

Identifying Vulnerabilities in the Dutch Electricity Grid

A modelling approach to identify asset vulnerabilities and structural network criticalities under physical malicious threats

SEN2331: CoSEM Master Thesis

RAE Ploemen

Delft University of Technology

Identifying Vulnerabilities in the Dutch Electricity Grid

A modelling approach to identify asset
vulnerabilities and structural network
criticalities under physical malicious threats

by

RAE Ploemen

Robin Anne Elisabeth

Ploemen

Supervisor:	Dr. ir. P. Heijnen
Second supervisor:	Prof. dr. M. Warnier
Mentors:	H. Vergnes & W. Broos
Project Duration:	April, 2026 - September, 2026
Faculty:	Faculty of Technology, Policy and Management, Delft

Cover:	Iwan Baan, 2012
Style:	TU Delft Report Style, with modifications by R. Ploemen

This thesis is confidential and cannot be made public until September 24, 2026.

Preface

This thesis marks the completion of my Master's programme in Complex Systems Engineering and Management at Delft University of Technology. The research was conducted at Stedin and focuses on identifying vulnerable and critical assets in electricity distribution networks under physical malicious threats.

From the moment I started this project, I knew it was exactly the type of research I had hoped to work on. Throughout my studies I have always been fascinated by complex networks, the increasing geopolitical tensions surrounding critical infrastructure, and the societal challenge of electricity grid congestion. Having the opportunity to bring these interests together in a single research project made this thesis especially rewarding to work on.

I would first like to express my sincere gratitude to my supervisor Petra, who has been an exceptional supervisor throughout this thesis. Our weekly meetings were not only highly informative, but also genuinely enjoyable. I always left our discussions feeling motivated, enthusiastic and confident about the next steps in my research. Her guidance, constructive feedback and small talks from time to time made this project a pleasure to work on.

A special thank you goes to my supervisors at Stedin as well. Hugo has taught me an incredible amount about electricity networks. His door was always open, and whenever I got stuck, he would enthusiastically explain concepts by sketching the network on a whiteboard until everything clicked. Walter has been equally valuable throughout this project. His expertise in resilience continuously challenged me to think beyond the technical aspects of the network and helped me to place my research in a broader perspective.

I would also like to thank Martijn for his involvement and willingness to share his knowledge throughout the project. One moment that particularly stayed with me was when he immediately noticed that I had cited my cover page incorrectly. At that point I realized I was receiving feedback from someone who truly paid attention to every detail. His critical perspective has undoubtedly improved the quality of this thesis.

Finally, I would like to thank my family and friends, especially my room-mates, for their continuous encouragement and support throughout my studies. Their confidence in me has meant more than they probably realize. I couldn't have done this without them!

I hope this research contributes, even in a small way, to map vulnerability in electricity networks and provides useful insights for both academia and practice.

*RAE Ploemen
Delft, September 2026*

Executive Summary

Electricity networks are becoming increasingly critical to society, while recent sabotage incidents demonstrate their exposure to deliberate physical attacks. Distribution System Operators (DSOs) therefore need to understand both which assets are most vulnerable to malicious threats and which are structurally most critical to network functioning. Existing research predominantly approaches this problem from a structural perspective using graph-theoretic methods, while asset-specific vulnerability to malicious physical threats has received considerably less attention. Consequently, it remains unclear whether structurally critical assets are also the assets most vulnerable to deliberate physical attacks.

This research addresses this gap through a case study of Stedin's subtransmission network. Two complementary assessment frameworks were developed. First, an asset-specific Vulnerability Index (VI) was constructed based on characteristics relevant to deliberate physical threats, identified through literature and validated with Stedin experts. The VI combines asset exposure with the potential impact of asset failure. Second, a structural criticality framework was developed using graph-theoretic centrality measures, $N - 1$ contingency analysis and sequential targeted attack simulations. Applying both frameworks to the same network enabled direct comparison between asset-specific vulnerability and structural criticality.

The results demonstrate that vulnerability to malicious physical threats cannot be explained by network topology alone. Relevant asset characteristics were identified across four dimensions: Topological, Spatial/Physical, Operational and Security. These characteristics do not contribute equally to vulnerability. Physical redundancy and $N - 1$ capacity redundancy emerged as the strongest and most consistent determinants, demonstrating that the quality of redundancy is more important than its mere topological presence. In particular, an alternative connection may exist structurally without being physically independent or having sufficient capacity to supply the affected demand after an outage.

The structural assessment shows that criticality in a predominantly radial distribution network cannot be captured by a single decision rule alone. The $N - 1$ contingency analysis identifies individual assets whose removal causes islanding and loss of weighted connectivity, while the targeted attack simulations show how network performance deteriorates under sequential removals. Nodes generally represent more effective targets than individual connections, as their removal can disconnect multiple downstream branches simultaneously.

The targeted attack simulations further show that the most effective decision rule depends on the attacker's objective and available information. Impact-informed IEMDE causes the largest initial reduction in remaining weighted connectivity for both nodes and edges. However, additional impact information does not necessarily result in faster overall fragmentation. For nodes, IEMDE and OD-betweenness both achieve complete fragmentation after removing 27% of the candidate nodes. For edges, OD-betweenness reaches the maximum observed fragmentation after 41% of the candidate removals, compared with 54% for IEMDE. These findings demonstrate that source-to-demand topological information alone can already provide a strong indication of structural criticality, while impact information is particularly valuable when the objective is to maximise immediate weighted disruption.

The central finding of this research is that asset-specific vulnerability and structural criticality identify different priority assets. The degree of overlap between the highest-ranked assets depends on the assessment method and asset type. Assets that are structurally critical are therefore not necessarily the assets that are most vulnerable to deliberate physical attacks, and vice versa. The structural assessment identifies assets whose removal is particularly consequential for network connectivity, whereas the VI evaluates asset-specific characteristics associated with vulnerability to deliberate physical threats. Structural criticality should therefore not be used as a direct proxy for asset-specific vulnerability.

For Stedin, the findings show that the two frameworks should be used together rather than as substitutes. Particular attention should be given to ensuring that redundant connections are physically

independent and that sufficient $N - 1$ capacity redundancy is maintained. The VI additionally identifies vulnerable assets for which targeted protection measures can be considered, including physically separating redundant routes, assessing the feasibility of undergrounding exposed overhead sections, and addressing deficiencies in physical security measures.

Overall, this research demonstrates that vulnerable and critical assets in electricity subtransmission networks are best identified by assessing asset-specific vulnerability and structural criticality separately and subsequently combining their insights. The VI identifies which assets are particularly vulnerable based on their asset-specific characteristics and potential impact, while the structural assessment identifies which assets are most consequential when removed and how this changes under different attacker information assumptions. Together, the frameworks distinguish which assets are vulnerable, which are structurally critical, and where both characteristics coincide. This provides Stedin, and potentially other DSOs, with a more complete basis for prioritising resilience and security measures against deliberate physical threats. Ultimately, strengthening resilience starts with understanding where vulnerabilities exist, why they exist, and what their consequences may be.

Contents

Preface	i
Executive Summary	ii
Nomenclature	viii
1 Introduction	1
1.1 Reading guide	2
2 Literature review	3
2.1 Hazard perspective on electricity networks	3
2.1.1 Types of attackers	3
2.1.2 Attack strategies	4
2.2 Asset vulnerability and structural criticality	4
2.3 Vulnerability and criticality assessment methods	6
2.3.1 Asset vulnerability assessment	6
2.3.2 Structural criticality assessment	7
3 Research Design	9
3.1 Knowledge gap and research questions	9
3.2 Research approach	10
3.3 Research method	11
3.3.1 Data requirements and tools	11
3.3.2 Validation and verification	12
3.4 Case study: Stedin network	14
4 Vulnerability Framework	17
4.1 Vulnerability identification at DSOs	17
4.2 Identification of malicious threat specific characteristics	18
4.2.1 Weighting of characteristics	20
4.2.2 Operationalisation to vulnerability assessment framework and scores	21
5 Graph-theoretic assessment of structural criticality	25
5.1 Case-study network representation	25
5.1.1 Local node weights	25
5.2 Selection of assessment methods	26
5.2.1 Centrality-based decision rules	26
5.2.2 Weighted Connectivity Loss	28
5.3 Single-asset contingency analysis	28
5.4 Sequential targeted attack simulations	29
5.4.1 Iterated Most Central Element	30
5.4.2 Iterated Electrical Most Damaging Element	30
5.5 Global robustness consistency check	30
6 Results	32
6.1 Vulnerability Index (VI) scores	32
6.1.1 Sensitivity Analysis	35
6.2 Structural criticality scores	36
6.2.1 Targeted attack simulations	36
6.2.2 Contingency Analysis (N-1)	39
6.2.3 Robustness trajectory comparison	41
6.3 Overlap across frameworks	42

7 Discussion	44
7.1 Modelling purpose: bridging structural criticality and asset vulnerability	44
7.2 What the two frameworks reveal about the network	45
7.2.1 The VI covers a broad range of malicious-threat-relevant characteristics across different categories	45
7.2.2 Physical and capacity redundancy is the strongest and most consistent determinant of vulnerability	46
7.2.3 A subset of high-scoring edges share overhead and redundancy characteristics	46
7.2.4 The most effective attack strategy depends on the attacker's objective	47
7.2.5 Node removal causes more damage to the network than edge removal	48
7.2.6 The $N - 1$ assessment reveals a physical redundancy gap	48
7.2.7 The relation between vulnerability and structural criticality depends on the assessment perspective	48
7.3 Academic contribution and methodological reflection	49
7.3.1 Scope of the asset-specific vulnerability assessment	49
7.3.2 Choice of structural decision rules	49
7.3.3 Scope of the removal-based assessments	50
7.3.4 Positioning the methodology within broader power-grid vulnerability research	51
7.4 Limitations	52
7.4.1 Methodological limitations	52
7.4.2 Data limitations	53
7.4.3 Technical limitations	53
7.5 Implications for Stedin	53
7.5.1 Implications for DSOs	54
7.6 Future research	54
8 Conclusion	56
8.1 Societal relevance	57
References	58
A AI Statement	61
B Model Choices	62
C Network Characteristics	69
D Vulnerability Index	73
E Targeted Attacks	76
F Contingency Analysis	81

List of Figures

2.1	Radial/tree-like and meshed electricity-network topologies.	5
2.2	The upper-right quadrant represents assets that combine a high susceptibility to successful sabotage with severe consequences for network functioning.	6
3.1	Research methodology	13
3.2	Graph representation of the case study.	15
3.3	Primary substation of 150/50/13 kV to illustrate different compartments of the nodes (Phase to Phase, n.d.; own interpretations)	16
4.1	Continued functioning of the subtransmission network (ST) due to the presence of a redundant connection (N-1)	19
4.2	Shift in feed area following the redundancy recalculation. The failure of the line shown in red causes a subtransmission station to be fed through a different upstream 150 kV node.	23
5.1	Overall methodological framework for the vulnerability and structural criticality assessments.	31
6.1	Relative VI decrease within nodes and edges	34
6.2	Network F , represented as a graph constructed in NetworkX.	36
6.3	Targeted attack strategies and remaining weighted connectivity.	38
6.4	Node contingency analysis on Network F , showing (a) WCL and (b) the number of isolated nodes.	40
6.5	Edge contingency analysis on Network F , showing (a) WCL and (b) the number of isolated nodes.	41
6.6	Comparison of the normalised distance-based EGR proxy and $1 - \text{WCL}$ along the IEMDE node-removal sequence.	42
B.1	Meshed vs Radial busbar explanation for nodes (AI generated)	67
B.2	Overhead lines representation abstract and real-time	68

List of Tables

4.1	Overview of vulnerability categories considered in this research.	18
4.2	Overview of the characteristics included in the vulnerability index and their role in the assessment framework.	19
6.1	Top 10 nodes ranked by Vulnerability Index.	33
6.2	Top 10 edge groups ranked by Vulnerability Index.	33
6.3	Top 5 centrality rankings for Network F	37
6.4	Top 10 edges ranked by all vulnerability frameworks.	43
6.5	Top 10 nodes ranked by all vulnerability frameworks.	43
7.1	Overview of the main findings organised by discussion theme.	45
B.1	Overview of the main modelling choices and abstractions.	63
C.1	Overview of vulnerability characteristics used in this research.	70
C.2	Operationalisation of the vulnerability characteristics included in the Vulnerability Index.	72
D.1	Complete overview of node vulnerability scores.	73
D.2	Complete overview of edge vulnerability scores.	74
D.3	Statistical relationship between the node characteristics and the Vulnerability Index.	75
D.4	Statistical relationship between the edge characteristics and the Vulnerability Index.	75
D.5	Sensitivity analysis of the Vulnerability Index ranking.	75
E.1	Top 15 nodes by betweenness and degree centrality (Network F).	76
E.2	Top 15 branches by edge betweenness centrality (Network F).	76
E.3	IMCE results for node removal using betweenness centrality, Network F	77
E.4	IMCE results for node removal using degree centrality, Network F	77
E.5	IEMDE results for node removal, Network F	78
E.6	IMCE results for edge removal using edge betweenness centrality, Network F	78
E.7	IEMDE results for edge removal, Network F	79
F.1	Branch (edge) contingency results (Network F).	81
F.2	Station (node) contingency results (Network F).	82

Nomenclature

Definitions

Term	Definition
Transmission grid	High-voltage network that transports bulk electrical power over long distances and supplies regional substations.
Distribution grid	Medium- and low-voltage network (and, depending on the DSO, parts of the subtransmission network) that distributes electricity from the transmission grid to consumers.
Asset	A physical component of the electricity network considered in this study, consisting of either a station (node) or a connection (edge).
Critical asset	An asset whose loss would have the largest structural consequence to the network, regardless of how likely that loss is.
Exposure	The component of the Vulnerability Index describing how likely an asset is to be targeted, expressed through characteristics that determine an asset's accessibility to a potential attacker.
Impact	The component of the Vulnerability Index describing the severity of the consequences once an asset is successfully attacked, reflecting the asset's structural and topological importance and the redundancy available to it.
Islanding	The disconnection of one or more load-bearing nodes from all available sources following the removal of an asset, resulting in a Weighted Connectivity Loss greater than zero.
Malicious threat	A deliberate, physical attack carried out by a saboteur aiming to disrupt the (electricity) network.
Redundancy	The presence of one or more alternative paths or installations that allow the network to continue functioning after the failure of a single asset.
Redundant circuit	An additional circuit that provides an alternative connection within the network. This may represent an N-1 redundant path within a single forecast area or an energized circuit that normally carries no load and connects different forecast areas, thereby providing additional network flexibility.
Robustness	The ability of a system to continue performing its intended function under unforeseen disturbances.
Structural criticality	The extent to which the loss of an asset disrupts network connectivity due to its position within the network topology, used to identify which assets are structurally critical independent of the likelihood of an attack.
Vulnerability	A weakness or limitation within a system that can impair its ability to perform its intended function.
Vulnerable asset	An asset that is particularly susceptible to failing or being compromised, as identified through its Exposure and Impact scores in the Vulnerability Index.

Abbreviations

Abbreviation	Definition
DSO	Distribution System Operator
VI	Vulnerability Index
N-1	Redundant criterion requiring that the network remains operational after the failure of any single asset
OD	Origin-Destination
IEMDE	Iterated Electrical Most Damaging Element
IMCE	Iterated Most Central Element
WCL	Weighted Connectivity Loss
HILP	High-Impact Low-Probability
EGR	Effective Graph Resistance
OPF	Optimal Power Flow

Symbols

Symbol	Definition	Unit
Sets and general notation (Section 5.1)		
V	Set of nodes (stations)	—
v	Node (station)	—
E	Set of edges (connections)	—
e	Edge (connection)	—
A	Set of all assets, all nodes and edges ($A = V \cup E$)	—
S	Subset of source nodes in all nodes, $S \subseteq V$	—
D	Subset of load-bearing (demand) nodes in all nodes, $D \subseteq V$	—
F	Structural multigraph representation of the case-study network, $F = (V, E)$	—
M	Adjacency matrix of the multigraph	—
m_{ij}	Element of M : number of edges between nodes i and j	—
Local node weights (Section 5.1.1)		
w_i	Local weight of node i , used in the WCL calculation	GWh
$L_{\text{station},i}$	Cumulative annual energy value associated with station node i	GWh
$\text{Ch}(i)$	Set of direct downstream child nodes of node i	—
$L_{\text{child},j}$	Cumulative annual energy value of child node j	GWh
Centrality measures (Section 5.2.1)		
k_i	Degree centrality of node i	—
B_i^{OD}	Origin-destination betweenness centrality of node i	—
B_e^{OD}	Origin-destination betweenness centrality of edge e	—
$\sigma_{s \rightarrow t}$	Number of shortest paths from source node s to demand (target) node t	—
$\sigma_{s \rightarrow t}(i)$	Number of shortest $s \rightarrow t$ paths passing through node i	—
$\sigma_{s \rightarrow t}(e)$	Number of shortest $s \rightarrow t$ paths passing through edge e	—
Connectivity loss and contingency analysis (Sections 5.2.2, 5.3)		
Γ	Network state after removal of an asset	—
$R(\Gamma)$	Subset of nodes reachable from at least one source in state Γ , $R(\Gamma) \subseteq V$	—
$\text{WCL}(\Gamma)$	Weighted Connectivity Loss of network state Γ	%
$R(F - v)$, $R(F - e)$	Reachable node set after removing node v / edge e from F	—

Symbol	Definition	Unit
Sequential targeted attack simulation (Sections 5.4, 5.4.1, 5.4.2)		
A_{att}	Candidate asset set for the attack (V for node attacks, E for edge attacks)	–
Γ_r	Network state after r asset removals, $\Gamma_0 = F$	–
C_r	Candidate set of assets not yet removed at iteration r	–
a_r	Asset selected for removal at iteration r	–
x	Candidate asset evaluated during selection	–
$\text{Remove}(\Gamma_r, a_r)$	Removal of asset a_r from network state Γ_r	–
ϕ	Centrality metric used as the decision rule in IMCE	–
$\tilde{z}_r$	Marginal increase in WCL caused by removing a_r at iteration r	%
Vulnerability Index: Exposure and Impact (Section 4.2.2)		
c	Exposure characteristic index	–
q	Index of an individual crossing segment for edge e	–
$l_{e,q}$	Length of segment q of edge e coinciding with the relevant infrastructure	m
l_e	Total length of edge e	m
$E_{e,c}, E_{v,c}$	Normalised score (0–1) of exposure characteristic c for edge e / node v	–
w_c	Weight assigned to characteristic c	–
E_e, E_v	Exposure score of edge e / node v	–
I_e, I_v	Impact score of edge e / node v	–
VI_e, VI_v	Vulnerability Index of edge e / node v	–
L_e, L_v	Annual energy associated with edge e / supplied through node v	GWh
$\mathbb{I}(P, C)$	Indicator function: 1 if $P > C$, 0 if $P \leq C$	–
t	Hour of the year	–
V_{area}	Set of substations within forecast area	–
$P_{v,t}$	Forecast load of substation v at hour t	MW
C_{area}^{N-1}	Available N-1 capacity of the supplying 150 kV station for forecast area	MW
$R_{N-1,\text{area}}$	N-1 redundancy of forecast area a (share of the year without N-1 capacity exceedance)	–
d_e	Distance between redundant connections	m
s_v	Categorical scale for physically separated redundant switchgear (1 = physically separated, 2 = shared room, 3 = single installation)	–

1

Introduction

In the early hours of 3 January 2026, residents of southwest Berlin found themselves stuck in lift shafts as the power cut out. Elsewhere in the area, a student staring down a thesis deadline watched their laptop battery drain with no way to recharge it. Unknown criminals had set fire to a cable bridge carrying high- and medium-voltage lines over a canal, destroying five 110 kV and ten 10 kV cables and leaving more than 40,000 households and 2,000 businesses without electricity for nearly four days (Proximities, 2026).

What stands out is how little specialised knowledge such an attack required. Experts have pointed out that critical nodes in energy infrastructure are often relatively easy to identify, and that sabotage can be carried out with simple means such as firebombing. The impact did not come from technical complexity, but from striking physical connections that turned out to be far less robust than assumed. This single bridge, carrying a disproportionate share of Berlin's cabling, proved to be a critical point in the city's electricity network. The question this raises is not just how the attack was carried out, but why the network was vulnerable enough to allow it: what made this specific point so critical, and could it have been identified in advance? (Proximities, 2026)

The electricity grid is becoming increasingly critical for modern society. Essential infrastructures such as hospitals, transport systems, communication networks and water supply depend on a continuously functioning electricity system (Ministerie van Justitie en Veiligheid, 2025b). As a result, the consequences of disruptions extend far beyond the electricity sector itself (Lange, 2025). A failure in the grid can trigger cascading failures across connected sectors simultaneously (Ministerie van Justitie en Veiligheid, 2025a).

This growing dependency makes the electricity grid increasingly important to protect (Techniek Nederland, 2026). This is of societal relevance: the more daily life depends on the grid, the greater the impact of a disruption, and the greater the risk that the grid becomes a target for malicious actors. This shifts the focus not only towards expansion of the grid, but towards a resilient and robust grid (Ministerie van Justitie en Veiligheid, 2026).

Traditionally, research has focused on resilience against natural disturbances, such as extreme weather events and technical failures (Belcic & Smalley, 2026). However, recent incidents, including the sabotage event in Berlin, demonstrate that electricity networks are also vulnerable to deliberate physical attacks. Together with the sabotage of subsea power and data cables in the Baltic Sea in December 2024 (Lange, 2025), these incidents illustrate that the consequences of a disruption depend not only on the nature of the disturbance, but also on the structural and asset-specific characteristics of the affected network (Proximities, 2026). The Berlin incident is particularly illustrative in this regard, as the attack targeted an asset whose physical characteristics made it susceptible to sabotage.

This threat is not confined to Berlin or the Baltic Sea. In September 2026, a notebook recovered from the suspect behind a series of substation attacks in Brandenburg and North Rhine-Westphalia was found to list an electrical station in Doetinchem, in the Netherlands, as a potential target ("Duitse

klimaatactivist wil ook in Nederland toeslaan op energienet”’, 2026). The Dutch grid is therefore not a hypothetical case for this type of threat. Identifying why critical points exist, and where they are located within a network, is consequently not only an academic question but also a practical necessity for grid operators.

This research focuses on how to identify vulnerable and critical assets within the Dutch electricity grid. It does not start from possible disruptions. Instead, it examines which network components are critical to the grid’s functioning. It also examines which components are vulnerable to deliberate attack. Based on this, the research investigates how these assets can be identified and mapped.

This leads to the following main research question:

How can vulnerable and critical assets in an electricity network under physical malicious threats be identified?

The objective of this research is to analyse how vulnerability emerges in electricity networks, which characteristics contribute to it, and whether assets that are structurally critical are inherently more vulnerable to deliberate physical threats. By combining a structural criticality assessment with a vulnerability assessment, this research provides a more comprehensive understanding of network vulnerability and establishes a foundation for future research into risk reduction and resilience enhancement.

This objective aligns closely with the MSc CoSEM programme, which focuses on the current and future design of energy systems. The programme is concerned with analysing and improving existing systems so that they can keep pace with society’s growing dependency on them, as outlined above. This requires not only technical understanding of the physical network, but also insight into the actors, such as network operators, security teams and decision-makers, who must act on that understanding. By analysing how vulnerable and critical assets in an electricity network can be identified, this research addresses exactly the kind of system-level analysis the CoSEM programme is centred on, and forms the necessary basis on which any subsequent redesign, investment decision or policy response by these actors would have to build.

1.1. Reading guide

Chapter 2 presents the literature review. It introduces the hazard perspective on electricity networks, distinguishes between asset vulnerability and structural criticality, and reviews the methods used to assess both perspectives. Chapter 3 describes the research design. It identifies the academic knowledge gap, formulates the research questions, explains the research approach and methodology, and introduces the Stedin case-study network. Chapter 4 presents the asset-specific Vulnerability Index (VI) framework. It discusses the identification of malicious-threat-specific vulnerability characteristics, their classification into Exposure and Impact dimensions, and their operationalisation into a quantitative vulnerability assessment framework. Chapter 5 presents the graph-theoretic framework for assessing structural criticality. It describes the network representation, the selected centrality measures, the contingency-analysis approach, and the targeted-attack simulations used to identify structurally critical assets. Chapter 6 presents the results of both frameworks. It reports the Vulnerability Index scores, the sensitivity analysis, the structural criticality assessment results, and the overlap between vulnerable and structurally critical assets. Chapter 7 discusses the findings and their implications. It reflects on what the two frameworks reveal about network vulnerability and structural criticality, evaluates the methodological choices and limitations, discusses implications for Stedin and other DSOs, and outlines directions for future research. Chapter 8 concludes the thesis by answering the subquestions and the main research question, and reflects on the societal relevance of the research. The appendices provide supporting material, including the AI statement, modelling choices, detailed vulnerability scores, targeted-attack results and contingency-analysis results.

2

Literature review

To properly frame the problem introduced in Chapter 1, this chapter reviews how vulnerability and criticality in electricity networks are addressed in the literature. The review follows a funnel structure. Section 2.1 first defines the threat considered in this research by examining the relevant hazard, attacker types, and attack strategies. Section 2.2 then distinguishes between two perspectives on the network: asset vulnerability and structural criticality. Finally, Section 2.3 reviews how both perspectives can be assessed.

2.1. Hazard perspective on electricity networks

Vulnerability starts from a hazard to which a system is exposed. The literature considers different sources of disruption depending on the research context. Holmgren (2007a) provides an overarching classification for electricity networks and distinguishes natural disasters, technical failures, and intentional attacks by saboteurs. This research focuses on the latter category and, more specifically, on deliberate physical sabotage, as introduced in Chapter 1. A similar physical malicious-threat perspective is adopted by Bilis et al. (2013).

Other vulnerability studies consider different hazards. Salman et al. (2025), for example, focus on hurricane impacts and vegetation-related failures, while Montalà-Palau et al. (2026) consider geographically varying hazards such as forest fires, heatwaves, and underground machinery. These studies remain relevant because they illustrate that the characteristics determining vulnerability depend on the hazard under consideration. An asset that is vulnerable to a storm could also be vulnerable to deliberate sabotage. Take the example of a tree falling on power lines; the same effect would occur if overhead power lines were cut.

Since the present study fixes on physical sabotage as its threat, the remainder of this section focuses on the saboteur. First, different attacker types are distinguished according to the information available to them. Second, the attack strategies through which this information can be translated into target selection are discussed.

2.1.1. Types of attackers

A central distinction in malicious-threat studies concerns the knowledge available to an attacker. Holmgren (2007a) distinguishes between an *opportunistic* and a *determined* saboteur. The opportunistic saboteur has little knowledge of the power system and may act irrationally, whereas the determined saboteur acts rationally in selecting a target. Holmgren additionally considers an informed *antagonist* in a game-theoretic setting, in which the attacker strategically selects actions in response to the defender.

A related distinction is made by Bilis et al. (2013), who consider both *deterministic* targeted attacks and stochastic *random* attacks. In targeted attacks, critical network elements are deliberately selected, whereas random attacks do not rely on such prior identification.

For deliberate targeted attacks, being determined does not necessarily imply complete information

about the electricity system. Albarakati and Bikdash (2022) distinguish attackers according to their access to system information. Network topology and certain electrical characteristics may be obtainable or estimated, whereas detailed system tolerances are considerably harder to obtain. Attackers can therefore differ substantially in the information available for identifying the most damaging targets.

The distinction of Albarakati and Bikdash (2022) looks into two types of determined saboteurs. An *uninformed determined saboteur* deliberately seeks to maximise disruption but bases target selection predominantly on geographical and topological information about the network. Since “uninformed” here refers to knowledge of power flow, this attacker uses topological metrics as his decision rule to identify vital “hubs” or chokepoints in the graph. An *informed determined saboteur* additionally has access to asset- and system-specific information that allows the actual consequence and impact of different attacks to be evaluated. Here, the components are strategically selected based on which ones would cause the most electrical damage if they failed. Both attackers are therefore strategic; the distinction concerns the information available for determining which target is most damaging.

This knowledge is particularly relevant for physical attacks because network topology may be easier to reconstruct than detailed operational information. Cetinay et al. (2018), for example, note that information on the topology of electricity networks may be obtainable from publicly available sources. Structural information can therefore already provide an attacker with a basis for targeted asset selection, even without knowledge of system operation.

2.1.2. Attack strategies

Sabotage can also differ in how assets are attacked, differentiating from the attackers knowledge. Broadly, the literature considers single or simultaneous and sequential targeted attacks.

Single and simultaneous attacks

A saboteur may attack one asset at a time, or several assets simultaneously. In the literature, single-asset failures are commonly analysed through $N - 1$ contingency analysis, in which each component is removed independently to assess whether the network can continue functioning after its loss (Coelho et al., 2019; Narimani et al., 2020). This provides an indication of the network’s resilience to individual failures and reveals which single outages cause the largest consequences.

The same principle can be extended to $N - x$ contingencies, in which multiple assets fail simultaneously. Narimani et al. (2020) show that such combinations can cause severe system impacts even when a network is resilient to individual $N - 1$ failures.

Sequential attacks

A determined saboteur aiming to cause maximum damage may instead attack assets sequentially. After each removal, the network changes and the next most attractive target can therefore also change. Cetinay et al. (2018) model such targeted attacks by recalculating topological centralities of assets after every removal.

Albarakati and Bikdash (2022) further distinguish between a centrality-based strategy, in which the most central remaining asset is attacked, and an impact-based strategy, in which the asset causing the largest decrease in system functioning is selected. These strategies reflect different levels of attacker knowledge and are therefore directly related to the attacker types introduced above.

2.2. Asset vulnerability and structural criticality

The literature uses the term *vulnerability* in several ways. Particularly within complex-network research, the term *structural vulnerability* is frequently used to describe how the topology of a network creates critical points. Forsberg et al. (2023), for example, assess the structural vulnerability of the Nordic transmission grid through centrality measures with the specific objective of identifying its most critical nodes. Similarly, Koc et al. (2015) analyse structural vulnerability in electricity distribution grids based on the dependency created by the network topology.

For clarity, this thesis refers to this perspective as *structural criticality*. The terminology does not imply that studies referring to structural vulnerability assess a fundamentally different phenomenon. Instead, it is used here to maintain a clear distinction between two questions central to this research:

- **Asset vulnerability:** How susceptible is an individual asset to a successful physical malicious attack, given its asset-specific characteristics?
- **Structural criticality:** How severe are the network consequences if that asset is successfully taken out of service?

Structural criticality is therefore defined by the impact associated with losing an asset, rather than by the specific input used to identify that asset. A critical asset may be identified directly from its topological position, but also through contingency analysis or through an attack strategy that evaluates the resulting disconnected demand. This distinction is relevant because not every criticality assessment relies exclusively on centrality or topology as its decision variable.

Structural criticality is strongly influenced by network topology. Distribution grids are commonly organised in radial or tree-like structures, whereas transmission grids are generally more meshed. In a radial network, downstream assets depend on a limited number of upstream paths. The loss of an upstream component can therefore disconnect a substantial part of the network. In a meshed network, alternative routes may remain available after an individual component is lost. Koc et al. (2015) consequently emphasise the importance of the number and independence of paths connecting distribution-grid components to their source. Figure 2.1 illustrates this distinction of different network topologies, ranging from radial/tree-like structures (left) with minimal redundancy to increasingly meshed structures with multiple alternative paths, mostly used in transmission grids (right) (Phase to Phase, n.d.).



Figure 2.1: Radial/tree-like and meshed electricity-network topologies.

Asset vulnerability takes a different starting point. Rather than asking how the position of an asset influences network functioning, it examines characteristics of the component itself. Importantly, these characteristics depend on the threat being considered. Salman et al. (2025) and Zeh et al. (2021), for example, incorporate vegetation exposure and pole fragility when considering storms and tree failures. Other approaches consider characteristics such as asset condition, geographical exposure, or environmental hazards.

Montalà-Palau et al. (2026) illustrate that both perspectives can also be combined. Their approach assigns geographically varying hazard information to individual network elements and subsequently evaluates consequences using power-system modelling. Asset vulnerability and structural criticality should therefore not be interpreted as mutually exclusive properties. An individual asset can score high or low on both dimensions.

For malicious physical attacks, this distinction has a direct interpretation from the perspective of the attacker. Asset-specific characteristics influence whether a target can realistically be attacked or compromised, whereas network characteristics determine how consequential the successful loss of that target would be. An asset of particular concern is therefore one that combines high asset vulnerability with high structural criticality.

A comparison can be drawn with a railway network. Consider, for example, a station located underground. From an asset vulnerability perspective, its physical location may make it relatively difficult for a saboteur to access and successfully attack. However, this provides little information about its structural criticality: the consequences for the wider railway network if that station were taken out of service. Conversely, a highly connected station such as Utrecht may be structurally critical because its disruption affects many connections within the network, while this does not necessarily imply that it is particularly vulnerable in terms of physical accessibility, exposure, or security. This illustrates that asset vulnerability and structural criticality describe two distinct characteristics of an asset.




Figure 2.2: The upper-right quadrant represents assets that combine a high susceptibility to successful sabotage with severe consequences for network functioning.

2.3. Vulnerability and criticality assessment methods

The two perspectives introduced above require different assessment methods. Asset vulnerability assessments predominantly evaluate characteristics of individual components, whereas structural criticality assessments evaluate the importance of components from the perspective of the network. The literature provides several approaches for both.

2.3.1. Asset vulnerability assessment

Asset vulnerability assessments evaluate component-specific characteristics that influence how vulnerable an asset is to a particular hazard. These characteristics are inherently hazard-specific: assets occupying similar positions in a network may differ substantially in their vulnerability because of differences in physical characteristics, exposure, or security (Hromada et al., 2021). Consequently, such assessments generally combine multiple asset characteristics into a composite score or index. Liu et al. (2025), for example, distinguish Danger, Susceptibility, Resistance, and Recoverability as dimensions contributing to vulnerability.

Vulnerability and risk assessments commonly combine the likelihood of an adverse event with its consequences. Holmgren (2007a) formalises vulnerability as the probability that a disturbance results in a negative consequence Q exceeding a critical threshold q . Similarly, Montalà-Palau et al. (2026) assess individual network elements by combining hazard occurrence, vulnerability likelihood, and impact. In abstract form, their risk formulation can be represented as:

$$R = P(h) \times L \times I$$

Although these approaches conceptualise vulnerability and risk differently, they illustrate the same broader distinction between the likelihood of an adverse event and the consequences resulting from it. The characteristics determining this likelihood depend on the hazard under consideration. For deliberate physical sabotage, for example, physical exposure or protection may influence the feasibility of successfully compromising an asset, while the impact dimension represents the consequences if that asset is successfully taken out of service.

Asset vulnerability assessment therefore provides an asset-specific perspective that cannot be derived from network topology alone. It evaluates which assets are particularly vulnerable to the specified hazard and why, complementing the structural criticality assessments discussed in the following section.

2.3.2. Structural criticality assessment

Structural criticality assessments can broadly be divided into methods that infer criticality from properties of the existing network and methods that evaluate criticality by simulating component removals. Following the distinction described by Shafiei et al. (2025), these are referred to here as *a-priori* and *a-posteriori* approaches. In this context, *a-priori* metrics assess the inherent criticality and physical exposure of grid assets before an event, while *a-posteriori* metrics quantify the resulting systemic damage and connectivity loss as disruptions propagate through the network.

A-priori: topological, electrically informed, and spectral metrics

A-priori methods estimate the potential importance of network components from the network as it exists before a specific disruption is simulated (Shafiei et al., 2025). They provide baseline indicators that can be used to rank critical assets or to define target-selection rules for later contingency or attack simulations. In the context of power-grid vulnerability, these methods can be divided into three broad categories: purely topological metrics, electrically and operationally informed metrics, and global network metrics, including spectral metrics. These categories may overlap because some metrics combine more than one type of information.

Purely topological metrics represent the power system as an unweighted graph $F(V, E)$. Degree centrality measures the number of direct connections of a node and therefore captures local connectivity. Betweenness centrality measures how often a node lies on shortest paths between other nodes, whereas closeness centrality measures the distance between a node and the rest of the network (Cetinay et al., 2018; Pagani & Aiello, 2015). These metrics can identify structurally important components, but they do not directly account for electrical loading, power-flow conditions, or the physical characteristics of the grid.

Electrically and operationally informed metrics extend the network representation with information such as line reactance, susceptance, admittance, power flow, or electrical demand. For example, Electrical Node Significance (δ_i) weights a node according to the amount of power distributed by that node relative to the total power distributed in the grid. Nodes associated with larger amounts of distributed power therefore receive a higher electrical significance value (Koc et al., 2013). Other electrically informed metrics describe electrical distance or the quality of alternative paths. Electrical closeness centrality uses effective resistance as an electrically meaningful distance measure, whereas flow betweenness is based on the distribution of electrical flows through the network (Cetinay et al., 2018). Effective Graph Resistance (EGR) aggregates the effective resistance between all pairs of nodes and accounts for both the number and quality of alternative paths (Ellens et al., 2011). Forsberg et al. weight centrality measures using the series reactance of transmission lines and find that their weighted combined measure performs best for identifying critical nodes in the Nordic transmission grid (Forsberg et al., 2023).

A third group consists of global network metrics. Rather than ranking individual assets, these metrics produce a single score that describes the network as a whole, often by aggregating the same information used in the metrics above. Algebraic connectivity is one example: a higher value indicates a network that is harder to break apart (Koç et al., 2014). EGR, introduced above, can be read the same way, since it aggregates redundancy across all node pairs into one network-wide number.

The metrics groups discussed above also differ in their level of analysis. Centrality measures and Electrical Node Significance rank individual components. Upstream Robustness also evaluates individual assets by measuring the robustness of their connection to the source based on the number, disjointness, and length of alternative paths (Koc et al., 2015). EGR and algebraic connectivity instead describe the network as a whole. Because these metrics operate at different levels, a global network score should not be interpreted as a direct ranking of individual asset criticality without an additional component-level analysis.

Which metrics are most informative depends partly on network topology. In radial networks, upstream dependency, single points of failure, and downstream connectivity are particularly relevant because the loss of an upstream component may disconnect an entire downstream region (Koc et al., 2015). In meshed networks, alternative paths and electrically informed distance measures become more relevant because network redundancy and multiple paths can reduce the effect of a component outage (Cetinay et al., 2018; Ellens et al., 2011).

A-posteriori: removal-based simulations

A-posteriori approaches assess criticality by removing one or more components and evaluating the resulting network functioning. Rather than estimating importance from the original network, these approaches directly observe the consequence of component loss.

Contingency analysis evaluates predefined component outages. Under an $N - 1$ assessment, each component is removed individually and the resulting system state is evaluated (Coelho et al., 2019). In radial networks, relevant consequences include islanding and disconnected downstream demand, whereas in meshed networks an outage may instead redistribute power flows and create overloads elsewhere (Koc et al., 2015; Koç et al., 2014).

Targeted attack simulations extend this principle by introducing an attacker decision rule and, particularly in sequential attacks, repeatedly updating the network after each removal. Cetinay et al. (2018) use centrality measures to determine successive attack targets, while Albarakati and Bikdash (2022) compare such centrality-based strategies with an impact-based strategy that selects the element causing the largest increase in unsatisfied load.

A-priori and a-posteriori approaches are therefore complementary rather than independent. A-priori metrics provide an indication of which assets are expected to be critical, while removal-based simulations evaluate the actual consequences of their loss. Moreover, an a-priori metric can itself be used as the target-selection rule within an a-posteriori attack simulation. The term structural therefore refers to the network-level consequence of asset loss, and includes both topological metrics and electrically informed in simulation-based methods.

3

Research Design

Chapter 2 identified two complementary perspectives for analysing vulnerable and critical assets in electricity networks: asset-specific vulnerability and structural criticality. It also showed that no existing study combines both perspectives for the same network under the threat of deliberate physical sabotage. This chapter translates this knowledge gap into a concrete research design.

Section 3.1 formulates the knowledge gap and resulting research questions. Section 3.2 outlines the research approach adopted to answer these questions. Section 3.3 describes the research method, including the data requirements and tools used for each subquestion (Section 3.3.1) and the validation and verification strategy applied throughout the research (Section 3.3.2). Finally, Section 3.4 introduces Stedin's subtransmission network as the case study in which both frameworks are applied.

3.1. Knowledge gap and research questions

The literature review revealed two complementary perspectives for identifying vulnerable and critical assets in electricity networks. Structural criticality assessments identify assets that are important for network functioning. These range from a-priori centrality and network-based metrics to a-posteriori contingency and targeted attack simulations (Albarakati & Bikdash, 2022; Cetinay et al., 2018; Coelho et al., 2019; Forsberg et al., 2023; Shafiei et al., 2025). These approaches identify critical bottlenecks and evaluate the consequences of component removal. However, they do not account for the physical asset characteristics that determine whether an asset is particularly susceptible to a specific threat.

Asset-level vulnerability assessments provide this complementary perspective. Rather than starting from an asset's position within the network, they assess component-specific characteristics associated with a particular hazard. Existing frameworks identified in this review predominantly address natural hazards, such as storms and vegetation-related failures, technical deterioration, or broader geospatial hazards (Montalà-Palau et al., 2026; Salman et al., 2025; Zeh et al., 2021). Characteristics relevant to deliberate physical sabotage, such as physical exposure, accessibility, protection, and the physical independence of backups, do appear in the literature (Bilis et al., 2013). However, they appear scattered across different hazards and studies. No asset-level framework was identified that combines these characteristics into one quantitative framework specifically for deliberate physical sabotage.

This reveals two related gaps. First, it remains unclear which characteristics are most relevant and how they can be combined and operationalised into a quantitative Vulnerability Index (VI) for deliberate physical sabotage specifically. Second, and more fundamentally, structural criticality and asset-level vulnerability are typically studied separately. It remains unclear whether the assets identified as structurally critical are also the assets identified as most vulnerable, when both perspectives are applied to the same network under the same threat. No study was identified that explicitly compares both perspectives in this manner.

This research focuses specifically on deliberate physical sabotage, as illustrated by the Berlin scenario introduced in Chapter 1. Within this threat, two attacker types are considered: an *uninformed attacker*,

who selects targets based on publicly available network topology only, and an *informed attacker*, who additionally has access to confidential operational data such as load and energy demands (Albarakati & Bikdash, 2022). Random attackers, multiple simultaneous attackers, and a fully omniscient attacker, who would know every asset-level vulnerability characteristic of the Vulnerability Index in advance, fall outside this scope.

The main research question of this thesis is:

How can vulnerable and critical assets in an electricity network under physical malicious threats be identified?

To answer this question, the research is structured into five sub-questions, each representing a step in the modelling process:

- SQ1** Which asset characteristics of electricity connections and stations are associated with increased vulnerability to malicious disruptions, and how can these be operationalised into a quantifiable Vulnerability Index (VI)?
- SQ2** Which assets are most vulnerable according to their Vulnerability Index (SQ1)?
- SQ3** Which of the characteristics incorporated into the Vulnerability Index contribute most to the vulnerability across assets?
- SQ4** Which graph-theoretical metrics and removal-based assessments are suitable for assessing structural criticalities, and which assets do they identify as most critical?
- SQ5** To what extent do the assets identified as structurally critical (SQ4) coincide with the assets identified as most vulnerable by the Vulnerability Index (SQ2)?

Sub-question 2 is taken separately from sub-question 1, whilst in sub-question 4, two questions are combined. This is because sub-question 1 involves building a new index, whereas sub-question 4 reviews the current literature and selects an appropriate metric and simulation for determining criticality.

3.2. Research approach

This research follows a modelling approach to analyse structural criticality and asset vulnerability in electricity networks. The objective is to understand how the topology and asset-specific characteristics of the network relate to structural criticality and asset vulnerability under malicious disruptions.

A case-study will be examined in terms of its asset vulnerability characteristics. These characteristics will form a mathematical index, the *VI*, which will represent the vulnerability of the asset. For structural criticality assessment, the network will be represented as an undirected graph $F = (V, E)$, where V denotes the set of nodes (substations) and E the set of edges (electricity connections). By analysing the structure of this network, it becomes possible to identify structurally critical components whose failure could significantly disrupt the system.

The modelling framework will be applied to a part of the Stedins' electricity network, which functions as the case study. Using real asset data from the network allows the model to be tested on an actual infrastructure system and provides insights into how vulnerabilities arise in practice. The results of the modelling analysis will be used to identify structurally critical and vulnerable assets, based on both their topological position and their asset-specific characteristics.

The advantage of this quantitative modelling approach is that it allows structural criticality and asset vulnerability to be analysed systematically, enabling the identification of structurally critical and vulnerable components that may not be visible through pure structural vulnerability inspection of the network alone (Forsberg et al., 2023). Applying the model to a real-world case study increases the relevance of the research. However, this approach also carries its limitations. First, purely topological graph modelling simplifies the reality of the network: operational aspects such as intraday power flow and restoration capabilities are not fully captured. This is partially mitigated by the load-based removal-based assessments introduced in Section 2.3, but detailed power-flow behaviour remains out of scope (Saei et al., 2025). Second, the asset-specific analysis depends on the availability and completeness of asset data provided by Stedin; gaps in this data may affect the accuracy of the constructed network model.

3.3. Research method

This research combines literature and data analyses to identify asset vulnerability characteristics. The literature analysis is used to identify different types of characteristics that contribute to vulnerability to malicious disruptions. These characteristics are subsequently operationalised into a quantifiable Vulnerability Index (VI) per asset. Independently, structural criticality is assessed on a network model in Python, combining graph-theoretic metrics with removal-based simulations (Section 2.3). Graph theory allows electricity distribution networks to be analysed as complex systems by representing substations as nodes and electricity lines as edges. The resulting VI scores and structural criticality outcomes are compared to determine the extent to which the assets identified as structurally critical coincide with those identified as vulnerable by the VI.

3.3.1. Data requirements and tools

For the first subquestion, *"Which asset characteristics of electricity connections and stations are associated with increased vulnerability to malicious disruptions, and how can these be operationalised into a quantifiable Vulnerability Index (VI)?"*, relevant vulnerability characteristics will be identified through a combination of literature analysis and expert input from within Stedin. The literature analysis draws primarily from ScienceDirect and Google Scholar and is complemented by experts from Stedin's Asset Management team, who provide domain-specific knowledge on technical definitions and asset characteristics not yet covered in academic literature. This step is separated from the modelling phase to ensure that the selection of characteristics is grounded in both academic evidence and practical domain expertise before being operationalised. Characteristics of interest include properties such as overhead placement, bridge crossings, co-location with other infrastructure, line length, and load capacity. Each characteristic is operationalised by normalising its value onto a common scale (e.g. 0–1), a step performed in Excel. The overall structure used to combine these characteristics into a single VI is derived from a combination of academic literature on risk and vulnerability assessment and internal input from Stedin's risk analysis team, who work with their own risk assessment.

For the second subquestion, *"Which assets are most vulnerable according to their Vulnerability Index (SQ1)?"*, the normalised characteristics and VI structure established in SQ1 are applied to real asset data provided by Stedin to calculate a VI score for each asset, using Python and Excel. This dataset, supplemented by Stedin's GIS platform, PowerBI and Databricks, contains the spatial locations of stations and the connections across the distribution network in the electric area, as well as their spatial characteristics. The model is set up for a defined subset of the network, including the connection points to TenneT's transmission network, which functions as Stedin's upstream source. The resulting VI scores are used to rank assets and identify which substations and connections are most vulnerable within the case study network. These scores and the underlying characteristics are subsequently reviewed with Stedin's grid experts, to validate that the characteristics driving the ranking correctly reflect real-world vulnerability.

For the third subquestion, *"Which of the characteristics incorporated into the Vulnerability Index contribute most to the vulnerability across assets?"*, the VI scores and underlying characteristic values established in SQ2 are analysed in SPSS and Python to determine which individual characteristics show a statistically significant relationship with the overall VI. The statistical test used depends on the measurement level of each characteristic gained from data-collection: Pearson's r for ratio-level, linear relationships, Spearman's ρ for ordinal relationships, and Welch's t-test for characteristics with two categories, to compare the mean VI between them. Characteristics are selected as meaningful contributors where $p < 0.05$. No additional data collection is required for this subquestion, as it relies entirely on the outputs of SQ2. The resulting ranking of characteristics is again reviewed with Stedin's grid experts, to confirm whether they agree that the characteristics identified as most influential are indeed the ones that matter most in practice.

For the fourth subquestion, *"Which graph-theoretical metrics and removal-based assessments are suitable for assessing structural criticalities, and which assets do they identify as most critical?"*, structural criticality is assessed using both a-priori graph-theoretic metrics and a-posteriori removal-based simulations (Section 2.3), implemented in Python. Using NetworkX, a Python library, the case-study network is modelled and both types of assessment are computed on it. A literature analysis is first performed to determine which a-priori metrics are best suited to Stedin's radial network, since, as shown in the

literature review, different network topologies require different metrics. This literature search also specifically targets studies applying these metrics to fictional benchmark case-study networks, such as IEEE bus systems, since these networks are often used precisely to test structural metrics under controlled topological conditions. For the a-posteriori assessment, $N - 1$ contingency analysis and sequential targeted attack simulations are applied, using both the *uninformed* and the *informed* attacker decision rules established in Section 2.1.1. The tools used for this subquestion are Stedin's GIS platform, to obtain the network topology of the case study, and Google Scholar and ScienceDirect, to identify suitable graph-theoretic metrics and removal-based assessment methods from academic literature. The network model constructed of the case study functions as the primary structural input for both types of assessment. The output of this step is a ranked list of the most structurally critical nodes and edges in the network, which forms the last input for SQ5.

For the fifth subquestion, *"To what extent do the assets identified as structurally critical (SQ4) coincide with the assets identified as most vulnerable by the Vulnerability Index (SQ2)?"*, the ranked list of structurally critical assets from SQ4 and the ranked list of vulnerable assets from SQ2 are compared directly, by examining the overlap between the top-ranked assets in each list. Strictly speaking, this comparison is not required to answer which assets are critical or which are vulnerable, since SQ2 and SQ4 already identify these independently. However, examining the extent to which they coincide is valuable for Stedin's practical prioritisation: assets that combine high structural criticality with high asset vulnerability represent the clearest priority for resilience investment. No additional data collection or analytical tool is required for this step, as it relies entirely on the outputs of SQ2 and SQ4 placed side by side. The output of this step is an overview of which assets are identified as both critical and vulnerable, which forms the basis for the recommendations provided to Stedin regarding the prioritisation of its resilience investments.

In Figure 3.1, the research methodology is structured into phases and research steps.

3.3.2. Validation and verification

Validation and verification will be performed throughout the research process to increase the credibility of the methodology and its outcomes.

Validation will focus on ensuring that the selected vulnerability characteristics, data collection and modelling choices are technically realistic and relevant for the case study. During the different research phases, intermediate results will be discussed with experts from Stedin in the fields of electricity networks, asset management, risk management and resilience. These discussions will be used to evaluate the relevance of the selected characteristics, the interpretation of the available data and the assumptions underlying the modelling approach before proceeding to subsequent research steps.

Verification will focus on confirming that the proposed methods are implemented correctly. Intermediate calculations, network representations and model outputs will be checked against the available network data and expected network behaviour. Where necessary, assumptions and implementation choices will be refined before continuing with the analysis. These will be tracked and aligned in Appendix B.




Figure 3.1: Research methodology

3.4. Case study: Stedin network

To test these questions, a case study of the grid operator Stedin is used. Stedin is one of the few Dutch operators that also operates a high-voltage level within its network. Parts of Stedins network make use of the >25 kV electricity network, also referred to as the *transportnet* (subtransmission grid) or *tussenspanningsnet* (intermediate-voltage network) in the Dutch electricity grid. For consistency, the term subtransmission network will be used throughout the remainder of this thesis.


The case study focuses on two specific areas within Stedin's service region. These areas were selected for their hierarchical importance within Stedin's overall network, as well as on the recommendation of a Stedin expert, who identified this part of the grid particularly interesting to examine. The reason is that the network consists of areas that would not be fully redundant ($N-1$), and some of the connections would be laid as overhead lines. The subtransmission layer of this network consists of approximately 26 stations and 45 connections. For confidentiality reasons, the exact geographical location and station identities are not disclosed in this research; all asset data used is anonymised and constructed to be structurally representative of the real network, without revealing sensitive infrastructure information.

As outlined in the methodology, the network is modelled using graph-theoretic concepts of nodes (V) and edges (E). In the case of the grid, *nodes* (V) represent the *hoofdverdeelstations* (main distribution stations) and *transformatorstations* (transformer stations). *Opstijgpunten* (riser points) also form part of certain connections between stations but are not modelled as separate nodes (see further below). The edges (E) represent the connections between the main distribution station and transformer station within the research area. These represent underground cables or overhead lines. Together, these form the network F , defined as:

$$F = (V, E)$$

The figure below shows an abstract representation of a part of Stedins' case study network. As discussed earlier, different grids can have different structural configurations; this representation makes clear that Stedin's subtransmission grid is **radial**, or '**star-shaped**'. Adding the transmission grid above this network, a **hierarchical tree structure** is formed. As becomes clear from figure 3.2, multiple lines are placed between stations by Stedin. This is due to the redundancy principle applied within the grid: redundancy can be seen as a 'reserve' and is essential for the functioning of the grid if an asset within the network fails or if is under construction. This rule, commonly referred to as the $N-1$ criterion, requires that the failure of any single asset must never compromise the continued functioning of the surrounding network. Historically, this redundancy was primarily installed as a reserve against technical faults, driven mainly by cost-efficiency considerations, rather than as a deliberate resilience measure. Resilience against intentional sabotage was therefore not part of these original design decisions. In this research we also make a distinction between two types of redundant circuits. One redundant circuits represent the $N-1$ in a power load area, the other represents an 'extra' connection placed by the grid operated between different power load areas.

An example of a redundant circuit is the grid island in the lower-left corner of the geographical model shown below; for this area to remain functional, the network must be able to withstand the failure of a single asset. This illustrates the essence of $N-1$ redundancy. At the top of the figure, the corresponding geographical topology of the grid is shown, with the substations represented as *yellow* nodes and the lines between them as the case-study edges. This graph is converted into a tree structure, shown in the network representation in the lower part of the figure. The *red* and *blue* nodes fictitiously represent the high-voltage grids to which the primary substations are connected. The yellow nodes and black lines correspond to the case-study-specific areas. The grey lines fall outside the scope of this study and are therefore not examined.



Electrical substations consist of multiple compartments, an understanding of which is essential for defining several of the characteristics used later in this thesis. Figure 3.3 shows a fictitious primary substation (connected to the 150 kV grid) consisting of three voltage levels, used here purely for illustrative purposes.

This particular station has three voltage levels; most stations in the network have only two — for example, from 50 to 10 kV, which is referred to as an *onderstation* (secondary substation). The double circles between the black rectangles in the figure represent the transformers of the station: these convert the voltage down to the level required for the cables serving smaller customers. The black rectangles represent the *schakelinstallaties* (switchgear installations), which are essential to the overall functioning of the grid. Each voltage level therefore has its own installation. The figure also shows that each switchgear installation has a *dubbel railsysteem* (double busbar system), the two parallel rails of the same color within each installation, meaning the installation itself is redundant: if one rail fails, the other can take over its function ($N-1$). Which is not visible in this figure, is the physical separation of these redundant compartments.

These components (transformers and switchgear) form the most important station-level properties considered in this case study, as they are vital to the functioning of Stedin's substations. The transformer can be thought of as the heart of the station, while the switchgear acts as its brain.

The network is constructed in stations and connections. Transformers are not modelled as separate nodes, as they are considered part of the station rather than the network's connective structure. Appendix B.1 illustrates the conceptual difference between radial and meshed network topologies. In radial networks, transformers are typically located within the substation, allowing the station to be represented as a single node. In meshed networks, however, transformers may be positioned between switching locations or at intermediate points in the network rather than being confined to a substation. As a result, representing an entire station as a single node is less straightforward, since transformers can form independent network elements within the meshed topology. Similarly, riser points are excluded as nodes, since they form part of the overhead lines and are considered as a characteristic, rather than distinct network components, as illustrated in Appendix B.2a. However, the transformers and switchgear of a substation are included as station characteristics, as will be discussed in 4.




Figure 3.3: Primary substation of 150/50/13 kV to illustrate different compartments of the nodes (Phase to Phase, n.d.; own interpretations)

4

Vulnerability Framework

This chapter answers the first subquestion: which asset characteristics are associated with increased vulnerability to malicious disruptions, and how can these be combined into a quantifiable Vulnerability Index (VI). The general vulnerability landscape at Stedin, from which these characteristics are drawn, is introduced in Section 4.1. This chapter narrows that landscape down to the single threat considered in this research, deliberate physical sabotage, and classifies the resulting characteristics as either Exposure or Impact indicators (Section 4.2). Section 4.2.2 then operationalises these characteristics into VI scores.

4.1. Vulnerability identification at DSOs

As Stedin invests in its resilience strategy, the company is also interested in a broader understanding of the vulnerabilities present within its network. Before the scope of this research is narrowed down to physical sabotage in Chapter 4, a complete picture of all relevant vulnerabilities was mapped using internal Stedin data, resulting in four categories: Topological, Spatial/Physical, Operational, and Security (Table 4.1). Also partly covered by literature as shown in Appendix C.

Stedin's categorisation uses *vulnerability* in a broad sense, covering any aspect that could disrupt the network, regardless of cause, including telecommunication failures and natural hazards alongside malicious threats. This is broader than the asset vulnerability and structural criticality distinction introduced in Section 2.2. Within the VI, the Topological category considers only an asset's position within the network. This differs from how structural vulnerability is treated elsewhere in this study, where structural criticality is assessed through the network-wide consequence of an asset's loss, not through position alone. The Spatial/Physical and Security categories correspond to asset vulnerability, since they describe component-specific susceptibility. The Operational category describes consequence-based characteristics, unserved energy, grid overload, and restoration time, following the same probability-consequence logic that Holmgren (2007a) uses to define vulnerability itself.

Table 4.1: Overview of vulnerability categories considered in this research.

Topological	Spatial / Physical
• Position in network (tree) • Coupling to upstream transmission network	• Physical exposure (open air) • Cable burial depth • Co-location with other infrastructure • Distance to roads • Environmental factors • Physical separation of redundant components
Operational	Security
• Potential Unserved Energy • Grid overload • Restoration time	• Telecommunication disruption • Physical security accessibility • Security measure deficit • Transparency of asset location

This categorisation is supported and jointly drafted by grid experts at Stedin, and partly grounded in academic literature, for this full analytical table see Appendix C.1. The topological category is addressed as an asset's position within the overall topology (e.g., tree-like radial structure), which is shown to determine its criticality (Koc et al., 2015; Saei et al., 2025). The spatial category, including exposure, depth of cable, and co-location with other infrastructure, is reflected in studies that integrate GIS data to capture the physical placement and surroundings of assets (Montalà-Palau et al., 2026; Stedin, n.d.). The operational category, covering unserved energy, grid overload, and restoration time, reflects asset-level consequence, the severity of a disruption at that specific asset, rather than the network-wide consequence captured by structural criticality. This follows the same *probability times consequence* logic that Holmgren (2007a) uses to define vulnerability itself, and aligns with studies that quantify disruption severity in terms of power loss or recovery duration (Montalà-Palau et al., 2026). Finally, the security category, including aspects such as noticeability and physical accessibility, reflects the security-classification perspective provided by Stedin's internal expertise, which is less extensively covered in academic literature but essential for assessing vulnerability to malicious threats specifically (Hromada et al., 2021).

4.2. Identification of malicious threat specific characteristics

Following the Exposure and Impact distinction described in Section 2.3.1, the vulnerability characteristics identified in Section 4.1 are classified according to whether they primarily contribute to the susceptibility of an asset to a successful attack or to the severity of the consequences once an attack occurs. As the hazard in this research is fixed, the probability of the hazard ($P(h)$) is excluded from the characteristics assessment. This research conditions a sabotage attempt already occurring, the VI quantifies vulnerability given that an attempt takes place, rather than the likelihood that sabotage occurs at all.

Exposure is used here as a proxy for the asset-specific conditions that influence its accessibility, reachability, or susceptibility to a deliberate physical attack. Exposure should therefore not be interpreted as an empirically estimated probability that an asset will be attacked. Rather, it represents the characteristics that may affect the feasibility of reaching or compromising an asset under the fixed sabotage threat.

Exposure indicators are generally spatial or security-related in nature. For example, cable sections located on bridges or crossing waterways may be more accessible to a potential attacker than cables located in less accessible environments. Likewise, overhead lines are treated in this study as more physically exposed than underground cables because they are directly accessible from the surrounding environment. These classifications represent the operationalisation used in this study and do not imply that every overhead line or bridge crossing is equally likely to be attacked.

Impact indicators, in contrast, describe the consequences of a successful attack for the functioning of an asset. Characteristics such as capacity-based redundancy, energy associated with an asset, and physically separated redundancy determine the extent to which the network can continue functioning after an asset is removed.




Figure 4.1: Continued functioning of the subtransmission network (ST) due to the presence of a redundant connection (N-1)

An overview of all vulnerability characteristics is provided in Table 4.2 (SQ1). The underlying values, such as security-monitoring levels and detailed scoring criteria, are omitted from the main text because they contain confidential information provided by Stedin. The table presents the characteristics eventually selected for the VI. The operationalisation column describes how each characteristic is measured or classified. The Exposure and Impact columns indicate the primary role of each characteristic in the assessment framework. The Direction column indicates whether a higher value would increase or decrease the vulnerability score.

Table 4.2: Overview of the characteristics included in the vulnerability index and their role in the assessment framework.

Characteristic	Operationalisation	Exposure	Impact	Edge	Node	Direction
N-1 redundancy exceedance (based on yearly peak load)	Percentage of yearly peak-load instances exceeding the N-1 capacity threshold		✓	✓	✓	+
Physically separated redundant switchgear	1 = Two switchgear installations in two physically separated rooms 2 = Two switchgear installations in one room 3 = One switchgear installation only (Worst-case installation/ attack + impact is used)	✓	✓		✓	+
Potential Unserved Energy	Annual energy demand of station or area (GWh)		✓	✓	✓	+
Physical exposure of transformers	1 = Roof and walls 2 = Open air and walls	✓			✓	+
Distance between transformer and station fence	Distance (m) or N.A.	✓			✓	+
Physical security monitoring (real-time)	0 = None 1 = Low 2 = Medium 3 = High 4 = Vital 5 = Critical	✓			✓	-
Security measure deficit	Difference between required and implemented security measures (0-3 scaled)	✓			✓	+
Distance between redundant circuits/ connections	Percentage of redundant cable length located within 1 m of each other	✓	✓	✓		+
Percentage crossing critical infrastructure						
• Length crossing waterways						
• Length crossing bridges	Percentage of total line length	✓		✓		+
Overhead line exposure	Percentage of line length above ground	✓		✓		+

The *N-1 redundancy exceedance* is classified as an Impact indicator. It represents the extent to which

the remaining network does not have sufficient capacity to accommodate the relevant power flow after the loss of one component (Sperstad et al., 2020). Although this characteristic is related to the structural configuration of the network, it is included in the VI because it directly affects the consequences of a successful attack. A visualisation of the *N-1* principle is presented in Figure 4.1. The remaining connection is assumed to have sufficient capacity to carry the relevant power flow between station *i* and station *j* after the other connection has failed.

The *physical separation of redundant assets* requires additional clarification. Redundant switchgear or connections that are located in the same room, cable duct, or physical corridor may be exposed to a common-cause attack. In that situation, one physical attack may disable both the primary asset and the asset intended to provide redundancy. Physically separated redundancy therefore reduces the possibility that a single physical attack affects both redundant installations. In this study, physical separation is classified as both an Exposure-related and an Impact-related characteristic: it affects the physical susceptibility to a common-cause attack and the consequences of a successful attack. The assessment does not model the specific physical mechanism through which an attack, such as fire or an impact from a drone, would damage multiple circuits on the same support structure. The possibility of such common-cause damage is therefore represented through the physical-separation characteristic and remains a simplification of the actual attack process (Koc et al., 2015).

The characteristic *potential unserved energy* represents the magnitude of the demand that could be affected if an asset is successfully attacked. It is operationalised using the annual energy demand of the relevant station or area. This measure is therefore used as a proxy for unserved energy rather than as a direct calculation of energy not supplied after an outage. A direct calculation of unserved energy would require both the amount of power that cannot be delivered and the duration for which the asset remains out of service. Restoration time is not included in this characteristic because reliable asset-specific restoration times are not available within the scope of this study. The selected measure should therefore be interpreted as an impact-magnitude proxy: it indicates how much energy demand is associated with an asset, but not how long that demand would remain unserved. This distinction also prevents restoration time from being counted twice. If restoration time were included as a separate characteristic or in a future extension of the index, the resulting measure could be interpreted more directly as expected energy not supplied. As it is also referred to in many other literature reviews as power loss (Holmgren, 2007a). In the present study, restoration time is outside the operationalisation of this characteristic.

The characteristic describing *Percentage co-location with other infrastructure* apply specifically to co-location with other infrastructure, such as waterways or bridges (Liu et al., 2025; Montalà-Palau et al., 2026). Underground connections that are not conventionally buried, but instead run along or attached to critical infrastructure, for example mounted under a bridge, or laid across a waterbed using a sinker rather than being buried in the ground. Such placements leave the connection physically accessible from the surface or waterline. This type of exposure has real-world precedent: anchor dragging has damaged underwater cables elsewhere (Lange, 2025), and bridge-mounted lines have been deliberately set on fire. Because both scenarios have already occurred, this characteristic is included in the VI.

The *coupling to the upstream transmission network* is excluded, as indicated in Table 4.1, because its effect is already captured in the potential unserved energy per station. *Telecommunication disruption* is likewise excluded, as this falls outside the scope of the examined network.

4.2.1. Weighting of characteristics

Some characteristics were considered more important by Stedin and were therefore assigned a higher initial weighting factor. The Exposure and Impact indicators listed below receive an initial weight of 2, whereas the remaining characteristics receive a weight of 1. These weights represent expert-informed modelling assumptions based on Stedin's risk priorities. They were specified independently of any statistical analysis of the case-study data.

The following characteristics receive an initial weight of 2 as *Exposure* indicators because Stedin considered them particularly relevant to the accessibility, reachability, or susceptibility of an asset. They are aligned, once again, with the direction in which it would influence the Vulnerability Index:

- + Physical exposure of transformers;
- Physical security monitoring;
- + Overhead line exposure;
- + Percentage crossing critical infrastructure.

The following characteristics receive an initial weight of 2 as *Impact* indicators because Stedin considered them particularly relevant to the consequences of a successful attack:

- + *N-1* redundancy exceedance;
- + Physically separated redundant switchgear;
- + Distance between redundant connections.

Subquestion 3 addresses a related but distinct question. It evaluates the extent to which the selected characteristics contribute to the variation in VI scores across the case-study assets. Because the VI is constructed from these characteristics, a characteristic with a higher initial weight contributes more to the score by definition. However, the weight alone does not determine the observed contribution of a characteristic. The contribution also depends on how much the characteristic varies across assets and on its relationship with the other characteristics included in the index.

A characteristic with a high weight but little variation across the assets may therefore explain relatively little of the variation in VI scores. Conversely, a characteristic with a lower weight may have a substantial contribution if its values differ strongly between assets. The results of SQ3 should therefore be interpreted as an empirical assessment of how the Stedin-informed weighting choices interact with the characteristics observed in the case-study network, rather than as an independent re-derivation of the initial weights.

4.2.2. Operationalisation to vulnerability assessment framework and scores

Now that all characteristics have been identified and quantified, they must be operationalised into usable scores for the vulnerability framework. Since each characteristic is expressed in a different unit, every score is normalised to a value between 0 and 1.

For characteristics describing *Percentage co-location with other infrastructure*, such as waterways or bridges, the underlying length is first converted into a percentage of the total connection length, after which this percentage is normalised into a score. A larger proportion of a connection's length exposed in this way results in a higher exposure score.

For a given exposure characteristic c , here, crossing critical infrastructure, let q index the individual crossing segments for edge e , for example separate bridge or waterway crossings along the same connection. The score for characteristic c , c being water-crossings or bridge-crossings, is then the combined length of these segments as a proportion of the total edge length:

$$E_{e,c} = \frac{\sum_q l_{e,q}}{l_e} \quad (4.1)$$

where $l_{e,q}$ denotes the length of segment q of the edge that coincides with the relevant infrastructure (e.g., a bridge or waterway crossing), and l_e denotes the total length of edge e (Montalà-Palau et al., 2026).

All characteristics, together with the scoring methods by which they were collected and their underlying data sources, are presented in Appendix C.2. For several characteristics, it was possible to find a complementary source besides internal data of Stedin for extra validation.

As discussed in Section 4.1, Holmgren (2007a) and Montalà-Palau et al. (2026) have proposed risk-index formulations that combine probability and consequence. Since such risk indices typically also incorporate hazard probability, this research instead formulates a Vulnerability Index (VI) per asset, treating hazard as a constant and thereby focusing exclusively on the asset-dependent components of vulnerability: *Exposure to threats* and *Impact of outage*. The VI is formulated for both edges (e) and nodes (v) as follows:

$$VI_e = E_e \cdot I_e \quad \forall e \in E \quad (4.2)$$

$$VI_v = E_v \cdot I_v \quad \forall v \in V \quad (4.3)$$

where stations are modelled as nodes (v) and connections as edges (e), as introduced before. Since the exposure and impact characteristics were identified separately in Table C.2, each component is calculated as follows.

The exposure score is calculated as a weighted sum of the normalised exposure characteristics. Since some characteristics are considered more important contributors to vulnerability than others (Section 4.2), weighting factors w_c are assigned to selected indicators: the particularly important contributors get a w_c of 2, and all other exposure characteristics get a w_c of 1, see 4.2.1. The exposure score is therefore calculated as:

$$E_e = \frac{\sum_c w_c E_{e,c}}{\sum_c w_c} \quad \forall e \in E \quad (4.4)$$

$$E_v = \frac{\sum_c w_c E_{v,c}}{\sum_c w_c} \quad \forall v \in V \quad (4.5)$$

where w_c denotes the weighting factor assigned to characteristic c , and $E_{e,c}$ and $E_{v,c}$ denote the normalised score (0–1) of exposure characteristic c for edge e and node v , respectively. To keep the exposure score on a normalised scale between 0 and 1, the weighted sum is divided by the sum of all applied weighting factors.

For the impact score, the operationalisation is a little more complex. The annual energy supplied (GWh) approximates the unserved energy resulting from an asset's failure and therefore represents the *Impact* component of the VI. Unserved energy is conventionally calculated as the product of power loss (MW) and recovery time (h) (Appendix B). However, recovery time data for malicious disruptions is unavailable, due to the infrequent and unpredictable nature of such events. The annual energy supplied in the reference year 2030 is used directly for the Impact component instead, since load prognoses were available for that year. For nodes, L_v equals the annual energy supplied through the corresponding station. For edges, L_e represents the annual energy demand of the downstream supply area that would be affected by the failure of that connection.

As well as the annual energy supplied, the $N-1$ redundancy exceedance ($1 - R_{N-1,area}$) is also determined per asset, per 150 kV forecast area. This follows Stedin's internal methodology, in which the forecast loads of all substations within a forecast area are aggregated for every hour and compared with the available $N-1$ capacity of the supplying 150 kV station. Connections at the 380 kV level are excluded from this correction, since the analysis considers 150 kV forecast areas specifically.

The share of hours in which this capacity is exceeded is defined through an indicator function. For an aggregated hourly load P and an available $N-1$ capacity C :

$$\mathbb{I}(P, C) = \begin{cases} 1, & \text{if } P > C, \\ 0, & \text{if } P \leq C. \end{cases} \quad (4.6)$$

For each forecast area and hour t , the two inputs of the indicator function are defined as the aggregated forecast load $\sum_{v \in V_{area}} P_{v,t}$ and the available $N-1$ capacity C_{area}^{N-1} , respectively. The annual $N-1$ redundancy exceedance is therefore calculated as:

$$1 - R_{N-1,area} = \frac{1}{8760} \sum_{t=1}^{8760} \mathbb{I} \left(\sum_{v \in V_{area}} P_{v,t}, C_{area}^{N-1} \right) \quad \forall areas. \quad (4.7)$$

Here, V_{area} is the set of substations within forecast area, t indexes the 8760 hours of a year, $P_{v,t}$ is the forecast load of substation v , during hour t , and C_{area}^{N-1} is the available $N-1$ capacity of the supplying 150 kV station. For every hour, the indicator function returns 1 if the aggregated load of the area exceeds this capacity, and 0 otherwise. Summing the indicator function over all 8760 hours and dividing by 8760 gives the share of annual hours with an $N-1$ capacity exceedance. The resulting redundancy exceedance ($1 - R_{N-1,area}$) can subsequently be assigned to all assets belonging to that forecast area.

This area-based approach is required because the network is physically not entirely radial. Several forecast areas are interconnected through redundant connections between substations. Under normal

operating conditions, these redundant connections remain energized but carry no power flow; they are only intended to transfer load following the failure of another connection.

For edges, the appropriate redundancy exceedance cannot be assigned directly from the original forecast area. Instead, each edge failure is analysed individually. The failure of an edge may divide the network into a newly formed supply area, in which one or more substations become supplied through a neighbouring forecast area via a redundant connection. For every edge, this newly created area is identified manually, after which the total forecast load of all substations within that area is compared with the available $N-1$ capacity of the supplying 150 kV station, recalculated using Equation (4.7). The resulting $N-1$ redundancy exceedance of this newly formed area is subsequently used as the correction factor in the Impact calculation for that specific edge.

Figure 4.2 shows an example of a newly formed forecast area, with the feed areas indicated by dotted lines. If the red line fails in the lower part of the figure, the redundant connection is used instead, and the affected node is supplied by a different 150 kV feeder.




Figure 4.2: Shift in feed area following the redundancy recalculation. The failure of the line shown in red causes a subtransmission station to be fed through a different upstream 150 kV node.

Whether this redundancy correction is applied depends on a second condition: whether the redundant connection is physically separated from the connection it backs up. In practice, co-located redundant connections in the Stedin network typically run through the same cable duct, with a physical separation of approximately 50–100 cm. Connections within this range are therefore considered to share a common cause failure risk and do not provide genuine physical redundancy. See a further explanation in Appendix B. For edges, this condition is derived from the characteristic *distance between redundant connections* (d_e , in metres), using a threshold of 1 m, consistent with this practical co-location range.

This threshold is applied directly in the Impact calculation:

$$I_e = \begin{cases} L_e \cdot (1 - R_{N-1, \text{area}}), & \text{if } d_e \geq 1, \\ L_e, & \text{if } d_e < 1, \end{cases} \quad \forall e \in E \quad (4.8)$$

For nodes, the equivalent condition is derived from the characteristic *redundant switchgear physically separated* (s_v), which is recorded on Stedin's internal 1–3 categorical scale (1 = two installations in two physically separated rooms, 2 = two installations in one shared room, 3 = only a single installation present). Categories 2 and 3 both indicate the absence of genuine physical redundancy, either due to co-location within the same room or the absence of a redundant installation altogether, and are therefore both treated as *not physically separated*. This condition is likewise applied directly in the Impact calculation:

$$I_v = \begin{cases} L_v \cdot (1 - R_{N-1, \text{area}}), & \text{if } s_v = 1, \\ L_v, & \text{if } s_v \in \{2, 3\}, \end{cases} \quad \forall v \in V \quad (4.9)$$

The impact value is taken from the system within the substation that is not physically separated. This is determined by using the annual energy supplied (*GWh*) of the corresponding system within the substation (e.g., 13, 23, 50 or 150 kV).

5

Graph-theoretic assessment of structural criticality

Section 4.2 identified the asset-specific characteristics included in the Vulnerability Index (VI). While the VI assesses the susceptibility and potential consequences associated with individual assets under the fixed physical sabotage threat, this chapter addresses the complementary network-level perspective: which assets are most consequential when removed from the network. These assets are referred to as *structurally critical assets*. Structural criticality therefore does not describe the probability that an asset will fail or be attacked, but the network-level consequence conditional on its loss.

To assess structural criticality, the case-study network is first represented as a graph and local node weights are assigned (5.1). Subsequently, centrality-based decision rules and Weighted Connectivity Loss (WCL) are introduced as the main components of the assessment (5.2). Structural consequences are then evaluated through single-asset contingency analysis and sequential targeted attack simulations, representing different asset-removal and attacker-information assumptions. Finally, the resulting removal sequences are evaluated using a global robustness consistency check (5.5).

5.1. Case-study network representation

The case-study network is represented as a multigraph

$$F = (V, E),$$

where V denotes the set of nodes and E denotes the set of edges. Nodes represent substations or stations, while edges represent electricity connections between these stations. Parallel edges are retained explicitly because they represent physically distinct circuits that may provide redundancy between the same pair of nodes. Retaining parallel connections is important when assessing structural vulnerability, because the removal of one circuit does not necessarily disconnect the areas connected by the remaining circuits (Forsberg et al., 2023; Pagani & Aiello, 2015).

The set of assets considered in the structural assessment is defined as

$$A = V \cup E.$$

This allows both node failures and edge failures to be assessed within the same framework. The graph also contains a subset $S \subseteq V$ of source nodes. A node is considered reachable when a path exists between that node and at least one source node in S .

5.1.1. Local node weights

The Weighted Connectivity Loss metric requires a weight for each node that represents the annual energy locally associated with that node. The available station values are cumulative: the measured value of an upstream station includes the energy associated with downstream stations. Using these

cumulative values directly in a network-wide sum would therefore count downstream energy multiple times.

To avoid this double counting, cumulative station values are converted into local node weights. For node i , the local weight is calculated as

$$w_i \text{ [GWh]} = L_{\text{station},i} - \sum_{j \in \text{Ch}(i)} L_{\text{child},j}, \quad (5.1)$$

where $L_{\text{station},i}$ denotes the cumulative annual energy value associated with station node i , and $\text{Ch}(i)$ denotes the set of direct downstream child nodes of i . The term $L_{\text{child},j}$ represents the cumulative annual energy value of child node j .

The resulting weight w_i represents the local annual energy demand associated with node i , rather than the total cumulative energy supplied downstream of that node. These local weights are used in WCL as proxies for the potential annual energy that could become disconnected after an asset removal. They differ from the cumulative annual energy values used as an asset-specific impact proxy in the VI.

5.2. Selection of assessment methods

The selected assessment methods reflect the topology of the case-study network and the information available to different types of attackers. Centrality-based attack strategies use information about the network structure to rank assets and determine the order in which they are removed (Cetinay et al., 2018). In particular, centrality-based attack profiles have been used to model sequential attacks in which the most central remaining component is removed at each iteration (Albarakati & Bikdash, 2022).

Degree and origin-destination betweenness centrality were selected over Effective Graph Resistance (EGR) and power-flow-based methods. These electrically informed metrics derive their added value from meshed networks, where multiple alternative paths exist between a node pair and differ in electrical distance. In the case-study network, which is radial, only one electrical path exists between any pair of nodes i and j . As a result, origin-destination betweenness centrality already identifies the same critical paths that an electrical distance measure would identify, without requiring impedance data. EGR and power-flow-based methods were therefore not selected.

The present framework distinguishes between centrality-based decision rules and impact-based decision rules. Centrality-based decision rules use the current graph structure to select the next asset, without evaluating the impact of removal. In this study, they therefore represent a topology-informed attacker without asset-specific information about the resulting network impact, also known as the *uninformed attacker*. The electrical weight w_i per station, defined in the previous section, is also used as an a-priori decision rule, explained further in the attack strategies below. Impact-based decision rules temporarily remove every candidate asset and select the candidate that produces the largest deterioration in network performance. They therefore represent an *impact-informed attacker* with information about the consequences of asset removal. This distinction is central to the targeted attack simulations in this chapter, because it determines how the next target is selected and what the results reveal about the value of attacker information.

Two removal-based assessment procedures are used, both a-posteriori. First, single-asset contingency analysis evaluates the consequence of removing each node or edge independently. Second, sequential targeted attack simulations evaluate how network performance deteriorates when assets are removed according to a specified decision rule. These procedures answer different questions: contingency analysis evaluates individual asset outages, whereas targeted attack simulations assess the deterioration caused by a sequence of deliberate removals.

5.2.1. Centrality-based decision rules

Centrality measures are used as decision rules in the centrality-based targeted attack simulations. They do not directly measure the probability that an asset will fail. Instead, they rank assets according to their position in the graph and determine which asset is selected at each iteration of the attack.

Centrality measures are therefore used as a-priori indicators of structural importance, rather than as direct measures of the consequence of asset removal (Cetinay et al., 2018; Pagani & Aiello, 2015).

Degree centrality

Degree centrality measures the number of direct connections incident to a node. Let M denote the adjacency matrix of the multigraph, where m_{ij} represents the number of edges between nodes i and j . The degree of node i is defined as

$$k_i = \sum_{j \in V} m_{ij}. \quad (5.2)$$

In the multigraph representation, parallel circuits are counted separately. Consequently, degree centrality captures the local number of direct connections, including the presence of parallel or redundant circuits. It does not directly indicate how many downstream demand nodes depend on a node. Degree centrality is therefore interpreted as a local connectivity indicator. This makes it a complementary measure to betweenness centrality, particularly in networks containing locally well-connected hubs or redundant circuits (Forsberg et al., 2023; Pagani & Aiello, 2015).

Origin-destination betweenness centrality

Standard betweenness centrality considers shortest paths between all pairs of nodes. This is not fully aligned with the purpose of the present analysis, because electricity is supplied from source nodes to demand nodes. The relevant paths are therefore the paths between an origin in the source set S and a destination in the set of load-bearing nodes D .

The origin-destination betweenness of node i is defined as

$$B_i^{OD} = \sum_{s \in S} \sum_{t \in D} \frac{\sigma_{s \rightarrow t}(i)}{\sigma_{s \rightarrow t}}, \quad (5.3)$$

where $\sigma_{s \rightarrow t}$ is the number of shortest paths from source node s to demand node t , and $\sigma_{s \rightarrow t}(i)$ is the number of these paths that pass through node i .

The equivalent edge-based measure is

$$B_e^{OD} = \sum_{s \in S} \sum_{t \in D} \frac{\sigma_{s \rightarrow t}(e)}{\sigma_{s \rightarrow t}}, \quad (5.4)$$

where $\sigma_{s \rightarrow t}(e)$ denotes the number of source-to-demand shortest paths that contain edge e .

The restriction to source-to-demand paths reflects the directional purpose of the network: electricity is supplied from source nodes to load-bearing nodes. The measure therefore differs from standard all-pairs betweenness, which considers shortest paths between all node pairs irrespective of their role as source or demand nodes. The source-restricted formulation is more closely aligned with the structural dependency that is relevant in a distribution network.

In a strictly radial network represented as a tree, each source-to-demand pair is connected by a unique path. For an individual source-to-demand pair, $\sigma_{s \rightarrow t} = 1$ and $\sigma_{s \rightarrow t}(i) \in \{0, 1\}$, so the ratio becomes either zero or one. A node may nevertheless lie on the paths of many different source-to-demand pairs, which is reflected in the sum over all pairs.

The case-study network is represented as a multigraph and contains redundant connections. Consequently, multiple shortest paths may exist between some source and demand nodes. In these cases, the ratio represents the fraction of shortest source-to-demand paths that pass through the asset rather than a binary value. The resulting score therefore reflects the extent to which an asset participates in source-to-demand connectivity while also accounting for alternative shortest paths.

Origin-destination betweenness is therefore more directly aligned with the objective of the present study than standard all-pairs betweenness. In a predominantly radial network, it captures the dependence

of downstream demand nodes on the paths connecting them to their sources. Degree centrality and origin-destination betweenness consequently capture different aspects of structural importance: degree centrality captures local branching and the number of direct circuits connected to a node, whereas origin-destination betweenness captures the number of source-to-demand routes that depend on a component. These two measures are used as alternative centrality-based decision rules in the IMCE simulations.

5.2.2. Weighted Connectivity Loss

Centrality measures provide a decision rule for selecting assets, but they do not directly quantify the consequence of removing an asset. The consequence of a removal is therefore measured using Weighted Connectivity Loss (WCL). The measure builds on connectivity-loss approaches that assess the share of network demand that becomes disconnected after an asset failure, while extending this principle with locally weighted node demand (Albarakati & Bikdash, 2022; Galvan & Agarwal, 2020).

Let Γ denote a network state and let $R(\Gamma) \subseteq V$ denote the nodes that remain reachable from at least one source node after the relevant asset removals. The WCL of network state Γ is defined as

$$\text{WCL}(\Gamma) = 100\% \cdot \left(1 - \frac{\sum_{i \in R(\Gamma)} w_i}{\sum_{i \in V} w_i} \right). \quad (5.5)$$

The resulting value represents the percentage of locally associated node weight that is no longer reachable from any source. Because the node weights represent potential annual energy proxies, WCL can also be interpreted as the percentage of potential energy demand associated with nodes that becomes disconnected. WCL is therefore a load-weighted connectivity measure. It is not a power-flow calculation and does not represent directly simulated unserved load or unserved energy. Instead, it provides a topological consequence measure weighted by the energy associated with the affected nodes.

A value of $\text{WCL} = 0$ indicates that all weighted nodes remain reachable from at least one source. A value greater than zero indicates that one or more load-bearing nodes have lost source reachability. In this study, such an outcome is referred to as *islanding*.

The node weights used in WCL differ from the cumulative annual energy values used in the VI. The VI uses cumulative annual energy as an asset-specific potential unserved energy proxy, whereas WCL requires local weights so that energy is not counted multiple times when the reachable nodes are summed.

5.3. Single-asset contingency analysis

Contingency analysis evaluates the consequence of individual asset outages. In accordance with the $N - 1$ criterion, each node and edge is removed once while all other assets remain operational. The resulting network state is then evaluated using WCL and an islanding indicator. This type of graph-theoretical contingency analysis provides a computationally efficient way to compare the structural consequences of individual asset outages, particularly when detailed electrical information is unavailable (Coelho et al., 2019; Islam Khan et al., 2026).

For a node contingency, node $v \in V$ and its incident edges are removed. The reachable node set is determined as

$$R(F - v),$$

after which the corresponding WCL is calculated. For an edge contingency, each edge $e \in E$ is removed independently and the reachable node set

$$R(F - e)$$

is determined.

In addition to WCL, the analysis records whether the outage causes source-islanding with impact. This occurs when one or more load-bearing nodes are no longer reachable from any source, resulting in

$WCL > 0$. The number of disconnected load-bearing nodes is reported as a complementary structural indicator. Reporting both the weighted connectivity loss and the resulting structural fragmentation distinguishes the amount of disconnected demand from the number of disconnected nodes.

The contingency analysis distinguishes between outages that preserve source reachability and outages that disconnect part of the network. In particular, parallel edges may preserve connectivity after the removal of one circuit. Such an edge is redundant with respect to the considered source-reachability criterion if its removal results in $WCL = 0$.

5.4. Sequential targeted attack simulations

Contingency analysis evaluates each asset independently. It does not model how a deliberate attacker may adapt the next target after observing the effect of previous removals. Sequential targeted attack simulations address this issue by removing assets one after another and recalculating the decision rule after every removal. Sequential centrality-based and impact-based attack profiles have been used in previous power-grid research to compare the effect of different target-selection strategies (Albarakati & Bikdash, 2022).

Let $F = (V, E)$ denote the original network, where V is the set of nodes and E is the set of edges. The complete set of network assets is defined as $A = V \cup E$.

Because node and edge attacks are analysed separately, the set of candidate assets in a given simulation is defined as

$$A_{\text{att}} = \begin{cases} V, & \text{for node attacks,} \\ E, & \text{for edge attacks.} \end{cases} \quad (5.6)$$

Let Γ_r denote the current network state after r asset removals, with $\Gamma_0 = F$ representing the original network. The attack sequence is indexed by $r \in \{0, \dots, |A_{\text{att}}| - 1\}$.

The set of candidate assets available at iteration r is

$$C_r = A_{\text{att}} \setminus \{a_0, \dots, a_{r-1}\}, \quad C_0 = A_{\text{att}}. \quad (5.7)$$

Thus, C_r contains all assets that have not yet been removed. At each iteration, one candidate asset $a_r \in C_r$ is selected. The next network state is obtained by removing this asset:

$$\Gamma_{r+1} = \text{Remove}(\Gamma_r, a_r). \quad (5.8)$$

Here, $\text{Remove}(\Gamma_r, a_r)$ denotes the network state after removing asset a_r . If a_r is a node, its incident edges are removed as part of the same operation. If a_r is an edge, only that edge is removed.

Two sequential attack strategies are considered. The Iterated Most Central Element (IMCE) represents an attacker who uses observable network topology but has no asset-specific information about the resulting impact. The Iterated Electrical Most Damaging Element (IEMDE) represents an impact-informed attacker who selects the next asset based on the network deterioration caused by its removal. The two strategies therefore correspond to different assumptions about the information available to the attacker (Albarakati & Bikdash, 2022; Cetinay et al., 2018).

The distinction between the two strategies is not that one is structural and the other is non-structural. Both strategies remove assets from the graph, and the resulting network states can be evaluated using network-level consequence measures. They differ in how the next target is selected: IMCE uses a graph metric as its decision rule, whereas IEMDE uses the impact of temporarily removing each candidate asset.

5.4.1. Iterated Most Central Element

IMCE represents the topology-informed, impact-uninformed attacker. A centrality metric ϕ is used as the decision rule. At each iteration, the centrality value of every remaining candidate asset is calculated on the current network state. The asset with the highest value is then selected for removal:

$$a_r = \arg \max_{x \in C_r} \phi(x \mid \Gamma_r). \quad (5.9)$$

Here, r denotes the attack iteration, a_r denotes the asset selected for removal at that iteration, and x denotes a candidate asset being evaluated during the maximisation. The centrality score $\phi(x \mid \Gamma_r)$ is therefore calculated for each candidate $x \in C_r$, while only the candidate with the highest score is selected as a_r .

For node attacks, IMCE is evaluated using degree centrality and origin-destination betweenness centrality. These decision rules represent different types of topological information available to the attacker. Degree centrality uses local connection information, whereas origin-destination betweenness measures the dependence of source-to-demand paths on a node. For edge attacks, the equivalent edge-based betweenness measure is used.

After each removal, the relevant centrality values are recalculated because the network structure and the remaining source-to-demand paths may have changed. The network state is then updated according to Equation 5.8. This recalculation means that IMCE represents an iterated centrality strategy rather than a one-time ranking based only on the original network.

5.4.2. Iterated Electrical Most Damaging Element

IEMDE represents the impact-informed attacker. Unlike IMCE, IEMDE does not use a predefined centrality measure as its decision rule. At each iteration, every remaining candidate asset is temporarily removed from the current network state. The resulting WCL is calculated for each candidate, and the candidate producing the largest WCL is selected for permanent removal:

$$a_r = \arg \max_{x \in C_r} \text{WCL}(\text{Remove}(\Gamma_r, x)). \quad (5.10)$$

The network state is then updated by removing the selected asset according to Equation 5.8. The marginal deterioration caused by this removal is defined as

$$\tilde{z}_r = \text{WCL}(\Gamma_{r+1}) - \text{WCL}(\Gamma_r). \quad (5.11)$$

The original IEMDE formulation selects the Electrical Most Damaging Element according to the increase in unsatisfied energy (Albarakati & Bikdash, 2022). In the present study, this consequence criterion is replaced by WCL. Because WCL is calculated from local node weights representing potential annual energy proxies, the selected asset is the candidate whose removal causes the largest immediate increase in load-weighted connectivity loss. WCL therefore provides a structural proxy for the potential energy demand that becomes disconnected, rather than a direct calculation of unsatisfied load or unserved energy.

IEMDE requires more information about the network response than IMCE, because every remaining candidate asset must be evaluated before the next target is selected. It therefore represents an impact-based decision rule under the WCL criterion, whereas IMCE represents a centrality-based decision rule. The comparison between the two strategies consequently indicates how the assumed level of attacker information affects the selected attack sequence.

5.5. Global robustness consistency check

Effective Graph Resistance (EGR) is used as a complementary global network measure to examine whether the degradation identified by WCL is also reflected by an independent structural measure

(Ellens et al., 2011). EGR is not used as a decision rule for selecting attack targets and does not replace WCL in the contingency or IEMDE analyses.

For this consistency check, the IEMDE node-removal sequence is reused. After each removal, shortest-path distances between all reachable real-node pairs within the same connected component are calculated using physical cable length as edge weight. These pairwise distances are summed and normalised by the value of the intact network, resulting in a trajectory that starts at one.

Because physical distance is used instead of electrical impedance and effective resistance, this measure is interpreted as a distance-based EGR proxy rather than conventional electrical EGR. Conventional EGR accounts for electrical properties and alternative paths through effective resistance (Ellens et al., 2011; Koç et al., 2014). The approximation is considered suitable here as a structural consistency check because the case-study network is predominantly radial and source-to-node connections mainly follow unique paths.

The normalised EGR proxy is compared with the remaining weighted connectivity, $1 - \text{WCL}(\Gamma_r)/100$, along the same IEMDE removal sequence. This comparison is relevant because IEMDE selects, at each step, the asset that maximises WCL. A strong decline in remaining weighted connectivity is therefore inherent to the target-selection procedure itself. The EGR proxy is not used for target selection and is calculated afterwards over the same removal sequence. A comparable degradation pattern therefore provides a consistency check of whether the network degradation identified through WCL is also reflected by a different structural robustness measure.

Figure 5.1 summarises the methodological criticality framework applied in this study. The VI is included for completeness but is not used as a decision rule within the structural criticality assessment. Conceptually, however, asset-specific vulnerability information could form an alternative target-selection criterion under a different attacker-information assumption.




Figure 5.1: Overall methodological framework for the vulnerability and structural criticality assessments.

6

Results

This chapter presents the results of applying collected data on the frameworks as described in Chapter 3. Due to confidentiality reasons, the input data is not shared in this research. However, how data is applied, is explained in Chapter 4. Paragraph 6.1 presents the resulting Vulnerability Index (VI) scores per asset (**SQ2**). Paragraph 6.2 presents the structural vulnerability scores resulting from the graph-theoretic metrics applied to this network (**SQ4**). The last paragraph reveals some overlaps between the two (**SQ5**).

6.1. Vulnerability Index (VI) scores

The Vulnerability Index (VI) for each asset was calculated as

$$VI = E \cdot I, \quad (6.1)$$

following the framework presented in Section 4.2.2. This section presents the resulting VI scores for both nodes and edges within the case study network.

The multiplicative form reflects the conceptualisation of vulnerability as the combination of exposure and impact. An asset is considered highly vulnerable when it is both susceptible to the relevant threat and associated with substantial consequences if successfully disrupted. Multiplication therefore ensures that a low value in either component limits the resulting VI score. For example, an asset with high exposure but negligible impact does not receive a high vulnerability score solely because it is likely to be affected. An additive formulation would not impose this joint requirement, because a high value in one component could compensate for a very low value in the other.

Because the exposure and impact scores are normalised, the multiplicative form also produces a dimensionless score on a bounded relative scale. The VI should therefore be interpreted as a composite prioritisation score rather than as a probability or a directly measured quantity of physical damage. A high vulnerability score indicates that an asset is both highly exposed to potential failure and associated with substantial consequences when disrupted. The underlying exposure and impact indicators provide insight into the factors driving the final vulnerability score.

Table 6.1 presents the ten highest-ranked nodes together with their Exposure, Impact and resulting VI scores. The complete overview of all 26 node scores is provided in Appendix D.

Table 6.1: Top 10 nodes ranked by Vulnerability Index.

Rank	Node ID	Exposure (E_v)	Impact (I_v)	VI_v
1	5	0.238	0.544	0.130
2	1	0.190	0.431	0.082
3	4	0.286	0.272	0.078
4	12	0.198	0.198	0.039
5	9	0.198	0.192	0.038
6	23	0.238	0.115	0.027
7	3	0.143	0.079	0.011
8	6	0.071	0.143	0.010
9	7	0.071	0.123	0.009
10	25	0.143	0.056	0.008

For interpretation, for example, Node 5 achieves the highest vulnerability score ($VI_v = 0.130$), whereas Node 9 obtains a score of only $VI_v = 0.038$. Consequently, Node 5 can be considered approximately 4.4 times more vulnerable than Node 9 according to the adopted vulnerability assessment framework:

$$\frac{VI_5}{VI_9} = \frac{0.130}{0.038} \approx 3.4 \quad (6.2)$$

This comparison does not imply that Node 5 is 3.4 times more likely to fail. It indicates that Node 5 receives a substantially higher relative vulnerability score under the adopted assessment framework. The VI is therefore primarily intended as a relative measure for comparing and prioritising assets.

Similarly, Table 6.2 presents the ten highest-ranked edge groups. Since edges occur in bidirectional pairs with identical scores, paired edges are reported together. The complete overview of all edge scores is provided in Appendix D.2. For some edge pairs, the scores differ slightly because parallel redundant cables may follow slightly different physical routes, resulting in minor differences in their physical characteristics.

Table 6.2: Top 10 edge groups ranked by Vulnerability Index.

Rank	Edge IDs	Exposure (E_e)	Impact (I_e)	VI_e
1	29/30	0.600	1.000	0.600
2	17/18	0.517–0.519	0.821	0.424–0.426
3	25/26	0.200	0.638	0.128
4	9/10	0.491	0.253	0.124
5	7/8	0.457–0.469	0.253	0.116–0.119
6	3/4	0.200	0.265	0.053
7	34/36	0.200	0.233	0.047
8	42/43	0.200	0.233	0.047
9	44/45	0.232	0.160	0.037
10	38/39	0.200	0.141	0.028




Figure 6.1: Relative VI decrease within nodes and edges

Figure 6.1 shows how the VI score decreases across the ranked assets, for both nodes and edges. This illustrates two things: how quickly the VI score drops off within the ranks, and what proportion of all assets actually receives a non-zero VI score. The relative node and edges rankings have a negative exponential distribution, where the distribution of the edges has a bigger drop at edge 3.

As discussed in Section 4.2.2 and Figure 4.2, the *N-1* redundancy exceedance used in the Impact component is determined at the level of the 150 kV forecast areas. Analysis of the load profiles associated with the different forecast areas showed that two areas do not fully satisfy the *N-1* criterion under all operating conditions. The remaining forecast areas satisfy the *N-1* criterion, resulting in a redundancy exceedance of zero. Consequently, no additional Impact condition was assigned to assets located within those areas.

6.1.1. Sensitivity Analysis

The asset-specific vulnerability framework is validated in two complementary ways: the resulting ranking and its underlying characteristics were assessed by grid experts, while the robustness of the framework is evaluated through sensitivity analysis. To this end, the characteristics contributing most to the VI are first identified through a correlation analysis and subsequently tested for robustness through a sensitivity analysis **SQ3**. In accordance with verification and validation principles for analytical models, this combination evaluates both the sensitivity of the framework to variations in its input characteristics and the plausibility of its outcomes (Sargent, 2020). Sensitivity analysis provides insight into the robustness of the resulting rankings, while expert validation assesses whether the identified vulnerable assets are considered reasonable from an operational perspective.

In this section's notation, E denotes an exposure characteristic and I an impact characteristic. The numerical subscript indicates the characteristic number within the respective category, while v and e denote vertex (node) and edge characteristics, respectively. For example, $E_{1,v}$ refers to the first exposure characteristic of a vertex, whereas $I_{2,e}$ refers to the second impact characteristic of an edge.

Correlation characteristics

Before selecting a correlation coefficient, each characteristic was categorised by measurement level (nominal, ordinal, ratio, or binary), since the appropriate test depends on this classification. Pearson correlation was used for ratio-level characteristics assumed to have a linear relationship with the VI. Spearman correlation was used for ordinal characteristics, because it assesses whether the variables have a monotonic relationship without requiring normally distributed values. Welch's t -test was used for the binary characteristic, $E_{2,v}$, because the VI scores of the two categories could be compared while allowing for potentially unequal variances between the groups. The full results for node and edge characteristics are reported in Appendix D.3 and Appendix D.4, respectively. Only characteristics showing a statistically significant relationship with the VI were carried forward into the sensitivity analysis (Scribbr, n.d.).

The characteristics showing a statistically significant relationship with the VI were $E_{1,v}$, $E_{3,v}$, $I_{1,v}$, $E_{1,e}$ and $E_{2,e}$. These characteristics include both exposure and impact indicators. Their statistical association with the VI indicates that they contribute to the variation in vulnerability scores within the case-study network.

Robustness outcomes

The characteristics selected in Section 6.1.1 were varied one at a time: by 10%, 20% and 30% for continuous characteristics, and by one category level up or down for the ordinal exposure characteristic $E_{1,v}$. After each variation, the VI ranking was recomputed and compared to the original ranking using two indicators: the overlap between the original and the new top five assets, and whether the internal order of the top five changed. The full results are reported in Appendix D.5.

For the first node impact factor, $I_{1,v}$, the top five assets and their order remained fully unchanged across the entire -10% to +30% range. The third node exposure factor, $E_{3,v}$, retained a full top-five overlap across all six scenarios, though the order of ranks four and five changed under all three negative variations. The first node exposure factor, $E_{1,v}$, showed the largest change in composition, with three of the original top five assets retained at one category level below the original value and two retained at one category level above the original value.

Among the edge characteristics, $E_{2,e}$ retained a full top-five overlap in both directions, with only the internal order changing under positive variation. The characteristic $E_{1,e}$ retained a full overlap under positive variation but dropped to two of five under negative variation, with three new edges entering the top five.

This pattern follows directly from the structure of the underlying impact calculation. Some of the characteristics identified as significant function as dual-role terms within the VI: besides contributing a continuous, normalised value, they also act as threshold conditions inside the impact term. For the ordinal characteristic, shifting the value by a single category level can be enough to cross this threshold and change which impact condition applies. For a ratio characteristic that showed reduced top-five overlap under negative variation, moving the value from just above to just below the 100% threshold has the same effect. In both cases, this produces a different ranking rather than a proportional shift in scores.

The remaining significant characteristics do not have the same dual-role structure, which is consistent with their producing only minor reordering within an otherwise stable top five.

Overall, the sensitivity analysis shows that the main vulnerability ranking is generally robust to moderate changes in the input characteristics. The largest changes occur for characteristics that also determine which impact condition applies through a threshold. These characteristics can change the composition of the top five when their values cross the relevant threshold, whereas characteristics without this threshold effect mainly influence the relative order of assets that already appear in the top five.

6.2. Structural criticality scores

This section presents the results obtained by applying the structural vulnerability framework (**SQ4**). Before presenting these findings, the case-study network was constructed in NetworkX. As discussed in Section 5.2, the analysis is based on the multigraph F . Compared with the representation introduced in Chapter 4, this graph additionally represents the physical separation of redundant circuits. This characteristic is included because it helps distinguish redundant circuits that follow physically separate routes from circuits that run alongside one another through the same cable duct.

Network F is presented below:




Figure 6.2: Network F , represented as a graph constructed in NetworkX.

The figure shows that several additional nodes have been added to the network. These nodes do not represent stations, but rather points where cabling splits or converges. They are therefore excluded from the subsequent analysis.

6.2.1. Targeted attack simulations

This section presents the results of the targeted attack simulations. Before these results are interpreted, the centrality measures that serve as input for the IMCE attack strategies are presented first.

Centrality results

Table 6.3 reports the five highest values for each centrality measure. For edge-based measures, the identifiers refer to the numbered edge records in the graph. A complete edge may therefore be rep-

resented by two edge identifiers. Additional visualisations of the spatial distribution of the centrality measures are available in the confidential version of this thesis.


Table 6.3: Top 5 centrality rankings for Network F .

Rank	Node	b_i	Node	k_i	Edge	$b_{i,j}$
1	18	0.0086	18	0.12	29/30	0.0078
2	11	0.0082	2	0.12	13/14a	0.0078
3	8	0.0082	11	0.10	13/14b	0.0078
4	12	0.0078	1	0.10	17/18b	0.0078
5	2	0.0065	12	0.08	17/18d	0.0078


These values serve as the input for the IMCE attack strategies. The full centrality rankings underlying this table are reported in Appendix E.1 and E.2.

Results of targeted attack scenarios

The figures below show the results of the **IMCE** (Iterated Most Central Element) and **IEMDE** (Iterated Electrical Most Damaging Element) attack strategies. IMCE represents a topology-informed attacker who selects assets using a centrality-based decision rule. IEMDE selects the candidate asset whose temporary removal produces the largest increase in WCL. Because WCL is weighted by local annual energy values, IEMDE represents an impact-based decision rule that prioritises the loss of weighted connectivity.



(a) Edges



(b) Nodes

Figure 6.3: Targeted attack strategies and remaining weighted connectivity.

As shown in the figures, IEMDE and the betweenness-based IMCE cause a comparably fast deterioration of weighted connectivity for node attacks, whereas for edge attacks the betweenness-based IMCE outperforms IEMDE. IEMDE selects the asset whose temporary removal produces the largest increase in WCL, using local annual energy values to weight the disconnected network consequence.

Among the node-based strategies, IEMDE and the betweenness-based IMCE reach complete loss

of weighted connectivity after the same number of removals, namely 7 out of 26 nodes (27%). The degree-based IMCE performs considerably worse, requiring 18 node removals (69%) to reach the same outcome.

The degree-based curve temporarily stabilises at approximately 2% remaining weighted connectivity around the 11th node removal, before reaching zero after 18 removals. This plateau indicates that the remaining upstream structure still retains weighted connectivity to the source, even though a substantial part of the downstream network has already become disconnected.

The full fragmentation and islanding results are reported in Appendix E. For edge attacks, the betweenness-based IMCE reaches its final plateau after 19 branch removals (41%), whereas the IEMDE curve reaches the same plateau only after 25 removals (54%).

As shown in Figure 6.3a, edge removal does not reduce the remaining weighted connectivity to 0%. Instead, the curves stabilise at approximately 39%. This remaining share is associated with the nodes at the top of the tree. It represents the local energy associated with these upstream nodes, excluding the local weights of their downstream children. Consequently, edge removal disconnects the downstream network while the upstream node weights remain reachable from the sources.

6.2.2. Contingency Analysis (N-1)

The contingency analysis differs from the targeted attacks because it does not consider iterative, sequential failures. Instead, each node or edge is removed independently while all other assets remain operational. This follows the classical *N-1 principle*: the network is assessed under the failure of one asset at a time, and the resulting connectivity loss is recorded.

The analysis was implemented in Python using NetworkX. For each asset, the nodes that become isolated from the sources are identified, and the corresponding lost weight is expressed as a WCL percentage of the total network weight.




Figure 6.4: Node contingency analysis on Network F , showing (a) WCL and (b) the number of isolated nodes.

Figure 6.4 shows that the node causing the largest WCL is not the same as the node isolating the largest number of nodes. The removal of Node 1 produces the largest WCL, whereas the removal of Node 2 isolates the largest number of nodes. This difference illustrates that the number of isolated nodes and the amount of disconnected weighted energy are distinct consequence measures.

The results therefore do not support that a topology-based target selection necessarily causes the same impact as an impact-based target selection. A target selected using only structural information may isolate many nodes while disconnecting less weighted energy than another asset that isolates fewer nodes. The contingency results should consequently be interpreted as separate evidence on node fragmentation and weighted network consequence, rather than as a direct comparison between attacker types.




Figure 6.5: Edge contingency analysis on Network F , showing (a) WCL and (b) the number of isolated nodes.

Figure 6.5 shows the WCL and islanding consequences of individual edge removals. The results show that not every branch contingency causes islanding with impact. Redundant connections can preserve source reachability after the removal of an individual branch. Consequently, some edge removals result in $WCL = 0$, whereas the removal of non-redundant branches can disconnect downstream nodes and produce a positive WCL.

6.2.3. Robustness trajectory comparison

As explained in Section 5.5, the IEMDE removal sequence is reused for this comparison. After each node removal, the distance-based EGR proxy is calculated for the resulting network state. The proxy is normalised by its value in the intact network, so that the trajectory starts at one.

The normalised EGR proxy is compared with the existing $1 - WCL$ trajectory from the IEMDE analysis. Both curves therefore follow the same convention, where one represents the intact network and zero represents the loss of the measured connectivity or weighted demand. The EGR proxy is not used to select the removed assets and does not generate an independent attack sequence. The comparison tests whether both measures show a similar degradation pattern under the same IEMDE sequence.




Figure 6.6: Comparison of the normalised distance-based EGR proxy and 1 – WCL along the IEMDE node-removal sequence.

Figure 6.6 shows that the two trajectories follow a similar overall pattern but are not identical. The distance-based EGR proxy declines more rapidly during the first removals and reaches a short plateau before reaching zero. The 1 – WCL curve declines more gradually because it weights the disconnected nodes according to their local energy values. The EGR proxy, in contrast, is based on the sum of pairwise distances between reachable real nodes and does not apply these energy weights.

Both curves reach zero after the same number of removals in the IEMDE sequence. This indicates that they identify the same endpoint under this particular sequence, although they respond differently during the intermediate stages. However, this shared endpoint is not independent evidence, because both curves use the same IEMDE node-removal sequence and source-reachability criterion. Moreover, the distance-based proxy excludes node pairs that become disconnected, so its decline reflects both changes in shortest-path distance and the loss of reachable node pairs. The result should therefore be interpreted as a sequence-specific consistency check between WCL and a shortest-path distance proxy, rather than as an independent validation of WCL or of the IEMDE target selection.

6.3. Overlap across frameworks

Tables 6.4 and 6.5 present the top-10 rankings of the assessment methods side by side. These methods include the VI, betweenness centrality, targeted attacks, and contingency analysis. For the targeted attacks, the attack strategy producing the fastest deterioration of weighted connectivity was selected for the comparison. For each removal assessment method, the top-10 ranking is compared against the VI top-10 to identify how many assets overlap between the two.

For the top-10 edges, targeted attacks with the betweenness decision rule share 4 out of 10 assets with the VI, while targeted attacks with IEMDE share 5 out of 10. For the top-10 nodes, targeted attacks with the betweenness decision rule share 5 out of 10 assets with the VI, targeted attacks with the degree decision rule share 4 out of 10, and targeted attacks with IEMDE share 3 out of 10.

Bold entries indicate assets that appear in multiple rankings. They therefore highlight assets that are identified as important by more than one assessment approach, rather than implying that every bold asset appears in all rankings.

For the edge rankings, an identifier such as 25/26 refers to the numbered edge records representing one physical connection. It does not indicate a connection between Node 25 and Node 26.

Table 6.4: Top 10 edges ranked by all vulnerability frameworks.

Rank	VI	Targeted Attacks (Betweenness)	Targeted Attacks (IEMDE)	Contingency Analysis
1	29/30	13/14b	25/26	25/26
2	17/18	37	3/4	3/4
3	25/26	7/8b	42/43	34/36
4	9/10	23/24	44/45	42/43
5	7/8	17/18c	5/6	44/45
6	3/4	25/26	38/39	5/6
7	34/36	1/10	15/16	38/39
8	42/43	1/21	32/33b	15/16
9	44/45	29/30	21/22	32/33a
10	38/39	32/33b	40/41a	32/33b

Table 6.5: Top 10 nodes ranked by all vulnerability frameworks.

Rank	VI	Targeted Attacks (Betweenness)	Targeted Attacks (Degree)	Targeted Attacks (IEMDE)	Contingency Analysis
1	5	18	18	1	1
2	1	11	1	2	2
3	4	2	2	11	11
4	12	1	11	18	26
5	9	9	8	26	18
6	23	5	26	5	22
7	3	26	7	9	8
8	6	25	5	13	12
9	7	16	22	19	5
10	25	6	4	15	10

7

Discussion

This chapter discusses the findings of this research and places them in the context of the research objective, existing literature and the case-study network. Section 7.1 first reflects on the modelling purpose of the two frameworks and explains why asset-specific vulnerability and structural criticality were assessed. Section 7.2 then discusses what the frameworks and their results reveal about vulnerability and structural criticality in the case-study network, and compares these findings with existing literature and expert insights from Stedin. Section 7.3 reflects on the methodology as a whole, discussing how the characteristics of the case-study network shaped the methodological choices, how these choices affect the interpretation of the results, and how the developed approach relates to alternative vulnerability and structural-criticality assessments in the literature. The chapter subsequently addresses the limitations of the research (7.4), its implications for Stedin and other DSOs, and directions for future research (7.5, 7.6).

7.1. Modelling purpose: bridging structural criticality and asset vulnerability

The objective of this research, and consequently of the two assessment frameworks, is to analyse how vulnerability emerges in electricity networks, which characteristics contribute to it, and whether assets that are structurally critical are also more vulnerable to deliberate physical threats. This requires a distinction between two related but different concepts. In this research, asset vulnerability refers to the susceptibility of an asset to a deliberate physical attack and to the asset-specific characteristics that influence this susceptibility. Structural criticality refers to the network level consequence of losing an asset, conditional on its removal.

Existing graph-theoretic approaches identify critical assets based on their position within the network, but generally do not account for asset-specific characteristics such as physical security, accessibility or capacity redundancy. Existing asset-specific vulnerability approaches, in turn, primarily focus on component fragility under hazards other than deliberate physical attacks on electricity networks. Consequently, structural criticality and malicious-threat vulnerability have generally not been examined together for the same network and under the same threat framing.

Addressing this gap, this research develops two complementary assessment frameworks. The first is an asset-specific Vulnerability Index (VI), which combines physical, operational and security-related characteristics to identify assets that are vulnerable to deliberate physical threats (**SQ1**, **SQ2**). The second is a graph-theoretic structural criticality assessment. It combines decision rules representing different attacker information levels with removal-based assessments that quantify the resulting network-level consequences. Degree centrality and origin-destination betweenness are used as decision rules for a topology-informed attacker (*uninformed attacker*), while an impact-based decision rule represents an impact-informed attacker (*informed attacker*) with additional knowledge of the consequences of asset removal. Single-asset $N - 1$ contingency analysis evaluates the consequence of removing each asset independently, whereas sequential targeted attack simulations evaluate how net-

work performance deteriorates when assets are repeatedly removed according to these decision rules. In both removal-based assessments, Weighted Connectivity Loss is used to quantify the resulting loss of weighted demand.

By applying both frameworks to the same case-study network, this research directly examines whether asset-specific vulnerability and structural criticality identify the same assets (**SQ5**). The contribution is therefore not that one framework replaces the other, but that the two perspectives are assessed together within the same network and threat context. This provides a direct basis for evaluating where vulnerability and structural impact converge and where they remain distinct.

7.2. What the two frameworks reveal about the network

The table below summarises the main findings and noteworthy observations derived from the results. While some findings directly answer the research subquestions, others highlight observations that are particularly relevant because they differ from, or extend, existing research. The findings are organised according to the corresponding research subquestions and grouped into the overarching themes of asset vulnerability and structural criticality.

Theme	Finding	Summary
Understanding vulnerability	SQ1	Malicious-threat-relevant characteristics could be identified across all four vulnerability categories (Topological, Spatial/-Physical, Operational and Security), indicating that the index created provides a comprehensive indication of vulnerability.
	SQ1	A subset of edges score noticeably higher on the VI than others.
	SQ3	Physical separation and capacity redundancy emerge as the strongest and most consistent determinants of vulnerability.
Understanding structural criticality	SQ4	The most effective attack strategy depends on the attacker's objective
	SQ4	Node removal fragments the radial network more rapidly than edge removal, requiring a smaller proportion of assets to achieve impact.
	SQ4	The $N - 1$ assessment reveals a physical redundancy gap. In the context of physical redundancy, the N-1 philosophy does not yet apply everywhere.
Comparing the frameworks	SQ5	The degree of overlap between asset-specific vulnerability and structural criticality depends on the assessment method and asset type.

Table 7.1: Overview of the main findings organised by discussion theme.

Theme 1: Understanding vulnerability

The first two findings look at how vulnerability to malicious threats emerges: which properties make a station or connection vulnerable, and which of these properties matter most.

7.2.1. The VI covers a broad range of malicious-threat-relevant characteristics across different categories

Most existing studies build their characteristics around a single category. Section 2 showed that most academic literature describe vulnerability purely through the network's topology. Salman et al. (2025) and Zeh et al. (2021) work at the asset level, but each around another hazard. Montalà-Palau et al. (2026) comes closest to a multidimensional view, combining asset-level exposure with a structural,

power-flow-based assessment.

This thesis instead mapped characteristics across four categories at once: Topological, Spatial/Physical, Operational, and Security (Table C.1). Using part of this literature and expert interpretations at Stedin. Characteristics could be found in every category, not just one.

The idea that vulnerability needs more than one lens is therefore not new. Montalà-Palau et al. already show this for a different hazard. What is new in this research is the threat lens itself. None of the studies above build their characteristics around a malicious attacker. This is likely why the Security category leans more on Stedin's internal expertise than on academic literature: existing frameworks are built for storms or aging infrastructure, not for someone actively choosing a target.

Eventually, it was possible to develop a new framework specifically build for physical malicious threats, covering all threat-origin characteristics. Even given the specific structure of the network and the chosen unit of the impact factor in the VI, it has been possible to identify which points within the network are critical while using only the index.

7.2.2. Physical and capacity redundancy is the strongest and most consistent determinant of vulnerability

The correlation analysis in this thesis (Section 6.1.1) supports observations made by Bilis et al. (2013) and Holmgren (2007b), who emphasize that the effectiveness of redundancy depends not only on its presence, but also on its physical independence.

Bilis et al. (2013) and Koc et al. (2015) argue that what matters is whether a backup is truly physically independent. Koc et al. (2015) refers to this as the quality of redundancy. A co-located backup shares the same route as the asset it protects. It offers little real protection. A single local event can disable both at once. This is known as a common cause failure.

This is consistent with the findings of this research, where physical and capacity redundancy, specifically whether a redundant connection or installation is physically separated from the asset it backs up, correlates more strongly with the overall VI than any other characteristic ($p < 0.05$). Several other characteristics did not meet the statistical significance threshold.

This finding also highlights an important difference from topology-based approaches such as Forsberg et al. (2023), who intentionally assess structural vulnerability using network topology alone. From a purely topological perspective, two redundant cables installed in the same cable trench are structurally identical to two physically separated redundant cables: both are simply represented as parallel edges between the same pair of nodes. The results of this research demonstrate that this distinction becomes crucial when considering deliberate physical attacks. A structurally central asset is therefore only as resilient as its redundancy allows, whereas graph-theoretic centrality metrics alone cannot distinguish between physically independent and co-located redundancy.

7.2.3. A subset of high-scoring edges share overhead and redundancy characteristics

The results show that a small group of edges score clearly higher on the VI than the rest. These edges share two properties. They are mostly overhead lines and are co-located with a redundant circuit and in some cases the capacity redundancy is exceeded.

This pattern follows directly from how the VI is built. Two characteristics in Table C.1 target exactly this situation. Overhead line exposure increases vulnerability, since an overhead line is easier to reach and damage than an underground cable. Distance between redundant circuits increases vulnerability, since redundant lines placed close together share the same physical risk. When an edge is both overhead and running parallel to its backup, both characteristics score high at the same time. The VI then combines these two exposures into one high score.

This is consistent with Finding SQ3 (Section 7.2.2). Physical separation of redundancy already emerged as the strongest determinant of vulnerability. These edges show what that finding looks like in practice. Under a malicious threat, the co-located pair can fail together, from one single event.

Theme 2: Understanding structural criticality

The previous findings explain what drives vulnerability to malicious threats. The next three findings shift to the structural framework and look at how the network's topology shapes the consequences of an attack.

7.2.4. The most effective attack strategy depends on the attacker's objective

The distinction between topology-informed (*uninformed*) and impact-informed (*informed*) attackers is relevant to the interpretation of the targeted attack results. The original IEMDE formulation selects the Electrical Most Damaging Element according to the increase in unsatisfied energy, whereas IMCE selects the most central remaining element (Albarakati & Bikdash, 2022). In the present study, IEMDE is implemented using WCL as the impact measure. It therefore represents an impact-informed attacker who can evaluate the weighted connectivity consequence of removing each candidate asset. IMCE represents an attacker who uses information on the network topology, but does not use asset-specific impact information. The comparison between these strategies therefore provides insight into the added value of impact information for targeted attacks.

The sequential attack results show that there is no single decision rule that is most effective throughout the complete attack sequence. Instead, the relative effectiveness of impact-informed and topology-informed target selection depends on the attacker's objective: causing the largest reduction in connected demand with a limited number of removals, or reaching maximum network fragmentation with the fewest removals.

For node removal, IEMDE initially causes the strongest reduction in remaining weighted connectivity. During approximately the first 15% of node removals, the impact-informed strategy therefore identifies more damaging targets than OD-betweenness. After this point, however, the OD-betweenness-based IMCE catches up. Both strategies eventually reach complete fragmentation after removing 7 of the 26 candidate nodes (27%). Degree-based IMCE performs considerably worse, requiring 18 of the 26 candidate nodes (69%) to reach its fragmentation endpoint and temporarily stagnating with approximately 2% of weighted connectivity remaining. Thus, although impact information provides an advantage during the initial removals, OD-betweenness ultimately identifies an equally efficient sequence for complete node fragmentation.

A similar pattern is observed for edge removal, but the difference at the fragmentation endpoint is more pronounced. IEMDE again produces the stronger initial reduction in remaining weighted connectivity and remains the more damaging strategy during approximately the first 19% of edge removals. Beyond this point, OD-betweenness becomes more effective. The OD-betweenness-based IMCE reaches the maximum observed fragmentation after removing 19 of the 46 edges (41%), whereas IEMDE requires 25 removals (54%). Neither strategy reduces the remaining weighted connectivity below approximately 39%, as several top-of-tree nodes remain connected to the sources regardless of the subsequent edge removals.

These findings show that the value of additional attacker information depends on how attack effectiveness is defined. If the objective is to cause the largest possible reduction in connected demand during the early stages of an attack, the impact-informed IEMDE strategy is most effective for both nodes and edges. If the objective is instead to reach maximum fragmentation with the fewest removals, OD-betweenness performs equally well for nodes and better for edges. The results therefore do not support the general conclusion that an impact-informed attacker is necessarily more effective than a topology-informed attacker (Albarakati & Bikdash, 2022).

The strong performance of OD-betweenness can be related to the predominantly radial, tree-like structure of the case-study network. Radial distribution networks are characterised by constrained paths between supply and demand and a low degree of redundancy, making their functioning strongly dependent on the position of assets within these paths (Hartmann et al., 2025). OD-betweenness directly captures how frequently an asset occurs on paths between the defined source and demand nodes. In this network, this relatively limited topological information therefore already provides a strong indication of which assets are structurally critical.

7.2.5. Node removal causes more damage to the network than edge removal

Across both targeted attack strategies, node removal fragments the case-study network more rapidly than edge removal. As shown in the previous subsection, complete node fragmentation is reached after removing 27% of the candidate nodes under both IEMDE and OD-betweenness, whereas maximum edge fragmentation requires 54% and 41% of the candidate edges, respectively. This difference reflects the predominantly radial, tree-like structure of the network, in which downstream stations typically depend on a limited number of upstream paths. Removing a station also removes its incident connections and can therefore disconnect several downstream branches simultaneously.

The contingency analysis shows the same general asymmetry. The maximum WCL for node contingencies is approximately 17.5%, whereas the maximum WCL for edge contingencies is approximately 8%. The most damaging node contingency therefore causes substantially greater weighted connectivity loss than the most damaging edge contingency. This result is consistent with research showing that network performance can deteriorate more rapidly under node removal than under link removal, particularly in networks with hub-like or tree-like structures (Bompard et al., 2011; Saei et al., 2025).

The results nevertheless require careful interpretation. Neither edge-attack sequence reduces the remaining weighted connectivity to zero in this case study. Both stabilise at approximately 39% remaining weighted connectivity because the local weights associated with several top-of-tree nodes remain connected to the sources regardless of which edge is removed. The results therefore demonstrate that node removal is more effective at fragmenting the case-study network, but they do not imply that every individual node outage is more damaging than every individual edge outage.

7.2.6. The $N - 1$ assessment reveals a physical redundancy gap

The contingency analysis provides an additional finding concerning the physical interpretation of the $N - 1$ criterion. Of the 46 analysed subconnections, 12 cause islanding with impact when removed individually. This corresponds to approximately 26% of the assessed subconnections.

Under the source-reachability criterion used in this study, these subconnections therefore do not provide full physical $N - 1$ security. The specific criterion on which the DSOs' network designs are based. Their removal disconnects one or more load-bearing nodes from all available sources, resulting in a positive WCL. This finding does not imply that the network operator's $N - 1$ policy is generally invalid. It shows that capacity-based $N - 1$ compliance and physical source-reachability $N - 1$ security are not necessarily equivalent.

In particular, a network may contain nominally redundant connections while still remaining vulnerable to the loss of one connection if the redundancy does not provide an independent physical route or if the remaining connection cannot maintain source reachability. The result therefore identifies a specific gap between the presence of redundancy in the network design and the ability of that redundancy to prevent physical islanding.

Theme 3: Comparing the frameworks

The final finding brings the two frameworks together and addresses the last central question of this thesis: to what extent are assets identified as vulnerable by the asset-specific VI also identified as structurally critical by the removal-based assessment?

7.2.7. The relation between vulnerability and structural criticality depends on the assessment perspective

The comparison shows that asset-specific vulnerability and structural criticality identify partly overlapping, but different sets of priority assets. The extent of this overlap depends on the asset type and the structural assessment method.

The contingency analysis provides the most direct comparison with single-asset criticality because each asset is removed independently, without assuming a particular target-selection strategy. For nodes, only three occur in both the VI top-10 and contingency top-10. For edges, the overlap is considerably larger, with six of the ten highest-ranked VI edges also occurring in the contingency top-10. Asset-specific vulnerability and single-removal structural criticality therefore appear more closely aligned for edges than for nodes.

This difference can partly be explained by how physical redundancy is represented in Network F . Physically separate redundant circuits are represented as separate edges, meaning that part of the physical redundancy captured by the VI is also reflected in the structural edge assessment. Physically separated installations within a station, however, are not represented as separate nodes. Station-specific characteristics such as physical security, accessibility and internal redundancy can therefore influence the VI without directly affecting structural node criticality. A station can consequently be structurally critical while being relatively well protected, or vulnerable without occupying one of the most structurally critical positions in the network.

The targeted attack results further show that the relationship depends on the information used for target selection (Albarakati & Bikdash, 2022). For nodes, OD-betweenness shows the largest overlap with the VI (5/10), followed by degree (4/10) and IEMDE (3/10). For edges, IEMDE shows the largest overlap (5/10), compared with OD-betweenness (4/10). Greater attacker information therefore does not automatically result in greater agreement with the VI. Instead, different decision rules identify different aspects of criticality: topology-informed strategies prioritise structural position, while IEMDE prioritises weighted consequences after removal.

The results therefore show that structural criticality cannot be used as a direct proxy for asset-specific vulnerability. Their relationship depends on the asset type, the representation of redundancy and the method used to identify critical assets. The two frameworks should therefore be interpreted as complementary: structural analysis identifies assets whose removal is particularly consequential for network connectivity, while the VI identifies assets that combine potential consequences with characteristics that make them vulnerable to deliberate physical threats.

7.3. Academic contribution and methodological reflection

Unlike Section 7.2, which discusses what the findings reveal about vulnerability and structural criticality, this section reflects on the methodological choices underlying these findings. The focus is therefore on why the selected methods were suitable for the predominantly radial case-study network, which alternative approaches could have been used, and how these choices affect the interpretation and transferability of the results.

7.3.1. Scope of the asset-specific vulnerability assessment

Asset-specific vulnerability assessments are inherently context dependent. The characteristics that determine vulnerability depend on the threat mechanism, asset type and consequence considered. Existing multidimensional approaches therefore combine different categories of information depending on their application, including physical exposure, susceptibility, resistance, recoverability and operational system performance (Hromada et al., 2021; Liu et al., 2025; Montalà-Palau et al., 2026).

The threat-specific design of the VI consequently limits its direct transferability. Other hazards or network types may require different characteristics or different relative importance of those characteristics. The contribution of the VI is therefore not a universally applicable vulnerability score, but a structured method for combining asset-specific characteristics relevant to deliberate physical threats.

7.3.2. Choice of structural decision rules

The choice of structural decision rules followed from the predominantly radial structure. Origin-destination betweenness was selected as the main topological decision rule because the functioning of the case-study network depends primarily on source-to-demand paths. Restricting betweenness to these origin-destination pairs makes the measure more closely aligned with the functional role of the network than conventional betweenness over all node pairs.

Degree centrality was included as a complementary local decision rule. Degree captures the number of direct connections of an asset, whereas betweenness captures its position within the wider network. This distinction is also observed by Cetinay et al. (2018), who show that degree mainly reflects local structure, while betweenness contains information about global network position. Their targeted attack results also demonstrate that high-betweenness nodes can have a substantial effect on network performance.

The results of this research provide additional support for this choice. OD-betweenness reaches com-

plete node fragmentation after the same number of removals as the impact-informed IEMDE strategy and reaches the maximum observed edge fragmentation with fewer removals. Degree performs considerably worse for node fragmentation. For this network, source-to-demand dependency therefore appears to provide more useful information on structural criticality than local connectivity alone.

Several alternative decision rules could nevertheless have been applied. A particularly relevant example is Upstream Robustness proposed by Koc et al. (2015), specifically for radial distribution networks. It evaluates supply robustness based on the availability, disjointness and length of alternative paths to the sources.

Conceptually, Upstream Robustness and OD-betweenness capture related structural dependencies, as both consider the dependence of demand on assets located between sources and loads. Upstream Robustness provides a more detailed representation of redundancy by evaluating alternative paths and their quality, whereas OD-betweenness considers shortest source-to-demand paths. Importantly, Koc et al. (2015) show that conventional betweenness may overlook source-related criticality because it does not distinguish between source and load nodes. The OD-betweenness used here partly addresses this limitation by defining source and demand nodes explicitly. Given the predominantly radial structure of Network F , where most assets have only one effective source-to-demand route, the additional complexity of explicitly evaluating all alternative paths was not considered necessary. Upstream Robustness nevertheless provides a relevant alternative for future applications with greater meshing.

Electrical Node Significance, introduced by Koc et al. (2013), provides another possible comparison, but addresses a different aspect of criticality. It reflects the share of total distributed power associated with an individual node and is therefore primarily a local electrical importance measure. It does not directly represent whether removing that node disconnects other downstream assets. WCL instead evaluates the total weighted demand that loses source connectivity after a removal, thereby also capturing downstream consequences. WCL therefore provides a more suitable consequence measure for the radial dependencies investigated in this research.

Other electrically informed decision rules were also considered less suitable as primary selection metrics for this network. Flow betweenness and electrical closeness incorporate electrical flow behaviour and effective resistance, respectively, and can therefore provide additional information in meshed power networks where flows redistribute over several active paths (Cetinay et al., 2018). In an unweighted tree, however, effective resistance reduces to shortest-path distance, illustrating that the additional information provided by such measures becomes smaller as the network becomes more radial.

Effective Graph Resistance (EGR) of Koc et al. (2013) was therefore not used as a primary decision rule. EGR is naturally a global network-level robustness measure rather than a direct per-asset ranking. Using it for target selection would require evaluating the change in EGR after every candidate removal, resulting in a procedure comparable in computational structure to IEMDE. Moreover, because EGR primarily reflects the availability and quality of alternative electrical paths, its added value as a selection rule is more limited in a network with relatively little meshed redundancy. EGR was instead used as a global robustness consistency check after the removal sequences. This allows the WCL-based degradation to be compared with a structural robustness measure that is not itself used for target selection. Previous work similarly positions EGR as a network-level robustness metric rather than primarily as an asset-ranking measure (Ellens et al., 2011).

Taken together, the selected decision rules represent a deliberate balance between structural relevance, interpretability and required information. OD-betweenness captures global source-to-demand dependency, degree provides a simple local comparison, and IEMDE provides an impact-informed benchmark. More detailed measures such as Upstream Robustness or electrically informed centralities could enrich the analysis, particularly in networks with more alternative paths or stronger electrical meshing, but were not considered necessary as primary decision rules for the present network.

7.3.3. Scope of the removal-based assessments

The contingency analysis and targeted attack simulations use asset removal for different purposes. The $N - 1$ contingency analysis evaluates each asset independently and is therefore suitable for identifying assets whose individual loss produces a large structural consequence. Similar contingency approaches

remove individual buses or branches and evaluate the resulting structural or operational impact (Coelho et al., 2019).

An advantage of this assessment is that the resulting criticality does not depend on an assumed attacker decision rule. Each asset is subjected to the same removal test, making the contingency analysis particularly useful for comparing single-asset structural criticality with the asset-specific VI. Its limitation is that it does not represent how an attacker selects targets or how criticality changes after earlier removals.

The targeted attack simulations address this dynamic perspective. IMCE and IEMDE remove assets sequentially and recalculate the next target after the network has changed. Similar sequential attack strategies are used by Albarakati and Bikdash (2022) and Cetinay et al. (2018). The resulting damage therefore depends on both the selected decision rule and the changing network structure.

The distinction between IMCE and IEMDE is methodologically relevant. IMCE selects the next target from a structural metric, whereas IEMDE evaluates the WCL consequence of each possible next removal before selecting the most damaging candidate. IEMDE therefore requires more information and computation, from an attackers' perspective.

The removal-based assessments nevertheless represent only specific attack scenarios. Independent $N - 1$ contingencies do not capture combinations of failures, while sequential targeted attacks assume that assets are removed one after another and that the next target is selected after the resulting network state is known. Simultaneous $N - k$ failures could expose combinations of assets or shared dependencies that remain hidden under these approaches and would require a separate generalized contingency assessment (Narimani et al., 2020).

7.3.4. Positioning the methodology within broader power-grid vulnerability research

Recent power-grid vulnerability research increasingly combines multiple dimensions of system behaviour rather than relying on topology alone. The way these dimensions are combined, however, depends strongly on the threat considered, the available data and the purpose of the assessment.

A useful comparison is provided by Sperstad et al. (2020), who distinguish between a system's susceptibility to a threat and its ability to cope with the consequences after failure. Their framework therefore separates different parts of vulnerability rather than treating them as one network property. The present research follows a comparable logic by assessing asset-specific vulnerability through the VI and structural consequence separately, before comparing where both perspectives converge.

A more detailed alternative is presented by Gjorgiev and Sansavini (2021), who combine AC power-flow-based cascading failure simulation with optimization to identify critical transmission assets and high-impact outage combinations. Such an approach captures power redistribution, overloads and cascading behaviour that are not represented by WCL. It would therefore be relevant for a more meshed network or when sufficient operational data are available. For the predominantly radial case-study network, however, loss of source connectivity itself already represents an important failure mechanism, making a structural removal-based approach suitable for the present research objective.

Recent attack research also shows that critical targets depend on the assumed knowledge of the attacker. Maleki et al. (2025), for example, distinguish between informed and limited-knowledge attackers and show that the available system information affects how attacks can be formulated. Although their work concerns cyber-based load-altering attacks, the same methodological principle is relevant here: the assets identified as critical depend partly on the information used to select them.

The contribution of the present methodology therefore lies mainly in explicitly separating and comparing asset-specific vulnerability and structural criticality within the same network and threat context. The VI identifies characteristics that make assets vulnerable to deliberate physical threats, while the structural assessments identify the consequences of their removal under different target-selection assumptions. This approach is less physically detailed than full power-flow or cascading-failure models, but provides a transparent way to distinguish vulnerability, structural consequence and attacker knowledge in a predominantly radial network.

7.4. Limitations

As with any modelling study, not all excluded factors carry the same weight. Some were excluded by design, falling outside the scope of a physical, malicious-sabotage threat framing; others reflect genuine of data or methodological constraints. This section distinguishes between the two.

7.4.1. Methodological limitations

Relative, not absolute, scoring. The *VI* produces a relative ranking of assets within the case-study network rather than an absolute measure of vulnerability. No external benchmark exists against which the absolute severity of a vulnerability score can be calibrated. Consequently, the *VI* indicates which assets are more vulnerable than others within the studied network, but it does not quantify the overall vulnerability of the network itself. Moreover, such an absolute measure would depend on the scope of assessment, for example whether vulnerability is considered at the level of a single DSO, the Dutch electricity system, or the European grid, as well as on the objective of a potential attacker.

Hazard treated as a constant. The *VI* is formulated as $\text{Exposure} \times \text{Impact}$, with Hazard treated as a sector-level constant rather than as an asset-specific variable. A more comprehensive formulation could explicitly model the probability of attack for individual assets. The probability of such an attack to happen, would have given a better representation of the vulnerability scores.

Abstracted network representation. The structural framework represents the network as stations and the connections between them. Internal station components, such as busbars, transformers, circuit breakers and control systems, are not modelled individually. This abstraction was chosen deliberately to maintain a tractable graph representation that reflects the network primarily from an attacker's perspective. Possible extensions towards a more detailed network representation are discussed in Section 7.6.

Node-level redundancy not represented structurally. Network *F* explicitly represents physically separated redundant connections between stations as separate edges, but it does not explicitly represent physically separated redundant installations within substations. The *VI* does account for this distinction through its Impact calculation. Consequently, the two frameworks are not entirely equivalent: the *VI* incorporates both edge-level and node-level redundancy, whereas the structural framework represents only edge-level redundancy. Extending the graph representation to include internal station redundancy would therefore be a logical direction for future research.

Impact represented by annual energy demand. The Impact component of the *VI* is represented by a single annual energy value per asset (GWh), corrected where applicable with *N-1* redundancy. This representation has two consequences. First, temporal variations in demand, such as seasonal effects, daily load fluctuations or the exact timing of an attack, are not captured. An attack occurring during a winter peak may therefore result in a different impact than suggested by the annual value. Second, other dimensions of impact, such as restoration time, critical customers, operational effects or societal consequences, are not incorporated. The framework therefore represents one specific interpretation of impact, rather than its complete operational or societal consequences.

Impact assigned to vulnerable station compartments. For substations, the Impact component of the *VI* is assigned only to the station compartment that is not physically separated from its redundant counter compartment. Consequently, the annual energy demand is attributed to the vulnerable installation rather than to the station as a whole. This modelling choice reflects the objective of assessing asset-specific vulnerability, but assumes that the identified non-separated compartment represents the part of the station that is susceptible to failure. Alternative compartmentalisation or protection concepts could therefore lead to different impact scores.

Attacker types not modelled. The targeted attack simulations model a topology-informed attacker (IMCE) and an impact-informed attacker (IEMDE), but not an attacker who selects targets based on ease of attack using the *VI* itself, nor an attacker who executes simultaneous, coordinated multi-asset attacks rather than sequential ones (Section 7.3.3). Cyber-physical or load-altering attackers, who manipulate operational demand rather than physically damaging assets, are likewise outside the scope of this research (Section 7.3.4). The framework therefore covers a specific, deliberate physical-sabotage threat profile rather than the full spectrum of possible malicious actors.

7.4.2. Data limitations

Measurement-level sensitivity of nominal and ordinal characteristics. Several characteristics included in the VI are measured on a nominal or ordinal scale rather than a continuous (ratio) scale. Unlike continuous measures, where a small data-collection error typically produces a proportionally small error in the outcome, a misclassification of a nominal or ordinal characteristic changes its category outright, with no intermediate degree of error possible. This risk is not merely theoretical: the sensitivity analysis (Section 6.1.1) shows that some characteristics are unstable precisely because many assets sit close to a threshold boundary, meaning a single incorrect classification during data collection could shift an asset across that boundary and materially change its VI score. Accurate classification of these characteristics therefore had to be correct at the point of data collection, and could not be independently re-verified for every asset within the scope of this study.

Mixed time displays of input data. The Impact component of the VI is based on Stedin's 2030 load forecasts, whereas most exposure characteristics represent the current state of the network. Consequently, the framework assumes that these asset characteristics remain unchanged over the planning horizon. Future changes in physical infrastructure, redundancy or security measures may therefore alter the resulting vulnerability ranking.

Data collection through expert input. Much of the physical layout data (cable depth, routing, physical separation of redundant components) was poorly documented and had to be collected verbally from area-responsible experts across the organization. During this process, it became apparent that at least one area-responsible expert had misinterpreted the data request. This reflects a broader risk in expert-elicited data collection: experts tend to interpret requests through the lens of their own area of expertise, as does the researcher collecting it, which may introduce inconsistencies that are difficult to detect after.

Application to a limited case-study network. The proposed frameworks were applied to a relatively small subset of Stedin's subtransmission network rather than to the complete network. Although the selected case study captures the key characteristics of a predominantly radial network, the findings may not fully generalise to larger or more meshed network configurations. Applying the frameworks to the complete Stedin network would provide a stronger validation of their robustness and general applicability.

7.4.3. Technical limitations

Manual weight decomposition. The local station weights used in the Weighted Connectivity Loss (WCL) calculation are derived from cumulative station load profiles using a manually constructed parent-child hierarchy. Although this approach was verified with Stedin experts and is appropriate for the case-study network, it is labor-intensive and may become impractical for larger distribution systems. Future implementations would benefit from an automated decomposition procedure based directly on the network topology.

7.5. Implications for Stedin

Most importantly, the VI and the structural framework should be used together, not as substitutes. The VI already reflects Stedin's current physical, operational and security measures, but it does not include structural criticality as an exposure factor. An informed attacker aiming for maximum impact would rationally prefer structurally critical assets. Structural criticality could therefore serve as a proxy for how likely an asset is to be deliberately targeted, and combining both frameworks would give Stedin a more complete view of likely attacker targets within a resilience assessment. This link was not incorporated in the current VI, since attack probability was treated as fixed in this research, but it could be a concrete direction for future refinement.

Secondly, the findings of this study complicate a commonly heard phrase at Stedin that their network should simply become more meshed to improve resilience. The results cannot fully support this recommendation. As long as *N-1* policy continues to be maintained, increasing the degree of meshing is not the primary concern. Instead, what matters most is that redundancy is physically separated, rather than only structurally or capacity-wise redundant. This directly reflects the finding that physical and capacity redundancy, rather than topological position alone, are the dominant determinants of vulnera-

bility (Section 7.2.2). The results further show that, provided physical and capacity-based redundancy is ensured together with an adequate level of protection against physical exposure, asset vulnerability remains low. For Stedin specifically, the main practical recommendation is therefore to ensure that redundant connections are physically separated wherever possible and that $N-1$ capacity redundancy is consistently maintained throughout the network. The quality of redundancy is therefore as important as its presence.

Thirdly, the VI results point to a small, identifiable set of high-risk connections rather than a general vulnerability across the network. A subset of four edges score noticeably higher on the VI than the rest. Most of these are overhead lines that run in masts, so run without physically separated redundancy. For Stedin, this gives a concrete set of candidates for targeted protection: examine whether the redundant connections along these corridors can be physically separated, and assess whether converting these overhead sections to underground cabling is feasible. The same applies to overhead lines crossing critical infrastructure such as waterways and bridges, which are similarly exposed and accessible.

A further point, although less prominent in the results, concerns physical security. A limited number of substations do not yet meet the level of protection prescribed by Netbeheer Nederland, indicating a gap between the required and implemented security measures (Section 4.2, measures deficit). Addressing these deficiencies would further reduce asset vulnerability.

Although the findings originate from a single case-study network, the proposed methodology is not specific to Stedin. As discussed in Section 7.5.1, the same assessment approach could, in principle, be applied to other distribution system operators, provided that the vulnerability characteristics are re-evaluated for the local network and operational context.

7.5.1. Implications for DSOs

Both frameworks developed in this study, the Vulnerability Index and the structural criticality framework, are in principle applicable to other Distribution System Operators (DSOs), since neither is built around metrics or characteristics that are inherently specific to Stedin's network. Applying the frameworks elsewhere does, however, require revisiting the characteristics identified in Section 4.2, to assess whether the exposure characteristics included were relevant primarily because they vary meaningfully within Stedin's network specifically. Cable depth is an example of this: cable depth is largely uniform across Stedin's network, so it contributes little discriminating value to the VI in this case study, whereas for a DSO with more heterogeneous cable depths, the same characteristic could become a meaningful exposure indicator. The reverse can also occur: a characteristic that varies substantially within Stedin's network may turn out to be uniform elsewhere, and would then no longer contribute meaningfully to that DSOs Vulnerability Index. Re-evaluating which characteristics actually discriminate between assets is therefore a necessary step when applying this framework to a different DSO.

7.6. Future research

Several directions for future research follow directly from both the findings and the methodological limitations of this study.

The first concerns the level of abstraction used in the proposed frameworks. Throughout this research, substations are represented as single nodes connected by electricity lines. Future research could extend both frameworks by explicitly representing the internal topology of substations, including components such as busbars, transformers, circuit breakers and protection systems. This would not only allow the structural framework to identify critical components within substations, but would also extend the VI to additional asset types beyond substations and connections.

A related extension concerns the representation of redundancy. While Network F explicitly distinguishes physically separated redundant circuits between substations, physically separated redundant installations within substations are not represented in the network. Extending the graph representation to include these internal redundancies would make the structural framework more consistent with the VI and thus real-time vulnerability representation. Similarly, the fictional nodes introduced in Network F to represent physically separated cable routes sometimes correspond to actual physical riser points. Now these riser points exist as solely modelling constructs, future research could investigate whether these locations should be represented as assets in their own right as different characteristics could be

assigned to these as well and therefore incorporated in the *VI*.

From a practical perspective, the most immediate next step would be to apply both frameworks to the complete Stedin network rather than the case-study subset considered in this research. In addition, the influence of the more detailed representation adopted in Network *F* could be evaluated directly by comparing the attack simulations with those obtained from the original network representation (without fictional nodes). Such a comparison would provide further insight into the value of incorporating physical characteristics into structural network models.

If the scope is broadened to the complete network, the structural framework itself could also be further developed. For this predominantly radial case-study network, betweenness centrality and degree centrality were used during targeted attacks. A larger and more meshed network may therefore require additional graph-theoretic metrics or alternative attack strategies to better represent structural criticality.


The scope of the contingency and targeted attack assessments could also be extended. This research uses single-asset ($N - 1$) contingency analysis and sequential targeted attacks, but does not evaluate simultaneous $N - x$ combinations of asset failures (Section 7.3.3). Future research could apply a generalized $N - x$ contingency approach, for example combining group betweenness with electrical measures such as Line Outage Distribution Factors (Narimani et al., 2020), to identify combinations of assets whose simultaneous failure causes the most severe consequences.

The targeted attack simulations could similarly be extended with additional attacker profiles. The present research distinguishes a topology-informed attacker (IMCE) from an impact-informed attacker (IEMDE), but does not model an attacker who selects targets using the *VI* itself. Such an attacker would combine physical exposure, security and redundancy characteristics with structural or impact information, representing an adversary with the most complete knowledge that could realistically be obtained about the network. Comparing this attacker profile with the topology- and impact-informed strategies used in this research would show how much additional damage becomes possible once asset-specific vulnerability information is combined with structural knowledge.

The scope of the *VI* could also be broadened beyond deliberate physical sabotage. The characteristics developed in this study were selected specifically for malicious physical threats. Extending the framework to include other threat types, such as cyber threats or combined cyber-physical attacks, would provide a more comprehensive vulnerability assessment. Future research could additionally investigate incorporating restoration time into the impact formulation, allowing multiple restoration scenarios to be evaluated once specific attack scenarios have been defined.

Recent events in Germany illustrate a further direction for future research (“‘Duitse klimaatactivist wil ook in Nederland toeslaan op energienet’”, 2026). In at least one case, the motive for an attack was connected to nearby electricity generation from lignite. This characteristic, proximity to environmentally sensitive generation sources, was excluded from the current *VI* as a model choice (Appendix B.1). These events suggest that such environmental characteristics can, in practice, form part of an attacker’s motive. Future research could therefore re-examine whether environmental factors deserve a place in the Exposure component of the *VI*.

Finally, this thesis considers vulnerability primarily from a top-down perspective, following the hierarchical structure of the electricity network. Stedin, however, is increasingly investing in flexibility solutions through flexibility tenders. Future research could therefore investigate how these flexibility resources contribute to resilience from the bottom up, particularly during deliberate physical disruptions in which conventional supply paths become unavailable.



Conclusion

The objective of this research was to analyse how vulnerability emerges in electricity subtransmission networks, which characteristics contribute to it, and whether assets that are structurally critical are inherently more vulnerable to deliberate physical threats. This objective was addressed through the following research question:

How can vulnerable and critical assets in an electricity network under physical malicious threats be identified?

To answer this question, two complementary assessment frameworks were developed: an asset-specific Vulnerability Index (VI) and a graph-theoretic structural criticality assessment. The results demonstrate that vulnerability to deliberate physical threats cannot be explained by network topology alone. The structural assessment identifies assets whose removal has substantial consequences for network connectivity, whereas the VI additionally captures physical, operational and security-related characteristics that influence the vulnerability of individual assets. Structural criticality and asset-specific vulnerability therefore represent different, but complementary, perspectives.

The VI further shows that vulnerability is not determined equally by all characteristics. Physical and capacity redundancy emerged as the strongest and most consistent determinants of vulnerability. In particular, the results demonstrate the importance of whether redundancy provides a physically independent alternative. From a malicious-threat perspective, the quality of redundancy is therefore more informative than its mere presence.

The structural assessment shows that the identification of critical assets depends on both the type of asset and the information used for target selection. For sequential attacks, impact-informed IEMDE causes the largest initial reduction in remaining weighted connectivity for both nodes and edges. However, additional impact information does not necessarily result in faster overall fragmentation. For nodes, IEMDE and OD-betweenness both reach complete fragmentation after removing 7 of the 26 candidate nodes (27%). For edges, OD-betweenness reaches the maximum observed fragmentation after 19 of the 46 removals (41%), compared with 25 removals (54%) for IEMDE. Degree performs considerably worse for node fragmentation. These results show that source-to-demand topological information alone can already provide a strong indication of structural criticality in this predominantly radial network, while impact information is particularly valuable when the objective is to maximise immediate weighted disruption.

The comparison between the two frameworks forms the last contribution of this thesis. The rankings overlap only partially, and the extent of this overlap depends on the asset type and assessment method. Assets that are structurally critical are therefore not necessarily the assets that are most vulnerable to deliberate physical attacks, and vice versa. Moreover, the targeted attack results demonstrate that which assets become relevant targets also depends on the information available to the attacker. Structural criticality should therefore not be used as a direct proxy for asset-specific vulnerability.

The answer to the main research question is therefore that vulnerable and critical assets are best

identified by assessing these two perspectives separately and subsequently combining their insights. The VI identifies which assets are particularly vulnerable based on their asset-specific characteristics and potential impact, while the structural assessment identifies which assets are most consequential when removed and how this changes under different attacker information assumptions. Together, the frameworks distinguish which assets are vulnerable, which are structurally critical, and where both characteristics coincide. This provides a more complete basis for prioritising assets against deliberate physical threats than either perspective alone.

Societal relevance

The societal relevance of this research lies in supporting more resilient electricity distribution networks against deliberate physical attacks. As society becomes increasingly dependent on electricity, disruptions to distribution networks can have consequences beyond the network operator itself.

For Stedin, and for distribution system operators more broadly, the results show why resilience planning should not rely on structural criticality alone. Doing so may overlook assets that are particularly vulnerable because of their physical, operational or security characteristics. Conversely, prioritising vulnerability without considering structural consequences may direct protection towards assets whose failure has relatively limited network impact. Combining both perspectives provides a more targeted basis for prioritising protective measures.


Ultimately, improving resilience starts with identifying both which assets are vulnerable to attack and which assets matter most when they fail. This research provides a framework for making that distinction.

References

- 'Duitse klimaatactivist wil ook in Nederland toeslaan op energienet'. (2026, September). Retrieved September 8, 2026, from <https://nos.nl/artikel/2630085-duitse-klimaatactivist-wil-ook-in-nederland-toeslaan-op-energienet>
- Albarakati, A. J., & Bikdash, M. (2022). Empirical framework for identification of the most harmful malicious attacks on a smart grid. *Applied Network Science*, 7(1), 13. <https://doi.org/10.1007/s41109-022-00451-5>
- Belcic, I., & Smalley, I. (2026). Grid Resilience. <https://www.ibm.com/think/topics/grid-resilience>
- Billis, E. I., Kroger, W., & Cen Nan. (2013). Performance of Electric Power Systems Under Physical Malicious Attacks. *IEEE Systems Journal*, 7(4), 854–865. <https://doi.org/10.1109/JSYST.2012.2223512>
- Bompard, E., Wu, D., & Xue, F. (2011). Structural vulnerability of power systems: A topological approach. *Electric Power Systems Research*, 81(7), 1334–1340. <https://doi.org/10.1016/j.epsr.2011.01.021>
- Cetinay, H., Devriendt, K., & Van Mieghem, P. (2018). Nodal vulnerability to targeted attacks in power grids. *Applied Network Science*, 3(1), 34. <https://doi.org/10.1007/s41109-018-0089-9>
- Coelho, E. P. R., Paiva, M. H. M., Segatto, M. E. V., & Caporossi, G. (2019). A New Approach for Contingency Analysis Based on Centrality Measures. *IEEE Systems Journal*, 13(2), 1915–1923. <https://doi.org/10.1109/JSYST.2018.2881558>
- Ellens, W., Spieksma, F., Van Mieghem, P., Jamakovic, A., & Kooij, R. (2011). Effective graph resistance. *Linear Algebra and its Applications*, 435(10), 2491–2506. <https://doi.org/10.1016/j.laa.2011.02.024>
- Forsberg, S., Thomas, K., & Bergkvist, M. (2023). Power grid vulnerability analysis using complex network theory: A topological study of the Nordic transmission grid. *Physica A: Statistical Mechanics and its Applications*, 626, 129072. <https://doi.org/10.1016/j.physa.2023.129072>
- Galvan, G., & Agarwal, J. (2020). Assessing the vulnerability of infrastructure networks based on distribution measures. *Reliability Engineering & System Safety*, 196, 106743. <https://doi.org/10.1016/j.ress.2019.106743>
- Gjorgiev, B., & Sansavini, G. (2021). Identifying and assessing power system vulnerabilities to transmission asset outages via cascading failure analysis. *Reliability Engineering & System Safety*, 217, 108085. <https://doi.org/10.1016/j.ress.2021.108085>
- Hartmann, B., Soha, T., Cirunay, M. T., & Erdei, T. (2025). Inherent structural vulnerabilities in power distribution networks: A longitudinal complex network analysis. *Applied Network Science*, 10(1), 31. <https://doi.org/10.1007/s41109-025-00724-9>
- Holmgren, Å. J. (2007a). A Framework for Vulnerability Assessment of Electric Power Systems. In *Critical Infrastructure* (pp. 31–55). Springer Berlin Heidelberg.
- Holmgren, Å. J. (2007b). A Framework for Vulnerability Assessment of Electric Power Systems [Series Title: Advances in Spatial Science]. In A. T. Murray & T. H. Grubessic (Eds.), *Critical Infrastructure* (pp. 31–55). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-68056-7_3
- Hromada, M., Rehak, D., & Lukas, L. (2021). Resilience Assessment in Electricity Critical Infrastructure from the Point of View of Converged Security. *Energies*, 14(6), 1624. <https://doi.org/10.3390/en14061624>
- Islam Khan, M. Z., Ferdowsi, F., & Zappi, M. E. (2026). Port Resilience Index for Service-Centric DER Optimization in Distribution Systems. *2026 IEEE Texas Power and Energy Conference (TPEC)*, 1–6. <https://doi.org/10.1109/TPEC67884.2026.11513127>
- Iwan Baan. (2012). City and the Storm [photo]. Retrieved July 28, 2026, from <https://iwan.com/tag/city-and-the-storm/>
- Koc, Y., Warnier, M., Kooij, R. E., & Brazier, F. M. T. (2013). A robustness metric for cascading failures by targeted attacks in power networks. *2013 10th IEEE INTERNATIONAL CONFERENCE ON*

- NETWORKING, SENSING AND CONTROL (ICNSC)*, 48–53. <https://doi.org/10.1109/ICNSC.2013.6548709>
- Koc, Y., Raman, A., Warnier, M., & Kumar, T. (2015, June). Structural Vulnerability Analysis of Electric Power Distribution Grids [arXiv:1506.08641 [eess.SY]]. <https://doi.org/10.48550/arXiv.1506.08641>
- Koç, Y., Warnier, M., Miegheem, P. V., Kooij, R. E., & Brazier, F. M. (2014). The impact of the topology on cascading failures in a power grid model. *Physica A: Statistical Mechanics and its Applications*, 402, 169–179. <https://doi.org/10.1016/j.physa.2014.01.056>
- Lange, P. d. (2025, January). Sabotage onder water: De Russische schaduwoorlog legt kwetsbaarheid van het Westen bloot [Section: Kijk Verder]. Retrieved September 8, 2026, from <https://www.volkskrant.nl/kijkverder/v/2025/rusland-datakabels-sabotage-olietankers~v1243286/>
- Liu, L., Wang, H., Li, J., & Ren, X. (2025). DSR model for safety vulnerability assessment of power grid infrastructure in Henan Province China. *Scientific Reports*, 15(1), 29283. <https://doi.org/10.1038/s41598-025-07805-w>
- Maleki, S., Pan, S., Lakshminarayana, S., & Konstantinou, C. (2025, April). Survey of Load-Altering Attacks Against Power Grids: Attack Impact, Detection and Mitigation [arXiv:2410.22007 [eess.SY]]. <https://doi.org/10.48550/arXiv.2410.22007>
- Ministerie van Justitie en Veiligheid. (2025a). Critical Infrastructure Protection. <https://english.nctv.nl/topics/c/critical-infrastructure-protection>
- Ministerie van Justitie en Veiligheid. (2025b). Increasing and Changing Threats to Critical Infrastructure. <https://english.nctv.nl/latest/news/2025/11/14/increasing-and-changing-threats-to-critical-infrastructure>
- Ministerie van Justitie en Veiligheid. (2026). Vitale infrastructuur en Wet weerbaarheid kritieke entiteiten. <https://www.nctv.nl/onderwerpen/v/vitale-infrastructuur>
- Montalà-Palau, M., Mañé, M. C., & Gomis-Bellmunt, O. (2026). How resilient is a power system? An open-source tool integrating GIS and Optimal Power Flow for resilience assessment. *Energy Nexus*, 21, 100630. <https://doi.org/10.1016/j.nexus.2026.100630>
- Narimani, M. R., Huang, H., Ummnakwe, A., Mao, Z., Sahu, A., Zonouz, S., & Davis, K. (2020, July). Generalized Contingency Analysis Based on Graph Theory and Line Outage Distribution Factor [arXiv:2007.07009 [math.OC]]. <https://doi.org/10.48550/arXiv.2007.07009>
- Pagani, G. A., & Aiello, M. (2015). A complex network approach for identifying vulnerabilities of the medium and low voltage grid. *International Journal of Critical Infrastructures*, 11(1), 36. <https://doi.org/10.1504/IJCIS.2015.067394>
- Phase to Phase. (n.d.). Phase to Phase – Netten voor distributie van elektriciteit, hoofdstuk 2. Retrieved June 25, 2026, from https://www.phasetophase.nl/boek/boek_1_2.html#_2.4.2
- Proximities. (2026, February). Stroomuitval in Berlijn legt kwetsbaarheid van infrastructuur bloot. <https://www.proximities.com/insight/sabotage-in-berlijn-legt-kwetsbaarheid-van-kritieke-verbinding-en-bloot>
- Saei, S., Barker, K., Tajik, N., & Ermagun, A. (2025). On topological measures and network vulnerability patterns: A comparative analysis. *Reliability Engineering and System Safety*, 266, 111608. <https://doi.org/10.1016/j.ress.2025.111608>
- Salman, U., Wang, Z., Chen, F., Acquah, M., & Stevens, D. (2025). Risk-Aware Vulnerability Assessment: Mesh-View Grid Mapping for Distribution System Resilience. *2025 IEEE Kiel PowerTech*, 1–6. <https://doi.org/10.1109/PowerTech59965.2025.11180691>
- Sargent, R. G. (2020). Verification And Validation Of Simulation Models: An Advanced Tutorial. *2020 Winter Simulation Conference (WSC)*, 16–29. <https://doi.org/10.1109/WSC48552.2020.9384052>
- Scribbr. (n.d.). Statistiek voor Beginners | 5 Stappen & Voorbeelden. Retrieved August 3, 2026, from <https://www.scribbr.nl/category/statistiek/>
- Shafiei, K., Zadeh, S. G., & Hagh, M. T. (2025). Robustness and resilience of energy systems to extreme events: A review of assessment methods and strategies. *Energy Strategy Reviews*, 58, 101660. <https://doi.org/10.1016/j.esr.2025.101660>
- Sperstad, I. B., Kjølle, G. H., & Gjerde, O. (2020). A comprehensive framework for vulnerability analysis of extraordinary events in power systems. *Reliability Engineering & System Safety*, 196, 106788. <https://doi.org/10.1016/j.ress.2019.106788>

- Stedin. (n.d.). Open liggingsdata kabels en leidingen Stedin. <https://www.stedin.net/zakelijk/open-data/liggingsdata-kabels-en-leidingen>
- Techniek Nederland. (2026). *Sterk in tijden van crisis: Hoe de technieksector haar weerbaarheid vergroot* (tech. rep.). echniek Nederland.
- Zeh, A., Dollichon, J., & Spitzer, H. (2021). Integrated Electrical and Long-Term Analysis for the Medium and Low Voltage Grids of KWH Netz. *CIREN 2021 - The 26th International Conference and Exhibition on Electricity Distribution*, 2487–2491. <https://doi.org/10.1049/icp.2021.2041>



AI Statement

Artificial Intelligence (AI) was used as a supporting tool during this research. At all times, care was taken to ensure that confidential or sensitive information from Stedin was not disclosed to AI systems. All AI-generated outputs were critically reviewed, verified and, where necessary, adapted by the author before being incorporated into this thesis. The author remained fully responsible for the content, methodology, analyses and conclusions presented in this thesis.

AI was used for the following activities:

- Improving the clarity, grammar and academic style of written text;
- Summarising, interpreting and comparing scientific literature;
- Assisting with the development, debugging and interpretation of Python code;
- Supporting brainstorming on methodological choices and the interpretation of results.

B

Model Choices

This appendix summarises the main modelling choices and abstractions made during the development of the Vulnerability Index and the structural criticality framework.

Table B.1: Overview of the main modelling choices and abstractions.

Model choice	Implementation	Justification	Consequence or limitation
Station-load calculation	For the local node weight w_i used in the WCL calculation, the throughput of a station's direct child is subtracted from its own station profile: $L_{node,t} = L_{station,t} - L_{child,t}.$	The available Stedin planning profiles are cumulative: an upstream station profile already includes the demand of its downstream stations. Summing these profiles directly would therefore double count the same demand. This differs from GWh_n used in the Vulnerability Index, which uses the cumulative (non-decomposed) throughput.	The decomposition depends on a manually constructed parent-child hierarchy. An incorrect hierarchy would affect the local weights and therefore the Weighted Connectivity Loss.
Transformers included within station nodes	Transformers are not modelled as separate nodes or edges; they are included as station-level characteristics in the Vulnerability Index (e.g. physical exposure of transformers, distance to station fence).	Because the case-study network is radial rather than meshed, transformers are physically located within the substation and do not form separate branching points in the topology. In a meshed network, transformers can sit between switching locations rather than being confined to one substation, so this simplification would not hold there. See B.1.	The frameworks cannot distinguish between the failure of an individual transformer and the loss of the station as a whole.
Riser points included within connections	Riser points (<i>opstijgpunten</i>), where an underground cable transitions into an overhead line, are represented as a characteristic of the corresponding edge rather than as separate nodes.	Riser points form part of the physical structure of an overhead line rather than a distinct switching or transformation point; their relevance lies in the physical exposure of the connection (overhead line exposure). They are therefore double weighted (see Figure B.2a.	No separate criticality or vulnerability score is produced for individual riser points; their vulnerability is captured only through the characteristics of the complete connection.

Continued on next page

Table B.1 – continued from previous page

Model choice	Implementation	Justification	Consequence or limitation
Different level of station detail between the two frameworks	The Vulnerability Index evaluates individual station compartments (e.g. transformers, switchgear) that could be targeted separately. The structural framework treats each station as one complete node.	The two frameworks answer different questions: the VI assesses asset-specific exposure and impact at the compartment level relevant to an attacker's choice of target, while the structural framework assesses the network-wide connectivity consequence of losing a station as a whole.	Node-level (within-station) redundancy is represented in the Vulnerability Index's Impact score but not in the structural graph. The two frameworks are therefore not framed identically, which should be kept in mind when comparing their rankings.
Cable burial depth	Cable burial depth was mapped as a candidate spatial exposure characteristic but was not included with meaningful weight in the final Vulnerability Index.	Cable depth is uniform across Stedin's network in the case-study area, so it contributes little discriminating value between assets here.	For a DSO, or a part of the network, with more heterogeneous cable depths, this characteristic could become a meaningful exposure indicator; its exclusion here reflects a case-study-specific finding rather than a general judgement on its relevance.
Weighted Connectivity Loss instead of Unserved Load	The IEMDE targeted-attack strategy substitutes Weighted Connectivity Loss (WCL) for the Unserved Load metric used in the original formulation of Albarakati and Bikdash (2022).	Among other reasons, WCL was chosen because, for confidentiality reasons, the original load behind each asset could not be shown directly in the results. WCL expresses impact as a relative percentage of the network's total local weight rather than as an absolute energy figure.	As a relative measure, WCL does not convey the absolute scale of unserved energy behind an asset, only its share of the network's total local weight.

Continued on next page

Table B.1 – continued from previous page

Model choice	Implementation	Justification	Consequence or limitation
Physical separation threshold for redundant connections ($\tau = 1$ m)	Redundant connections whose GIS-derived distance from each other falls below 1 metre are classified as not physically separated ($d_e < \tau$).	In practice, co-located redundant connections in the Stedin network typically run through the same cable duct, with a physical separation of approximately 50 to 100 cm. Connections within this range are considered to share a common failure risk, since damage to the shared duct or a physical attack on it would affect both simultaneously. This was confirmed by Stedin staff early in the data collection process. Later in the process, however, a different source indicated that the failure of one cable does not necessarily imply the failure of the other, since a connection can be laid using two separate drillings rather than one, which would place the two cables in physically independent ducts even along the same route. This research takes the first scenario into account, in which one duct serves multiple routes.	Data on the number of drillings used per connection was not collected. Because this information could not be verified for individual assets, the model conservatively assumes a single drilling, i.e. two cables sharing the same trench section, even though some connections may in practice have been laid with two separate drillings and could therefore be more physically independent than assumed. The 1 metre threshold is a fixed, uniform cutoff based on typical duct-separation practice; it does not capture graded degrees of common-cause risk beyond this distance, nor site-specific variation in duct or drilling construction

Continued on next page

Table B.1 – continued from previous page

Model choice	Implementation	Justification	Consequence or limitation
Environmental factors such as proximity to a coal-fired power station and/or other potentially sensitive environments surrounding assets have been excluded	This characteristic was considered as a candidate spatial characteristic during the identification phase (Section 4.1), but was not included in the final Vulnerability Index.	The VI is scoped to characteristics that affect either the probability of an asset being deliberately targeted (Exposure) or the severity of the consequences for the electricity network itself (Impact). Proximity to a coal-fired power station or other sensitive environment primarily relates to broader environmental or collateral safety consequences, rather than to the accessibility of the asset or the severity of the disruption to the network.	The VI does not capture assets whose disruption could trigger secondary environmental or safety incidents, even though such proximity could, in principle, increase the wider consequences of a malicious attack beyond the electricity network alone.
Overhead line towers treated as independent, not co-located	Circuits sharing the same overhead line tower (mast) are modelled as independent connections. The exposure indicators do not represent a shared tower as a form of physical co-location.	The framework assumes that overhead lines are not physically co-located with one another, so a common cause failure arising from physical proximity falls outside the scope of the exposure indicator. In reality, however, two circuits mounted on the same tower are physically co-located.	A shared-tower attack, such as sabotage, fire or a drone strike, could disable both circuits simultaneously. This risk is not captured by the current exposure indicators. Explicitly modelling shared towers as co-located assets is a concrete direction for future refinement.





Figure B.1: Meshed vs Radial busbar explanation for nodes (AI generated)



(a) Overhead line representation Google Earth (f7-f8)



(b) Riser points (author's own photograph)

Figure B.2: Overhead lines representation abstract and real-time

C

Network Characteristics

Table C.1: Overview of vulnerability characteristics used in this research.

Category	Attribute	Source	Explanation
Topological	Line length	(Forsberg et al., 2023; Koç et al., 2014)	The physical length of a connection contributes to the length of paths through the network and may therefore affect both the structural role of a connection and the quality of alternative routes.
	Position in the network	(Cetinay et al., 2018; Forsberg et al., 2023; Pagani & Aiello, 2015; Saei et al., 2025)	A station's position as a cut vertex, hub or end node determines its structural criticality, independently of its physical condition. Centrality measures are commonly used to identify such structurally important nodes and connections.
	Coupling to the upstream power grid	(Bilis et al., 2013; Koc et al., 2015)	A station connected directly to the 150 or 380 kV network may represent a primary supply point and can therefore be critical for supplying downstream parts of the network.
Spatial / Physical	Physical exposure	(Bilis et al., 2013; Montalà-Palau et al., 2026; Saei et al., 2025; Sperstad et al., 2020)	Indicates how physically accessible an asset is. Transformers may be located in open-air installations, while overhead lines and externally accessible facilities are exposed to physical damage and environmental hazards.
	Cable depth	<i>Case-study data and operational knowledge</i>	Deeper buried cables are more difficult to access and damage than cables located close to the surface. This characteristic is therefore based on case-study data and operational knowledge rather than on a directly matching source in the reviewed literature.
	Co-location with other infrastructure	(Liu et al., 2025; Montalà-Palau et al., 2026)	Sharing a corridor with roads, bridges or railways may increase exposure to accidental and intentional damage. Infrastructure-crossing types and surrounding hazards are therefore relevant spatial vulnerability characteristics.
	Distance to roads	(Liu et al., 2025; Montalà-Palau et al., 2026)	Straight-line distance to roads can represent accessibility for a potential attacker as well as accessibility for inspection, repair and emergency response.
	Environmental context	(Liu et al., 2025; Montalà-Palau et al., 2026; Sperstad et al., 2020)	Environmental and spatial conditions, such as exposure to flooding, heat, vegetation, machinery or other hazards, may increase the likelihood that an asset is affected by an external event.
	Physical separation of redundant components	(Koc et al., 2015; Pagani & Aiello, 2015)	Redundant switchgear or cable circuits should be physically separated. Otherwise, a single local disruption or common-cause event may disable both components simultaneously. The quality of redundancy depends not only on the number of alternative paths, but also on their disjointness and configuration.
Operational	Grid overload	(Albarakati & Bikdash, 2022; Bilis et al., 2013; Koc et al., 2013; Koç et al., 2014)	Overload sensitivity is relevant because an initial outage can redistribute power flows and trigger subsequent line failures or cascading effects.
	Restoration time	(Holmgren, 2007a; Liu et al., 2025; Montalà-Palau et al., 2026; Sperstad et al., 2020)	Restoration time determines how long the consequences of a disruption persist. Restoration duration is therefore directly related to the amount of energy or service that remains unavailable after an asset failure.
	Unserved energy	(Albarakati & Bikdash, 2022; Holmgren, 2007a; Montalà-Palau et al., 2026)	Energy not supplied following an asset failure. Depending on the method, this can be calculated through power-flow analysis or approximated using the affected demand and the duration of the outage.
Security	Telecommunication disruption	(Bilis et al., 2013); Techniek Nederland Weerbaarheidsrapport (2026)	Loss of internet and telecommunications caused by sabotage is relevant because power-grid operation depends on monitoring, communication and control systems such as SCADA.
	Physical security and accessibility	(Bilis et al., 2013; Holmgren, 2007a); Techniek Nederland Weerbaarheidsrapport (2026)	Physical access control, fencing, cameras and sabotage protection may reduce the probability of a successful physical attack. The accessibility and protection level of an asset therefore influence its exposure to deliberate damage.

Continued on next page

Table C.1 continued from previous page

Category	Attribute	Source	Explanation
	Security measure deficit	<i>Internal and sector-specific security information</i>	Represents the difference between the security measures required by Netbeheer Nederland and the additional measures currently implemented by Stedin. This characteristic is based on internal and sector-specific security information.
	Transparency of asset location	(Bilis et al., 2013; Cetinay et al., 2018; Holmgren, 2007a)	Assets that can easily be located through public GIS, infrastructure data or knowledge of the network layout may be more exposed to targeted attacks. The availability of topological information can support the selection of critical targets.

Table C.2: Operationalisation of the vulnerability characteristics included in the Vulnerability Index.

Characteristic	Node	Edge	Unit / score type	Scale	Scoring methodology	Data source
Redundancy exceedance (N-1, based on peak-load forecasts)	✓	✓	Ratio (%)	0–1	Percentage of annual peak-load hours during which the forecasted load exceeds the available N-1 capacity. The value is already expressed as a proportion between 0 and 1.	PowerBI and Stedin internal load forecasts
Physical separation of redundant switchgear	✓	–	Ordinal (3 categories)	0 or 1	A score of 1 is assigned when redundant switchgear is physically separated and fire-compartmentalised. Co-located switchgear receives a score of 2 and in case of no redundant component present, the score is 3.	Stedin internal data and expert assessment
Distance to redundant connection	–	✓	Ratio (m)	0–1	Physical distance between a connection and its redundant connection. Smaller distances indicate a lower probability that both connections are affected by the same malicious disruption. The normalised score is therefore inverted so that 1 results in affecting both circuits.	Stedin GIS
Annual energy supplied behind the station or connection	✓	✓	Ratio (GWh)	0–1, min–max	Total annual energy dependent on the station or connection. Higher energy values represent a greater potential impact and therefore receive a higher normalised score.	PowerBI and Stedin internal load profiles
Physical exposure of transformers	✓	–	Ordinal (2 categories)	0–1	Transformers are classified according to their physical exposure to the surroundings. More openly accessible installations receive a higher score than enclosed or protected installations.	Stedin internal data and expert assessment
Distance between transformers and station fencing	✓	–	Ratio (m)	0–1, min–max	Physical distance between the transformer and the station boundary. Shorter distances indicate greater physical accessibility and therefore receive a higher vulnerability score. The normalised score is inverted.	Stedin internal station data
Real-time physical security monitoring	✓	–	Ordinal (5 categories)	0–1, rescaled	Security monitoring classification established by Netbeheer Nederland. The original ordinal scale is rescaled to a value between 0 and 1. Detailed categories are omitted for security reasons.	Stedin internal security data
Security measure deficit	✓	–	Ordinal (3 categories)	0–1, rescaled	Difference between the required security measures defined and implemented by Stedin. A larger deficit results in a higher vulnerability score. Detailed measures are omitted for security reasons.	Stedin internal security data
Exposure to critical infrastructure (waterways and bridges)	–	✓	Ratio (% of line length)	0–1	Percentage of the total connection length located along, beneath or across waterways, bridges or other critical infrastructure. A larger exposed share results in a higher vulnerability score.	Derived from Stedin GIS
Overhead exposure of lines	–	✓	Ratio (% of line length)	0–1	Percentage of the connection installed above ground. A larger overhead share represents greater physical accessibility and therefore receives a higher vulnerability score.	Stedin internal data and GIS

D

Vulnerability Index

Table D.1: Complete overview of node vulnerability scores.

Node ID	Exposure (E_v)	Impact (I_v)	VI_v
1	0.190	0.431	0.082
2	0.095	0.000	0.000
3	0.143	0.079	0.011
4	0.286	0.272	0.078
5	0.238	0.544	0.130
6	0.071	0.143	0.010
7	0.071	0.123	0.009
8	0.286	0.002	0.001
9	0.198	0.192	0.038
10	0.000	0.000	0.000
11	0.000	0.000	0.000
12	0.198	0.198	0.039
13	0.238	0.004	0.001
14	0.000	0.000	0.000
15	0.000	0.060	0.000
16	0.000	0.000	0.000
17	0.000	0.000	0.000
18	0.000	0.007	0.000
19	0.159	0.000	0.000
20	0.238	0.000	0.000
21	0.071	0.065	0.005
22	0.000	0.000	0.000
23	0.238	0.115	0.027
24	0.000	0.000	0.000
25	0.143	0.056	0.008
26	0.143	0.000	0.000

Table D.2: Complete overview of edge vulnerability scores.

Edge ID	Exposure (E_e)	Impact (I_e)	VI_e
1	0.000	0.000	0.000
2	0.000	0.000	0.000
3	0.200	0.265	0.053
4	0.200	0.265	0.053
5	0.076	0.002	0.000
6	0.076	0.002	0.000
7	0.469	0.253	0.119
8	0.457	0.253	0.116
9	0.491	0.253	0.124
10	0.491	0.253	0.124
11	0.030	0.000	0.000
12	0.034	0.000	0.000
13	0.524	0.000	0.000
14	0.524	0.000	0.000
15	0.200	0.136	0.027
16	0.200	0.136	0.027
17	0.517	0.821	0.424
18	0.519	0.821	0.426
19	0.108	0.000	0.000
20	0.092	0.000	0.000
21	0.200	0.113	0.023
22	0.200	0.113	0.023
23	0.229	0.000	0.000
24	0.229	0.000	0.000
25	0.200	0.638	0.128
26	0.200	0.638	0.128
27	0.000	0.000	0.000
28	0.000	0.000	0.000
29	0.600	1.000	0.600
30	0.600	1.000	0.600
31	0.216	0.028	0.006
32	0.082	0.001	0.000
33	0.078	0.001	0.000
34	0.200	0.233	0.047
35	0.020	0.000	0.000
36	0.200	0.233	0.047
37	0.400	0.028	0.011
38	0.200	0.141	0.028
39	0.200	0.141	0.028
40	0.088	0.028	0.002
41	0.028	0.028	0.001
42	0.200	0.233	0.047
43	0.200	0.233	0.047
44	0.232	0.160	0.037
45	0.232	0.160	0.037

Table D.3: Statistical relationship between the node characteristics and the Vulnerability Index.

Characteristic	Measurement level	Test	Statistic	p-value	Interpretation
$E_{1,v}$	Ordinal	Spearman	$\rho = 0.740$	< 0.001	Strong positive relationship.
$E_{2,v}$	Ordinal (2 categories)	Welch's t-test	$d = 0.82$	0.057	Higher VI in category 2, but not statistically significant.
$E_{3,v}$	Ratio	Pearson	$r = -0.751$	0.005	Strong negative relationship.
$E_{4,v}$	Ordinal	Spearman	$\rho = 0.235$	0.247	Weak positive relationship.
$E_{5,v}$	Ordinal	Spearman	$\rho = 0.277$	0.170	Weak positive relationship.
$I_{1,v}$	Ratio	Pearson	$r = 0.666$	< 0.001	Strong positive relationship.
$I_{2,v}$	Ratio	Pearson	$r = -0.005$	0.979	No relationship observed.

Table D.4: Statistical relationship between the edge characteristics and the Vulnerability Index.

Characteristic	Measurement level	Test	Statistic	p-value	Interpretation
$E_{1,e}$	Ratio	Pearson	$r = 0.619$	< 0.001	Strong positive relationship.
$E_{2,e}$	Ratio	Pearson	$r = 0.347$	0.019	Moderate positive relationship.
$E_{3,e}$	Ratio	Pearson	$r = -0.211$	0.164	Weak negative relationship; not statistically significant.
$I_{1,e}$	Ratio	Pearson	$r = 0.121$	0.427	Very weak positive relationship; not statistically significant.
$I_{2,e}$	Ratio	Pearson	$r = 0.253$	0.094	Weak positive relationship; not statistically significant.

Table D.5: Sensitivity analysis of the Vulnerability Index ranking.

Characteristic	Scenario	Top 5 overlap	Order changed	Remarks
$E_{1,v}$	-1 level	3/5	Yes	Two new assets entered the top five.
$E_{1,v}$	+1 level	2/5	Yes	Three new assets entered the top five.
$E_{3,v}$	-10%	5/5	Yes	Only the order of ranks 4 and 5 changed.
$E_{3,v}$	-20%	5/5	Yes	Only the order of ranks 4 and 5 changed.
$E_{3,v}$	-30%	5/5	Yes	Only the order of ranks 4 and 5 changed.
$E_{3,v}$	+10%	5/5	No	No change observed.
$E_{3,v}$	+20%	5/5	No	No change observed.
$E_{3,v}$	+30%	5/5	No	No change observed.
$I_{1,v}$	-10% to +30%	5/5	No	No change observed.
$E_{1,e}$	-10% to -30%	2/5	Yes	Three new edges entered the top five.
$E_{1,e}$	+10% to +30%	5/5	No	No change observed.
$E_{2,e}$	-10% to -30%	5/5	No	No change observed.
$E_{2,e}$	+10% to +30%	5/5	Yes	Only the ranking order changed.

E

Targeted Attacks

Table E.1: Top 15 nodes by betweenness and degree centrality (Network F).

Betweenness centrality			Degree centrality		
Rank	Node	Value	Rank	Node	Value
1	18	0.0086	1	18	0.1200
2	11	0.0082	2	2	0.1200
3	8	0.0082	3	1	0.1000
4	12	0.0078	4	11	0.0800
5	2	0.0065	5	8	0.0600
6	1	0.0020	6	22	0.0600
7	26	0.0016	7	7	0.0600
8	9	0.0010	8	26	0.0600
9	22	0.0008	9	6	0.0600
10	5	0.0006	10	5	0.0400
11	20	0.0006	11	4	0.0400
12	7	0.0004	12	23	0.0400
13	4	0.0002	13	9	0.0400
14	23	0.0000	14	20	0.0400
15	13	0.0000	15	17	0.0400

Table E.2: Top 15 branches by edge betweenness centrality (Network F).

Rank	Edge	Value
1	29/30	0.0078
2	13/14a	0.0078
3	13/14b	0.0078
4	17/18b	0.0078
5	17/18d	0.0078
6	17/18a	0.0078
7	17/18c	0.0078
8	23/24	0.0051
9	25/26	0.0016
10	19/20	0.0012
11	37	0.0010
12	15/16	0.0008

Continued on next page

Table E.2 – continued from previous page

Rank	Edge	Value
13	32/33a	0.0008
14	32/33b	0.0008
15	5/6	0.0008

Table E.3: IMCE results for node removal using betweenness centrality, Network F.

Step i	Removed node a_i	$\phi(a_i \Gamma_i)$	WCL(Γ_{i+1}) (%)
1	18	0.008571	8.383234
2	11	0.001701	30.538922
3	2	0.001773	47.487633
4	1	0.001850	73.496485
5	9	0.001208	77.219474
6	5	0.002020	90.757615
7	26	0.001057	100.000000
8	25	0.000000	100.000000
9	16	0.000000	100.000000
10	6	0.000000	100.000000
11	10	0.000000	100.000000
12	22	0.000000	100.000000
13	4	0.000000	100.000000
14	19	0.000000	100.000000
15	17	0.000000	100.000000
16	21	0.000000	100.000000
17	3	0.000000	100.000000
18	12	0.000000	100.000000
19	8	0.000000	100.000000
20	20	0.000000	100.000000
21	13	0.000000	100.000000
22	15	0.000000	100.000000
23	14	0.000000	100.000000
24	24	0.000000	100.000000
25	7	0.000000	100.000000
26	23	0.000000	100.000000

Table E.4: IMCE results for node removal using degree centrality, Network F.

Step i	Removed node a_i	$\phi(a_i \Gamma_i)$	WCL(Γ_{i+1}) (%)
1	2	0.120000	16.948711
2	1	0.102041	42.957563
3	18	0.104167	51.340797
4	11	0.085106	73.496485
5	26	0.065217	82.738870
6	7	0.066667	82.738870
7	8	0.068182	82.738870
8	22	0.046512	82.738870
9	4	0.047619	85.758917
10	5	0.048780	91.017964
11	9	0.050000	100.000000
12	6	0.025641	100.000000

Continued on next page

Table E.4 – continued from previous page

Step i	Removed node a_i	$\phi(a_i \Gamma_i)$	WCL(Γ_{i+1}) (%)
13	19	0.026316	100.000000
14	25	0.027027	100.000000
15	3	0.027778	100.000000
16	12	0.028571	100.000000
17	13	0.029412	100.000000
18	23	0.030303	100.000000
19	20	0.031250	100.000000
20	24	0.000000	100.000000
21	14	0.000000	100.000000
22	21	0.000000	100.000000
23	17	0.000000	100.000000
24	10	0.000000	100.000000
25	15	0.000000	100.000000
26	16	0.000000	100.000000

Table E.5: IEMDE results for node removal, Network F.

Step i	Removed node a_i	$\tilde{z}_i$ (%)	WCL(Γ_{i+1}) (%)
1	1	17.833897	17.833897
2	2	25.123666	42.957563
3	11	10.830513	53.788076
4	18	19.708409	73.496485
5	26	9.242385	82.738870
6	5	5.259047	87.997917
7	9	12.002083	100.000000
8	13	0.000000	100.000000
9	19	0.000000	100.000000
10	15	0.000000	100.000000
11	6	0.000000	100.000000
12	7	0.000000	100.000000
13	25	0.000000	100.000000
14	10	0.000000	100.000000
15	4	0.000000	100.000000
16	24	0.000000	100.000000
17	16	0.000000	100.000000
18	3	0.000000	100.000000
19	23	0.000000	100.000000
20	22	0.000000	100.000000
21	20	0.000000	100.000000
22	21	0.000000	100.000000
23	12	0.000000	100.000000
24	14	0.000000	100.000000
25	17	0.000000	100.000000
26	8	0.000000	100.000000

Table E.6: IMCE results for edge removal using edge betweenness centrality, Network F.

Step i	Removed edge a_i	$\phi(a_i \Gamma_i)$	WCL(Γ_{i+1}) (%)
1	13/14a	0.007843	0.000000

Continued on next page

Table E.6 – continued from previous page

Step i	Removed edge a_i	$\phi(a_i \Gamma_i)$	WCL(Γ_{i+1}) (%)
2	23/24	0.000980	0.000000
3	37	0.000980	0.000000
4	7/8b	0.001569	8.279094
5	25/26	0.000784	16.219734
6	17/18d	0.000784	22.207758
7	21/22	0.000392	24.420724
8	19/20	0.000392	24.420724
9	11/12	0.000392	28.299922
10	3/4	0.000392	33.506899
11	44/45	0.000392	36.631086
12	38/39	0.000392	39.416819
13	42/43	0.000392	43.998959
14	32/33b	0.000392	46.602447
15	5/6	0.000392	49.648529
16	29/30	0.000392	54.985681
17	2	0.000196	54.985681
18	1	0.000392	59.281437
19	27	0.000196	59.281437
20	28	0.000392	61.390263
21	32	0.000000	61.390263
22	5	0.000000	61.390263
23	12	0.000000	61.390263
24	6	0.000000	61.390263
25	33	0.000000	61.390263
26	9/10	0.000000	61.390263
27	11	0.000000	61.390263
28	40/41c	0.000000	61.390263
29	7/8c	0.000000	61.390263
30	34/35/36b	0.000000	61.390263
31	20	0.000000	61.390263
32	7/8a	0.000000	61.390263
33	13/14b	0.000000	61.390263
34	31	0.000000	61.390263
35	19	0.000000	61.390263
36	17/18b	0.000000	61.390263
37	34/35/36a	0.000000	61.390263
38	40/41b	0.000000	61.390263
39	40/41a	0.000000	61.390263
40	32/33a	0.000000	61.390263
41	41	0.000000	61.390263
42	34/35/36	0.000000	61.390263
43	15/16	0.000000	61.390263
44	40	0.000000	61.390263
45	17/18a	0.000000	61.390263
46	17/18c	0.000000	61.390263

Table E.7: IEMDE results for edge removal, Network F.

Step i	Removed edge a_i	$\tilde{z}_i$ (%)	WCL(Γ_{i+1}) (%)
1	25/26	7.940640	7.940640
2	3/4	5.206977	13.147618

Continued on next page

Table E.7 – continued from previous page

Step i	Removed edge a_i	$\tilde{z}_i$ (%)	WCL(Γ_{i+1}) (%)
3	42/43	4.582140	17.729758
4	44/45	3.124186	20.853944
5	5/6	3.046082	23.900026
6	38/39	2.785733	26.685759
7	15/16	2.681593	29.367352
8	32/33b	2.603489	31.970841
9	21/22	2.212965	34.183806
10	40/41a	1.588128	35.771934
11	17/18b	0.000000	35.771934
12	29/30	8.643582	44.415517
13	7/8a	0.000000	44.415517
14	37	6.690966	51.106483
15	40/41b	0.000000	51.106483
16	34/35/36	0.000000	51.106483
17	1	0.000000	51.106483
18	11/12	0.000000	51.106483
19	2	4.295756	55.402239
20	19/20	3.879198	59.281437
21	12	0.000000	59.281437
22	13/14b	0.000000	59.281437
23	33	0.000000	59.281437
24	31	0.000000	59.281437
25	11	0.000000	59.281437
26	32/33a	0.000000	59.281437
27	20	0.000000	59.281437
28	34/35/36b	0.000000	59.281437
29	41	0.000000	59.281437
30	5	0.000000	59.281437
31	27	0.000000	59.281437
32	28	2.108826	61.390263
33	40	0.000000	61.390263
34	19	0.000000	61.390263
35	7/8c	0.000000	61.390263
36	17/18c	0.000000	61.390263
37	40/41c	0.000000	61.390263
38	7/8b	0.000000	61.390263
39	34/35/36a	0.000000	61.390263
40	6	0.000000	61.390263
41	13/14a	0.000000	61.390263
42	32	0.000000	61.390263
43	23/24	0.000000	61.390263
44	17/18d	0.000000	61.390263
45	9/10	0.000000	61.390263
46	17/18a	0.000000	61.390263

F

Contingency Analysis

Table F.1: Branch (edge) contingency results (Network F).

Edge nr	Redundant	Islanding	Isolated stations	Lost weight (%)
25/26	False	True	2	7.94
3/4	False	True	1	5.21
34/35/36	False	True	1	4.58
42/43	False	True	1	4.58
44/45	False	True	1	3.12
5/6	False	True	1	3.05
38/39	False	True	1	2.79
15/16	False	True	1	2.68
32/33a	False	True	1	2.60
32/33b	False	True	1	2.60
21/22	False	True	1	2.21
40/41a	False	True	1	1.59
17/18c	False	False	0	0.00
34/35/36b	False	False	0	0.00
2	False	False	0	0.00
1	False	False	0	0.00
40	True	False	0	0.00
7/8c	False	False	0	0.00
41	False	False	0	0.00
11/12	True	False	0	0.00
40/41c	True	False	0	0.00
19/20	False	False	0	0.00
12	True	False	0	0.00
11	True	False	0	0.00
32	False	False	0	0.00
17/18a	False	False	0	0.00
7/8a	False	False	0	0.00
34/35/36a	False	False	0	0.00
13/14a	True	False	0	0.00
9/10	False	False	0	0.00
40/41b	False	False	0	0.00
37	False	False	0	0.00
31	True	False	0	0.00
27	False	False	0	0.00

Continued on next page

Table F.1 – continued from previous page

Edge nr	Redundant	Islanding	Isolated stations	Lost weight (%)
28	False	False	0	0.00
29/30	False	False	0	0.00
13/14b	True	False	0	0.00
7/8b	False	False	0	0.00
17/18b	False	False	0	0.00
5	False	False	0	0.00
6	False	False	0	0.00
17/18d	False	False	0	0.00
23/24	True	False	0	0.00
19	False	False	0	0.00
20	False	False	0	0.00
33	False	False	0	0.00

Table F.2: Station (node) contingency results (Network F).

Asset	Islanding	Isolated stations	Lost weight (%)
1	True	3	17.83
2	True	4	16.95
11	True	2	10.83
26	True	3	9.24
18	True	3	8.38
22	True	2	7.94
8	True	2	5.99
12	True	1	5.34
5	True	1	5.26
10	True	1	5.21
16	True	1	4.58
19	True	1	4.58
6	True	1	4.30
7	True	1	3.88
9	True	1	3.72
24	True	1	3.12
3	True	1	3.05
4	True	1	3.02
21	True	1	2.79
15	True	1	2.68
25	True	1	2.60
14	True	1	2.21
23	True	1	2.16
17	True	1	2.11
13	True	1	1.59
20	True	1	1.51