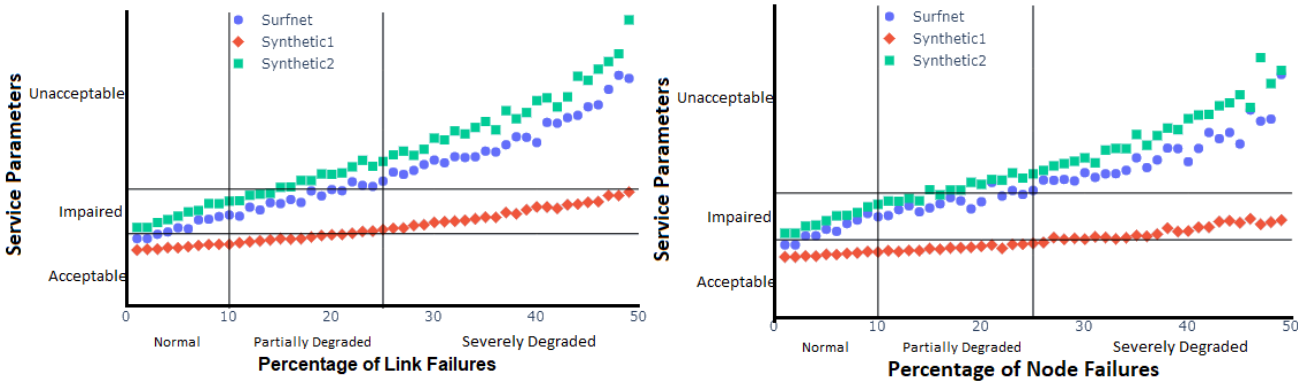

(a) Resilience after Random Link Failure.

(b) Resilience after Random Node Failure.

Figure 5: Resilience after Random Failures.

point for future research.

7. CONCLUSION

Given the increasing importance of communication networks, it is important to understand the resilience of a network and possible risks to a network. Hence, this paper provides a literature survey on resilience metrics and potential risks that might lead to degradation on a network's service. After overviewing the literature and identifying the resilience metrics, we have evaluated the resilience of three selected network topologies to develop insights on how link and node failures affect the resilience of these networks. Our analysis via simulations shows that specifically the link placement in networks has a large influence on the resilience and therefore should be considered carefully when designing new resilient wired networks. Further research still needs to be done with regards to wireless networks to determine the influence of both link and node level resilience for these types of networks.

8. ACKNOWLEDGEMENTS

I would like to express my sincere gratitude to Suzan Bayhan, my supervisor, for her valuable advice during the course of this research.

9. REFERENCES

[1] Avižienis, A., Laprie, J. C., and Randell, B. Dependability and its threats: A taxonomy. In *IFIP Advances in Information and Communication Technology* (2004), vol. 156, Springer New York LLC, pp. 91–120.

[2] Bafna, S., Pandey, A., and Verma, K. Anatomy of the Internet Peering Disputes. *Networking and Internet Architecture* (sep 2014).

[3] Çetinkaya, E. K., Broyles, D., Dandekar, A., Srinivasan, S., and Sterbenz, J. P. Modelling communication network challenges for Future Internet resilience, survivability, and disruption tolerance: A simulation-based approach. *Telecommunication Systems 52*, 2 (feb 2013), 751–766.

[4] Çetinkaya, E. K., and Sterbenz, J. P. A Taxonomy of Network Challenges. In *9th International Conference on the Design of Reliable Communication Networks (DRCN)* (2013).

[5] Department of Homeland Security. A Roadmap for Cybersecurity Research. Tech. rep., Department of Homeland Security, 2009.

[6] ENISA Europe. Measurement Frameworks and Metrics for Resilient Networks and Services-Technical report. Tech. rep., ENISA, 2011.

[7] ENISA Europe. Measurement Frameworks and Metrics for Resilient Networks and Services: Challenges and Recommendations. Tech. rep., ENISA, 2011.

[8] Hosseini, S., Barker, K., and Ramirez-Marquez, J. E. A review of definitions and measures of system resilience. *Reliability Engineering and System Safety 145* (jan 2016), 47–61.

[9] Hutchison, D., and Sterbenz, J. P. Architecture and Design for Resilient Networked Systems. *Computer Communications 131* (oct 2018), 13–21.

[10] Ibrahim, M., and Alsheikh, A. Assessing Level of Resilience Using Attack Graphs. In *Proceedings of the 10th International Conference on Electronics, Computers and Artificial Intelligence, ECAI 2018* (apr 2019), Institute of Electrical and Electronics Engineers Inc.

[11] Jabbar, A., Shi, Q., Çetinkaya, E., and Sterbenz, J. P. G. KU-LocGen: Location and Cost-Constrained Network Topology Generator The University of Kansas. Tech. rep., Univeristy of Kansas, 2008.

[12] Jabbar, A., Sterbenz, J. P., Duncan, T., Frost, V. S., Hutchison, D., Minden, G., Scoglio, C., and Wyglinksi, A. M. A Framework to Quantify Network Resilience and Survivability. Tech. rep., University of Kansas, 2010.

[13] Kumar, S., Dalal, S., and Dixit, V. The OSI Model: Overview on The Seven Layers of Computer Networks. *International Journal of Computer Science and Information Technology Research 2*, 3 (sep 2014), 461–466.

[14] Les, M. The New Front Line: Estonia under Cyberassault. *IEEE Security and Privacy 5*, 4 (jul 2007), 76–79.

[15] Lummen, D. ResiliSim Github Repository, 2020.

[16] Mauthe, A., Hutchison, D., Cetinkaya, E. K., Ganchev, I., Rak, J., Sterbenz, J. P., Gunkelk, M., Smith, P., and Gomes, T. Disaster-resilient communication networks: Principles and best practices. In *Proceedings of 2016 8th International Workshop on Resilient Networks Design and Modeling, RNDM 2016* (oct 2016), Institute of Electrical and Electronics Engineers

Inc., pp. 1–10.

[17] Mayer, M. Official google blog: "this site may harm your computer" on every search result?!?!, jan 2009.

[18] Modarresi, A., and Ymons, J. Modeling and graph analysis for enhancing resilience in smart homes. In *Procedia Computer Science* (2019), vol. 160, Elsevier B.V., pp. 197–205.

[19] Mohammad, A. J., Hutchison, D., and Sterbenz, J. P. G. Poster: Towards Quantifying Metrics for Resilient and Survivable Networks. Tech. rep., University of Kansas, 2006.

[20] Mueller, P., and Yadegari, B. The Stuxnet Worm. Tech. rep., University of Arizona, 2012.

[21] Naldi, M. Connectivity of Waxman topology models. *Computer Communications 29*, 1 (dec 2005), 24–31.

[22] NetworkX. NetworkX — NetworkX documentation.

[23] Partridge, C., Barford, P., Clark, D. D., Donelan, S., Paxson, V., Rexford, J., Vernon, M. K., and Eisenberg, J. The Internet under crisis conditions: Learning from september 11, aug 2003.

[24] Plotly. Plotly Python Graphing Library | Python | Plotly.

[25] Rosenkrantz, D. J., Goel, S., Ravi, S. S., and Gangolly, J. Structure-based resilience metrics for service-oriented networks. In *Lecture Notes in Computer Science* (2005), vol. 3463, Springer, Berlin, Heidelberg, pp. 345–362.

[26] Sterbenz, J. P., Hutchison, D., Çetinkaya, E. K., Jabbar, A., Rohrer, J. P., Schöller, M., and Smith, P. Resilience and Survivability in Communication Networks: Strategies, Principles, and Survey of Disciplines. *Computer Networks 54*, 8 (jun 2010), 1245–1265.

[27] SURFNet. SURF.nl | Samen aanjagen van vernieuwing.

[28] University of Adelaide. The Internet Topology Zoo.