Universiteit Leiden

Identifying Structural Vulnerabilities in the European Electricity Transmission Network

Reinout Offringa (s3455408)

Supervisors:
Frank Takes & Olga Gadyatskaya

Bachelor Thesis– Informatica & Economie

Leiden Institute of Advanced Computer Science (LIACS)
liacs.leidenuniv.nl

16/07/2026

Abstract

This thesis presents a network-driven method for identifying the most vulnerable regions of the European electricity transmission grid. Network studies of the European grid have characterised its structure and its response to failures, but no study combines a recent pan-European dataset, multiple removal strategies, and a mapping of the results back onto geography; power-flow models capture more electrical detail but remain computationally infeasible at continental scale. This thesis instead models a recent openly available pan-European dataset, PyPSA-Eur, as an unweighted graph of 6,795 substations spanning 36 countries, of which 25 are analysed, and simulates the progressive removal of its nodes under random, degree-based and betweenness-based strategies. From the resulting fragmentation we then propose a vulnerability score for each country. Our main finding is that the grid is highly tolerant of random failure, but much more vulnerable to targeted attack, with betweenness-based removal the most destructive. The two targeted strategies furthermore reveal that vulnerability to a degree-based attack is governed by degree assortativity, and vulnerability to a betweenness-based attack by the concentration of shortest-path traffic, while a few sparse grids such as Norway's are fragile under any failure at all. Interestingly, we find that, geographically, this vulnerability is not spread evenly but concentrates in a western corridor running from the Paris basin to the Spanish border, where the most central substations and the densest cross-border interconnectors coincide. Overall, our results show that our approach is able to locate the structurally critical regions of the European grid at continental scale, without the electrical physics of the grid.

Contents

1 Introduction

This chapter sets out the motivation for the thesis and the questions it addresses. Section 1.1 gives the background on electricity transmission grids and the network-based approach used to study them, Section 1.2 presents the research question and its sub-questions, and Section 1.3 outlines the structure of the remainder of the thesis.

1.1 Background

Electricity transmission grids are critical infrastructures, as illustrated by large-scale blackouts such as the 2003 outage in southern Sweden and eastern Denmark [GKL17]. Such events demonstrate how failures in the transmission system can disrupt electricity supply over large regions. Climate change-related natural disasters further increase the risk of power outages by placing additional stress on grid infrastructure, therefore testing its fault tolerance. The economic consequences of these disruptions can be severe; for example, the storm Gudrun caused an estimated economic damage of approximately €3 billion [GKL17].

Unlike natural disasters, there is a second type of threat that is deliberate rather than accidental. Between October 2022 and March 2023 Russia targeted almost all high-voltage substations in Ukraine that perform the task of transforming voltages from power plants to long-distance transmission lines. These attacks were not random, but specifically aimed at the most critical parts of the power grid. Some substations were retargeted due to their importance to the overall connectivity of the grid; within six months, 42 of Ukraine's 94 high-voltage substations had been damaged [SK24]. Targeted attacks make it necessary to understand which parts of the electricity grid are most critical, and how resilient the grid is to both random failures and deliberate attacks.

Modelling the grid as a network is a deliberate simplification of the actual power grid. By modelling the grid as a network in which nodes represent substations, and edges represent transmission lines [PA13], its structural properties can be assessed more directly. Research shows that modelling power grids as networks provides insight into their connectivity and reveals components that strongly influence system behaviour [AAN04]. By switching from a physical representation of the spatial grid to an abstract network model, structural patterns can be analysed that are difficult to observe using only geographical or engineering-based representations.

Representing the grid as a network does more than simplify it. Once the grid is modelled as a graph, the full set of tools developed in network science can be applied to it [Bar16]; without this step, the analysis would amount to plotting substation coordinates on a map, which says little about how the system is actually connected. Modelling the grid this way also brings together two fields: the engineering view of how power grids are operated and optimised, and network science, the study of complex networks. Network science has been used to analyse many different systems, from social ties to biological interactions to the spread of epidemics, and has produced general insights that apply across all of them, such as the small-world property [WS98] and the way load spreads to the rest of a network when a heavily used element fails [CLM04]. Because these patterns recur across so many systems, the field offers a common language for describing and comparing them, power grids included [AJB00].

Within this network view, topological measures describe how the grid is organised. The most basic is the degree of a node: the number of transmission lines that meet at a substation, so a high-degree substation acts as a local hub. A second, complementary measure is betweenness

centrality, which captures how often a substation lies on the shortest paths between other parts of the grid; substations with high betweenness carry much of the grid's connectivity, so their loss can cut many indirect routes at once [PA13]. These measures, together with the others used in this thesis, are defined formally in Chapter 2. What matters here is that they make it possible to judge how important each component is from the structure of the network alone, without modelling the electrical physics of the grid. Following most topological studies, the high-voltage network is treated as an undirected, unweighted graph, so the analysis stays focused on this structure [PA13].

The vulnerability and robustness of electricity transmission grids have been widely studied using network-based approaches. In this context, vulnerability refers to the sensitivity of the grid to failures of specific components, while robustness describes the ability of the network to remain connected after such failures. Studies on real power grids show that the removal of certain critical nodes can lead to a substantial loss of connectivity, revealing a weakness in the network [AAN04]. Similar analyses of the European power grid show that targeted failures can significantly reduce network performance, which highlights the importance of identifying vulnerable components [RCVS07].

Even though network models provide useful insights into the structure of electricity transmission grids, they are a simplified abstraction of the physical grid. Topological analyses ignore electrical properties, like power flows and capacities, which limit their ability to capture all operational aspects of the grid [HCSB10]. Consequently, results obtained at the network level must be carefully interpreted and linked back to the physical grid. By mapping structurally important nodes and connections to their real-world locations within the electricity transmission network, network-based analyses can provide insights into potential vulnerabilities and critical components [PA13].

A network analysis like this is only possible because the necessary data is now available. Earlier European grid studies relied on a few older datasets, such as the UCTE reference grid, that no longer reflect the current network [RCVS07, SRCCMV08]. This thesis instead uses the PyPSA-Eur model, a recent open dataset of the European transmission system that is described in Section 4 [HHSB18]. In this dataset the grid is represented at the level of high-voltage substations and the transmission lines between them at 220 kV and above, not at the level of individual households or local distribution. This sets the scale, or granularity, at which structural vulnerability is studied here (Figure 1).

In summary, the importance of the European transmission grid is clear, it remains difficult to determine where the grid is most structurally vulnerable. Failures in critical infrastructure can have severe consequences, yet those responsible for planning and protecting the grid often have to act without a clear picture of where its weaknesses lie. In particular, there is limited insight into how vulnerability differs between countries, which components of the grid are most critical to its overall connectivity, and where those components are located. Network analysis offers a way to study these questions, but its results are often abstract and difficult to place in the geographic context in which grid decisions are made. The problem this thesis addresses is therefore the absence of a structural assessment that identifies the most vulnerable regions of the European grid and locates them geographically.

1.2 Research Question

The European electricity transmission grid is a complex system in which failures can affect large regions. Existing network studies tend to treat its structure, its response to failures, and its geography separately: most either characterise structure or simulate failures on an abstract network without mapping the result back to the territory, and several influential studies rely on grid data

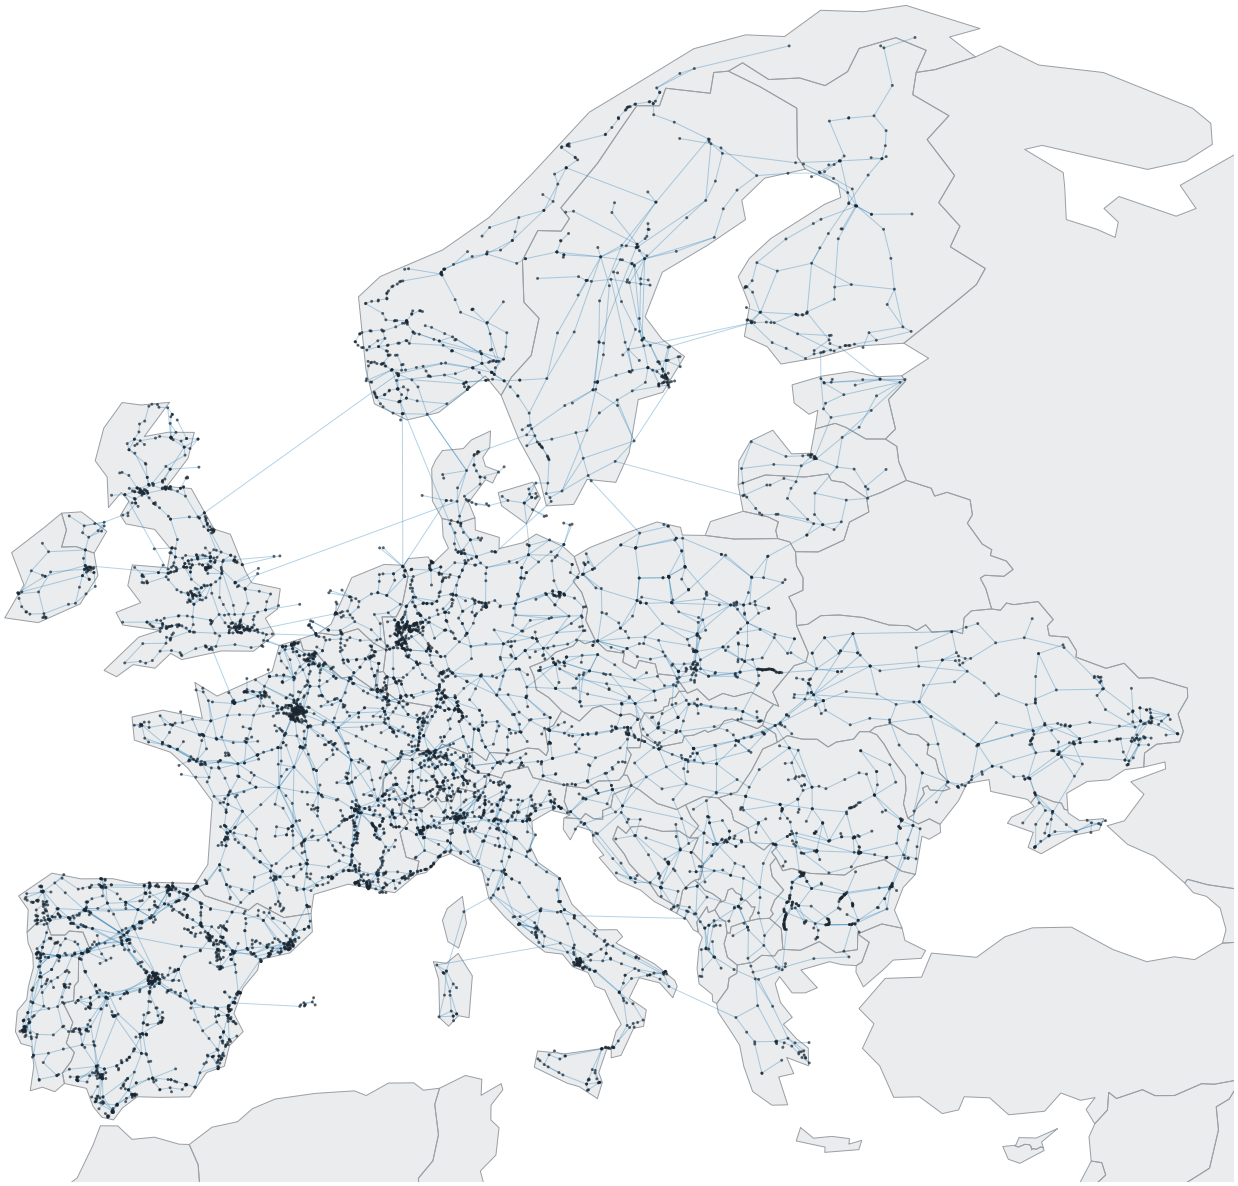

Figure 1: The European high-voltage transmission grid modelled as a network. Nodes are substations and power plants; edges are transmission lines at 220 kV and above

that is now more than a decade old. Three choices distinguish the analysis in this thesis. First, it uses a recent, openly available dataset, the most up-to-date model of the European grid, whose pan-European scope allows a consistent comparison across countries. Second, rather than relying on static centrality values, it uses those values to guide a node-removal simulation, so that vulnerability is judged from how the network actually fragments under attack rather than from centrality scores alone. Third, it does not stop at the abstract network but maps the results back onto geography, comparing the European grid as a whole with individual countries to locate where vulnerability is concentrated. Together these choices lead to the following research question:

How can the European power transmission grid be modelled as a network to identify structurally vulnerable regions?

Subquestions

1. What are the structural and topological characteristics of the European power transmission grid when modelled as a network?
2. How does the European transmission network fragment under the simulated removal of nodes, and how does this differ between random failures and targeted attacks?
3. Which regions of the European electricity transmission grid emerge as most structurally vulnerable?

1.3 Thesis Structure

The remainder of this thesis is organised as follows. Chapter 2 reviews complex network theory and the existing literature on power-grid vulnerability, and further identifies the research gap this thesis addresses. Chapter 4 describes the dataset, the criteria used to select countries for comparison, the network model, the three node-removal strategies, and the tools used to implement the analysis. Chapter 5 presents the results of the network analysis. Sections 5.2 through 5.4 address Sub-question 1 by reporting the topological characteristics, centrality distributions, and path-length properties of the European grid and its 25 country subgraphs. Section 5.5 addresses Sub-question 2 through robustness simulations under three node-removal strategies. Section 5.6 addresses Sub-question 3 by mapping the vulnerability findings back onto the geographic system. Chapter 6 interprets the findings in relation to the three subquestions, discusses the implications for infrastructure resilience, and identifies the limitations of the topological approach. Chapter 7 summarises the main conclusions and outlines directions for future research.

2 Preliminaries

This chapter provides the foundation for the analysis in the remainder of the thesis. Section 2.1 introduces complex network theory, the framework within which the transmission grid is studied. Sections 2.2 to 2.5 define the structural measures used throughout the thesis: node degree, betweenness centrality, the clustering coefficient and the average path length. Section 2.6 completes the framework with the connectivity concepts of components, bridges and degree assortativity.

2.1 Complex Network Theory

Complex network theory provides a formal framework for analysing the structural properties of networks. A system is reduced to a graph $G = (V, E)$, where V is the set of nodes and E is the set of edges; throughout this thesis, $n = |V|$ denotes the number of nodes, $m = |E|$ the number of edges, and angle brackets $\langle \cdot \rangle$ an average over all nodes. A *path* is a sequence of edges that connects two nodes, and a *shortest path* between two nodes is a path that uses the fewest edges. The structure of the graph is then described by a set of quantitative measures.

2.2 Node Degree

The degree k_i of a node $i \in V$ is the number of edges incident to it:

$$k_i = |\{j \in V : \{i, j\} \in E\}|. \tag{1}$$

Here, a high degree means that a node is directly connected to many neighbours. The degree distribution $P(k)$ describes the fraction of nodes in the network with degree k and characterises the general connectivity. In some networks a small number of nodes have a very high degree while most have only a few connections; when the degree distribution follows a power law in this way, the network is called scale-free.

2.3 Betweenness Centrality

Centrality measures rank how structurally important a node is to the network as a whole. Betweenness centrality measures the extent to which a node lies on paths between other pairs of nodes in the network. Betweenness can be defined over different sets of paths; this thesis uses the form in which only shortest paths are counted. Let σ_{st} be the total number of shortest paths from node s to node t, and let $\sigma_{st}(i)$ be the number of those shortest paths that pass through node i. The betweenness centrality b_i of node i is then defined as:

$$b_i = \sum_{\substack{s,t \in V \\ s \neq t,\ i \notin \{s,t\}}} \frac{\sigma_{st}(i)}{\sigma_{st}}. \tag{2}$$

Here, a high value indicates that a node is an important intermediary through which much of the network's traffic must pass. Throughout this thesis, betweenness is reported in normalised form: each b_i is divided by the number of node pairs not involving i, $(n-1)(n-2)/2$, so that $b_i \in [0, 1]$ and values are comparable across networks of different size. A value of $b_i = 1$ would mean every path in the network passes through node i.

The inequality of betweenness centrality across nodes can be summarised by the *betweenness Gini coefficient* $\mathrm{Gini}_b(G)$, the standard Gini inequality measure applied to the distribution of b_i values across all nodes:

$$\mathrm{Gini}_b(G) = \frac{\sum_{i=1}^{n} \sum_{j=1}^{n} |b_i - b_j|}{2n \sum_{i=1}^{n} b_i}. \tag{3}$$

Here, $\mathrm{Gini}_b(G) = 0$ indicates that betweenness is distributed evenly across all nodes and $\mathrm{Gini}_b(G) = 1$ indicates that a single node carries all shortest-path traffic.

2.4 Clustering Coefficient

The clustering coefficient measures the degree to which the neighbours of a node are also connected to each other, capturing local density in the network. The local clustering coefficient c_i of node i is defined as:

$$c_i = \frac{|\{\{i, j\} \in E : i, j \in N(i)\}|}{\binom{k_i}{2}}, \tag{4}$$

Here, $N(i)$ denotes the set of neighbours of node i and $\binom{k_i}{2} = k_i(k_i-1)/2$ is the maximum number of edges that could exist among them. The global clustering coefficient $\langle C \rangle$ of the network is the mean of all local values:

$$\langle C \rangle = \frac{1}{n} \sum_{i \in V} c_i. \tag{5}$$

Here, a high clustering coefficient is an indicator of the network containing tightly connected local groups. Watts and Strogatz introduced this metric to separate real-world networks from solely random graphs, showing that many real networks exhibit clustering coefficients far above those of a random graph with the same number of nodes and edges [WS98].

2.5 Average Path Length

The distance $d(i, j)$ between two nodes i and j is the number of edges on the shortest path connecting them. The average shortest path length $\langle \ell \rangle$ of the network is defined as:

$$\langle \ell \rangle = \frac{1}{n(n-1)} \sum_{\substack{i,j \in V \\ i \neq j}} d(i, j). \tag{6}$$

Here, a low average path length indicates that any two nodes can be reached through a few steps in between. Together with the clustering coefficient, average path length forms the basis for looking at small-world behaviour [WS98]. A network is considered to be a small-world network when its clustering coefficient is much higher than that of an equivalent random graph ($\langle C \rangle \gg \langle C \rangle_{\text{random}}$), while its average path length stays comparable ($\langle \ell \rangle \approx \langle \ell \rangle_{\text{random}}$). Because larger networks naturally have longer paths, comparisons across networks of different size use the *size-normalised path length* $\langle \ell \rangle / \sqrt{n}$. The $\sqrt{n}$ reference comes from geometry: a network laid out in two dimensions spans roughly $\sqrt{n}$ nodes from side to side, so crossing it takes on the order of $\sqrt{n}$ steps. The ratio $\langle \ell \rangle / \sqrt{n}$ expresses path lengths relative to what a grid-like network of the same size would show.

2.6 Components, Bridges and Assortativity

A set of nodes in which every node can be reached from every other node by some path forms a *connected component*. A graph may split into several disconnected components; the largest of these, the one containing the most nodes, is the *largest connected component* (LCC), used in Section 4.4 to measure robustness.

An edge whose removal increases the number of connected components is called a *bridge*. A bridge is a single point of failure: it is the only connection between the two parts of the graph it joins, so deleting it splits the network in two. A graph with many bridges is locally tree-like and fragile, whereas a graph whose nodes are linked by several independent paths has few bridges. Removing a bridge disconnects the graph abruptly, in contrast to the gradual fragmentation of the largest connected component that occurs when nodes are removed one by one.

Degree assortativity measures whether nodes tend to connect to other nodes of similar degree. It is the Pearson correlation coefficient $r \in [-1, 1]$ between the degrees of the two endpoints of an edge, taken over all edges. A positive value (assortative) means high-degree nodes tend to attach to other high-degree nodes; a negative value (disassortative) means high-degree nodes tend to attach to low-degree nodes; and a value near zero indicates no degree correlation.

3 Related Work

This section 3 reviews existing work on the topological analysis of power grids: the foundational node-removal studies, more recent European and regional analyses, and the work that extends or challenges the purely topological view. Section 3.4 draws these together into the research gap that this thesis addresses. A recurring theme is how far structural network measures can characterise the vulnerability of a real transmission system due to abstracting away from electrical physics.

Complex network science was first applied to power grid vulnerability analysis through a series of foundational studies that introduced the node removal framework and found recurring structural properties across power grids. Two categories of node removal are used throughout this literature and form the base for the simulations in this thesis. The first is fault tolerance, which tests the network's resilience under random node removal, representing accidental failures such as natural disasters. The second one is attack tolerance, which concerns resilience under targeted removal of structurally important nodes [AAN04].

3.1 Topological Vulnerability of Power Grids

The theoretical foundation for fault and attack tolerance in networks was established by Albert, Jeong, and Barabási [AJB00], who show that scale-free networks, in which degree follows a power law, are highly robust to random node failure but collapse rapidly when hubs are targeted. Because hubs are rare, random removal is unlikely to hit them; targeted removal seeks them out deliberately. Networks with more uniform degree distributions do not show this asymmetry: random and targeted removal produce similar fragmentation curves. Power grids fall between these two cases because their degree distributions are neither scale-free nor uniform. This raises the question of how much the fault attack asymmetry applies to them.

Albert et al. [AAN04] conduct one of the first structural vulnerability studies of a real-world power grid, analysing the North American transmission network by comparing three node-removal strategies: random failure, degree-based targeted attack, and load-based targeted attack, where node load is based on betweenness centrality. Connectivity loss is measured as the fraction of distribution substations that lose their connection to the network. Their main finding is that generating substations are highly redundant under all strategies, while transmission substations represent the true structural vulnerability: targeted removal of just 4% of the highest-load transmission hubs causes approximately 60% connectivity loss. Random failure, by contrast, produces gradual and proportional degradation in networks with an exponential degree distribution. This is due to the fact that only a small number of nodes have a high degree. A randomly removed node is therefore unlikely to be structural [AAN04]. The difference between fault tolerance and attack tolerance becomes one of the main results of the power grid vulnerability literature.

Building on this framework, Crucitti et al. apply a cascading failure model [CLM04] to the Italian high-voltage transmission grid, a network of 341 nodes and 517 edges at 220 kV and 380 kV. Node load is defined as betweenness centrality and node capacity as a fixed multiple of its initial load; removing a high-load node causes remaining nodes to redistribute load to neighbouring nodes, which could trigger further overloads in a cascade. The Italian grid exhibits an exponential degree distribution and the same threshold as the North American case: removal of 4% of the highest-load nodes causes a 60% drop in global network efficiency. The convergence of this threshold across two structurally different grids on different continents and different sizes suggests it may reflect

a general property of power grid networks. Together, these two studies establish the exponential degree distribution, the fault-attack difference, and the 4% threshold as the reference point against which later European grid studies are evaluated.

3.2 Network Studies of the European Grid

The first application of the node-removal framework to the European grid was made by Rosas-Casals et al. [RCVS07], who analyse the UCTE interconnected European transmission system. In their analysis 3,000 nodes and 200,000 km of transmission lines are covered, both as a whole and at the individual country level. Their main methodological contribution is to summarise each grid by a single parameter describing how quickly its degree distribution decays, and to use a threshold on this parameter to classify grids as robust or fragile under targeted node removal. Applying this methodology country by country, they find that European grids split into two structural groups, with differences in fragility due to the historical, geographical and political circumstances of each country's grid development. All European grids have small-world characteristics relative to equivalent random graphs, which is consistent with findings on the North American and Italian grids [AAN04, CLM04].

Solé et al. [SRCCMV08] extend this analysis using percolation theory. Percolation theory describes the conditions under which a network loses its global connectivity as nodes are progressively removed. This theory is used to derive a theoretical prediction of the fraction of nodes whose removal causes the largest connected component to collapse. The observed robustness curves do not always match these theoretical predictions exactly. The countries where the largest deviations occur tend to also perform worse on real-world reliability metrics such as energy not supplied. This suggests that the degree distribution captures meaningful differences between countries in how their grids behave in the real world. Solé et al. also point out that grid operators plan for the failure of any single component at a time, the n−1 criterion, but not for the simultaneous or sequential loss of multiple critical nodes. The fast collapse of the LCC under targeted and progressive node removal suggests that this planning standard can leave the grid exposed to deliberate attacks.

Rosas-Casals and Corominas-Murtra [RCCM09] build on Rosas-Casals et al. [RCVS07] and Solé et al. [SRCCMV08] testing whether this topological classification actually corresponds to differences in real-world grid reliability, using operational data from the UCTE. Grids classified as robust, such as Belgium, the Netherlands, Germany, Italy, Romania, and Greece, show lower values of energy not supplied and shorter interruption lengths in real operational records. Grids classified as fragile, such as France, Hungary, and Spain, show higher interruption metrics. This positive correlation between a topological measure and real-world reliability provides research-based support for the use of structural network metrics in assessing the vulnerability of power grids.

More recent studies extend the European analysis using updated datasets and larger country samples. Espejo et al. [ELR18] analyse 15 European transmission networks at 220 kV and 400 kV using a homogeneous ENTSO-E dataset, which enables direct cross country comparison that earlier studies could not support due to heterogeneous national data. They confirm exponential degree distributions across all 15 countries and find that all networks exhibit small-world properties, but only when both voltage levels are analysed together. When a single voltage level is analysed in isolation, small-world properties go away. This suggests that the topology responsible for small-world properties emerges from combining transmission layers rather than an isolated layer. Espejo et al. also report that all 15 grids are disassortative, meaning that high-degree nodes often connect to

low-degree nodes. This structural property has direct implications for how targeted attacks on hubs propagate through the network.

Hartmann and Sugár [HS21] explain why small-world properties only emerge when multiple voltage levels are included, through a 70 year longitudinal study of the Hungarian power grid from 1949 to 2019. By tracking structural properties year by year, they show that small-world behaviour emerges when higher-voltage transmission lines are introduced. These long-distance lines connect substations that are geographically far apart. Therefore they dramatically reduce average path lengths while the local clustering coefficient, which is already high from the dense regional grid, remains almost unchanged. This combination of high clustering and short average path length is the condition for defining small-world behaviour introduced by [WS98]. The degree distribution never follows a power law at any point in the grid's development. This means that the Hungarian grid does not develop a small number of dominant high-degree hubs. Most nodes consistently have only a few connections, and the number of highly connected nodes drops off rapidly. This is consistent with the findings of Rosas-Casals et al. [RCVS07] on the European grid.

When looking at the regional level, Forsberg et al. [FTB23] analyse the Nordic transmission grid using degree, betweenness, and closeness centrality. Their analysis shows that the highest-betweenness nodes are geographically concentrated in south-eastern Norway and central Sweden. These are the nodes through which most shortest paths in the network run, making them structural bottlenecks. Failure of these nodes would therefore disrupt connectivity across the entire Nordic grid. By combining all three centrality measures into a single ranking, they identify around a dozen nodes that stand out as most critical, located primarily in central Sweden and the Stockholm region. They also find that analysing the network with and without edge weights, where weights represent line capacity, leads to different conclusions about which nodes are most vulnerable. This shows that line capacity is an important factor when identifying critical nodes, and an unweighted analysis alone might miss part of the full picture. The study shows that centrality-based analysis can identify specific, geographically meaningful locations in the transmission network, which is a direct motivator for the approach taken in this thesis.

3.3 Beyond Pure Topology: Power Flow and Cascades

While the studies reviewed above demonstrate that topological analysis can identify structurally critical nodes and characterise grid vulnerability, other research raises important questions about how well these structural findings are an indicator of the real physical behaviour. Hines et al. [HCSB10] test this by comparing topological vulnerability rankings against blackout sizes modelled using DC power flow simulations. They find that targeted attacks cause more damage than random failures regardless of which measure is used. This confirms the difference of impact between fault-tolerance and deliberate-attack strategies identified in earlier work. However, the topological and power-flow measures differ substantially in the nodes they identify as most critical and the correlation between topological rankings and actual blackout size is mild. Cuadra et al. [CSSDS+15] survey the literature and reach a similar conclusion. They find that purely topological analysis is widely used and computationally tractable, but it does not capture the electrical properties.

In response to these limitations, Bompard et al. [BNX09, BPD12] extend the topological metrics by incorporating DC power flow into the network analysis. Rather than treating all shortest paths as equally likely routes of power, their approach weights paths according to how power distributes across transmission lines in the real world. This produces an electrical version of betweenness

centrality and an efficiency metric called net-ability, which includes line capacity limits. When this method is applied to the European grid, their extended metrics identify different critical nodes and lines than purely topological metrics. This suggests that standard betweenness centrality can put vulnerable components in the incorrect ranking when the power flow deviates from the shortest path.

Other research addresses the dynamic nature of real grid failures. Wenli et al. [WZPS16] propose a cascade model in which node load is based on actual power flow distribution rather than betweenness centrality. They also look at situations where circuit breakers on neighbouring lines trip incorrectly in response to a nearby overload, also called hidden failures. They model how these hidden failures disconnect lines that were not themselves overloaded. This adds a second version of failures on top of the initial overload cascade, making total collapse more severe. Their results show that the nodes most likely to trigger large cascades are not necessarily the highest-degree or highest betweenness nodes, which limits how much basic centrality rankings can say about real cascade risk.

Schäfer et al. [SWTL18] show that power spikes in the seconds immediately after a line failure can overload neighbouring lines before the system reaches a new steady state. This is a dynamic cascade that static topological models are not able to capture. At the line level, Maity and Panigrahi [MP24] extend the framework of centrality from nodes to edges by defining line centralities based on effective electrical resistance. They identify transmission lines whose removal causes the largest drop in network efficiency.

Despite these extensions with electrical properties, purely topological analysis remains the most widely applied approach in large-scale power grid research [CSSDS+15, WKK+23], because it requires only network topology data and is computationally possible for European-scale networks of thousands of nodes. This thesis adopts the topological approach for the same reasons, while fully acknowledging the limitations that the studies above identify.

3.4 Research Gap

The studies reviewed above each address part of the problem this thesis takes on, but none combines them. The foundational node-removal studies establish the fault–attack asymmetry on individual grids [AAN04, CLM04], and the European studies that follow rely on the UCTE dataset from 2003 [RCVS07, SRCCMV08, RCCM09], which predates major changes to the European grid such as the expansion of renewable generation. More recent work brings the data up to date but treats the aspects separately: Espejo et al. [ELR18] enable cross-country comparison but characterise structure without simulating failures, and the most comparable study, Hartmann and Cirunay [HC26], compares European countries using structural metrics and a single fragmentation curve, but applies only one removal strategy and does not map its results back to geography. A separate line of work extends the purely topological approach with electrical power flow and cascade dynamics [BNX09, BPD12, WZPS16, SWTL18], which captures effects that topology misses but requires detailed electrical data and does not scale to a network of European size [CSSDS+15].

Against this background, three aspects distinguish the analysis in this thesis. First, it uses the PyPSA-Eur dataset [HHSB18], the most recent openly available model of the European transmission system, whose pan-European scope supports a consistent comparison across 25 countries that the heterogeneous national datasets of earlier studies could not. Second, rather than reporting static centrality values, it uses those values to drive a multi-strategy node-removal simulation, random, degree-based, and betweenness-based, and measures vulnerability from the resulting fragmentation

curves, so that centrality serves as an informed indicator of which nodes to remove rather than as the end result. Third, it does not stop at the abstract network but maps the vulnerability results back onto geography, comparing the European grid as a whole with individual countries to locate where structural vulnerability is concentrated. The analysis remains topological, which keeps it scalable and interpretable at European scale, while the limitations of abstracting away from electrical physics are acknowledged throughout [HCSB10]. Together, these three choices provide a structural assessment of the European power grid that extends and integrates what previous studies have addressed separately.

4 Data and Methodology

This chapter describes the steps that lead from the raw grid data to the vulnerability scores. Section 4.1 introduces the PyPSA-Eur dataset and the filtering applied to it. Section 4.2 sets out the criteria used to select the countries for the cross-country comparison. Section 4.3 defines how the data is represented as a graph. Section 4.4 describes the node-removal simulation and the three attack strategies, and Section 4.6 defines the vulnerability metric derived from the resulting fragmentation curves. Section 4.7 details the implementation.

4.1 Dataset

The analysis draws on the PyPSA-Eur dataset [HHSB18, BHS18], an open-source model of the European high-voltage transmission system derived from OpenStreetMap infrastructure data. PyPSA-Eur represents the transmission network as five component tables: buses (substations and converter stations), lines (alternating-current transmission lines), links (high-voltage direct-current connections), transformers (voltage-level couplings) and converters (AC/DC conversion).

The dataset was filtered to retain only extra-high voltage infrastructure at 220 kV and above, consistent with the transmission backbone studied in prior robustness analyses of European power grids [SRCCMV08, HC26]. After applying this voltage filter, the network comprises 6795 nodes and 8781 edges spanning 36 countries.

Each node retains its geographic coordinates, country code and nominal voltage level as attributes from the dataset. These attributes play no role in the topological analysis itself, but are used to locate and describe nodes when the results are interpreted.

The dataset distinguishes three types of edge: alternating-current transmission lines (7865 edges), transformers that couple different voltage levels within a substation (878 edges), and high-voltage direct-current (HVDC) links (38 edges), together accounting for the 8781 edges above. These types differ in their electrical role, but because the analysis in this thesis is purely topological it does not return to this distinction: all three are treated identically, as a single edge representing a physical connection between two substations. The graph that this produces is defined in Section 4.3.

4.2 Country Selection

For the country-level comparison, a country is included if the largest connected component of its national subgraph contains at least 39 nodes. This threshold ensures that the node-removal range used throughout this thesis (Section 4.4) spans enough removal steps for the fragmentation

curve to be meaningful, while keeping the geographic scope as wide as possible. Applying it yields 25 countries: France, Germany, Spain, Italy, Great Britain, Norway, Ukraine, Poland, Sweden, Switzerland, Romania, Bulgaria, Portugal, Austria, Finland, Belgium, Czechia, the Netherlands, Hungary, Ireland, Serbia, Denmark, Slovakia, Greece, and Bosnia and Herzegovina.

Because PyPSA-Eur is derived from OpenStreetMap, the completeness of its coverage varies from country to country. To make this variation explicit rather than leave it hidden, each country is assigned a coverage ratio, defined as the total circuit-kilometres of its transmission lines in the dataset divided by the circuit-kilometres reported by its national transmission system operator (TSO), compiled from TSO publications of differing reference years. Countries are then grouped into three coverage tiers. These tiers are not taken from prior work; they are defined here for this thesis as Tier A (ratio ≥ 0.65, good coverage), Tier B ($0.35 \leq$ ratio < 0.65, partial coverage) and Tier C (ratio < 0.35, sparse coverage). The cutoffs separate countries whose dataset reproduces most of the reported grid from those where a substantial share is missing, with Tier B as an intermediate band.

The thresholds are deliberately soft rather than hard decision rules. The Tier B/C boundary at 0.35 falls inside a wide gap in the data, between the lowest Tier B country at 0.45 and the highest Tier C country at 0.26, so the classification is insensitive to its exact value. The Tier A/B boundary at 0.65 is more arbitrary, as several countries cluster just below it; shifting it to 0.60 or 0.70 would reclassify a few borderline cases such as Slovakia (0.64) without altering the overall pattern. Coverage is used for interpretation, not exclusion. No country is removed on the basis of its tier, and only the 39-node rule above governs inclusion. As a result, a few low-coverage countries are retained, most notably the Netherlands, for which the dataset captures only about 14% of the grid length reported by its TSO (Tier C); its results are kept but flagged, since so large a share of the national network is absent that structural conclusions drawn from it remain uncertain. To keep this transparent, the relationship between coverage and vulnerability is examined directly in Section 5.5, so that results from low-coverage countries can be read in the appropriate context. (Great Britain and Ukraine are included by the size rule but lack comparable TSO reference statistics, and are therefore left untiered.) With the country scope and its data-quality caveats established, the following section defines how each network is represented as a graph.

4.3 Network Model

The transmission network is modelled as an undirected, unweighted graph as defined in Section 2.2. Each node corresponds to a high-voltage substation and each edge corresponds to a physical connection between two substations.

The graph is undirected because power in a transmission network flows in either direction depending on load conditions, so edge direction provides no relevant information for vulnerability analysis. The graph is unweighted because the analysis aims at the topological structure rather than electrical properties, like line impedance or thermal capacity. This is consistent with the topological approach used throughout the power grid vulnerability literature [AAN04, SRCCMV08, HC26].

HVDC links connect substations in electrically separate synchronous zones that share no alternating-current connection. For the topological analysis, HVDC links are treated as normal edges, such that the full pan-European network can be analysed as a single graph. This allows cross-zone connections to be captured in structural metrics such as betweenness centrality, which would otherwise be confined to one single synchronous zone. For country-level analyses, subgraphs are constructed by keeping only nodes and edges belonging to the country in question. This means

that cross-border edges connecting nodes in different countries are excluded.

4.4 Network Attack Simulation

The robustness of a network is measured by simulating the progressive removal of its nodes. Taking the graph G of Section 4.3 as input, the simulation removes nodes one at a time according to a chosen strategy and, after each removal, records how much of the network remains connected. The output is a fragmentation curve; the single vulnerability score derived from that curve is defined separately in Section 4.6.

The connected fraction is measured by the relative size of the largest connected component (LCC), defined in Section 2.6. Let $|\mathrm{LCC}(G, \varphi)|$ be the number of nodes in the largest connected component of G after a fraction φ of its nodes has been removed. The relative LCC size is

$$S(G, \varphi) = \frac{|\mathrm{LCC}(G, \varphi)|}{n}, \tag{7}$$

Here, $\varphi \in [0, 1]$ is the fraction removed. Thus $S(G, 0) = 1$ for a fully connected network and $S(G, \varphi) \approx 0$ once it has almost completely fragmented. The fragmentation curve is obtained by recording $S(G, \varphi)$ after every removal step, giving the relative LCC size as a function of the removed fraction φ. This curve is the primary robustness measure used in this thesis: a network is robust if S decreases gradually as φ increases, and fragile if S collapses rapidly at small φ. Because S is normalised by n, curves can be compared directly across networks of different sizes [AAN04].

4.5 Random Removal

Three removal strategies are used, corresponding to two kinds of threat. Fault tolerance is assessed with *random* removal, in which nodes are selected uniformly at random, simulating accidental failures such as those caused by natural disasters. Attack tolerance is assessed with two *targeted* strategies, degree-based and betweenness-based removal, which represent two different attackers [AAN04]. The degree-based attacker has only local information and prefers visibly well-connected substations; the betweenness-based attacker has full knowledge of the network and can identify its most critical nodes exactly. The two therefore bound a realistic range: degree-based removal is the conservative case, a fast attack using local information, and betweenness-based removal the worst case, a slower attack using complete information. Because random removal depends on the order in which nodes are drawn, it is repeated 100 times and the average curve is reported.

4.5.1 Degree-based removal

At each step a node is selected for removal with probability proportional to its current degree, so that a node of degree k is k times as likely to be chosen as a node of degree 1; after each removal the degrees of the affected neighbours are updated. Selecting nodes probabilistically, rather than always taking the single highest-degree node, reflects an attacker who can see which substations carry many connections but does not analyse the network as a whole. It also removes the need for a tie-breaking rule: degree ties are common in networks with exponential degree distributions [RCVS07] and are resolved naturally by the sampling weights. Because the selection is stochastic, the simulation is repeated five times and the mean fragmentation curve is reported, with shaded bands showing the standard deviation across runs.

4.5.2 Betweenness-based removal

At each step the node with the highest betweenness centrality is removed, representing an attacker with complete visibility of the network. Betweenness values are effectively continuous and rarely tied, so this deterministic rule is well defined and does not need averaging. Because betweenness can change substantially once a node is removed, the frequency at which it is recomputed affects both accuracy and computational cost. A preliminary experiment on the Ireland subgraph compares four recomputation frequencies: after every step, every two steps, every three steps, and once at the start. Recomputing after every step produces the fastest fragmentation at low φ and is therefore adopted for all betweenness-based simulations in this thesis, including the full European grid; the full results are shown in Figure 16 in Appendix A.

4.6 Network Vulnerability Metric

The gap between the random and targeted fragmentation curves measures how exposed a network is to a deliberate attack: if targeted removal fragments the network no faster than random failure, an attacker gains little from choosing nodes carefully, whereas a large gap means a few well-chosen removals are far more damaging than chance [AAN04]. To capture this in a single number, the vulnerability score R_{DEG} is defined as the difference in area under the fragmentation curve between random and degree-based targeted removal:

$$R_{\mathrm{DEG}} = \int_0^{\varphi_{\max}} S_{\mathrm{random}}(\varphi)\,\mathrm{d}\varphi - \int_0^{\varphi_{\max}} S_{\mathrm{degree}}(\varphi)\,\mathrm{d}\varphi. \tag{8}$$

Figure 2 illustrates this construction: each score is the area enclosed between the random curve and the corresponding targeted curve, so a larger enclosed area means a larger advantage for the attacker. The area under the curve is used rather than the connected fraction at a single threshold because it reflects the whole shape of the fragmentation process instead of its value at one point. A high R_{DEG} means that targeted removal fragments the network disproportionately faster than random failure, which signals a structure built around a small number of critical high-degree hubs and therefore high vulnerability. A low R_{DEG} means the two strategies fragment the network similarly, implying a more uniform structure that an attacker cannot exploit and therefore low vulnerability. Countries are ranked by R_{DEG}, and how it relates to other structural properties is examined in Section 5.5.

A second score, R_{BC}, applies the same definition to betweenness-based removal:

$$R_{\mathrm{BC}} = \int_0^{\varphi_{\max}} S_{\mathrm{random}}(\varphi)\,\mathrm{d}\varphi - \int_0^{\varphi_{\max}} S_{\mathrm{betweenness}}(\varphi)\,\mathrm{d}\varphi. \tag{9}$$

Here, R_{BC} measures exposure to an attacker who targets the network's most central nodes. Because betweenness-based removal is consistently the most destructive strategy, R_{BC} is generally larger than R_{DEG} and captures a distinct facet of structural vulnerability.

In their raw form the two scores occupy small and uneven numerical ranges (roughly 0 to 0.03 for R_{DEG} and 0 to 0.14 for R_{BC}), which makes them hard to compare. For the country comparison in Section 5.5, each score is therefore rescaled to $[0, 1]$ by dividing by its maximum across the 25

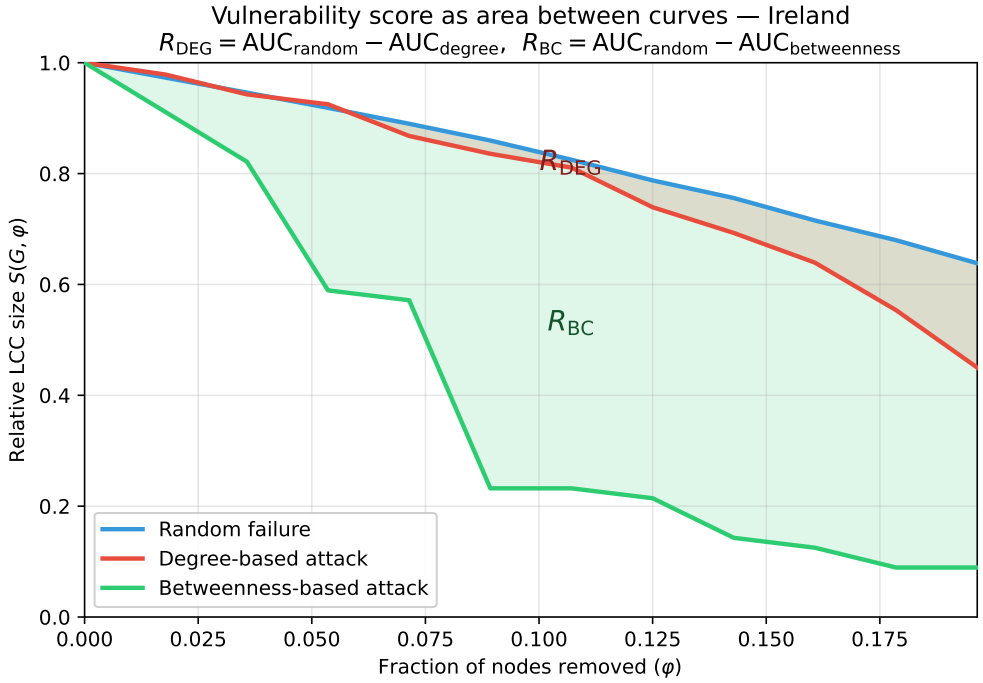

Figure 2: Construction of the vulnerability scores, illustrated on the Ireland subgraph. R_{DEG} (Equation 8) is the area between the random and degree-based fragmentation curves; R_{BC} (Equation 9) is the area between the random and betweenness-based curves.

countries:

$$R_{\mathrm{DEG,norm}}(c) = \frac{R_{\mathrm{DEG}}(c)}{\max_{c'} R_{\mathrm{DEG}}(c')}, \qquad R_{\mathrm{BC,norm}}(c) = \frac{R_{\mathrm{BC}}(c)}{\max_{c'} R_{\mathrm{BC}}(c')}. \tag{10}$$

On this scale the most vulnerable country scores 1.00 and the least vulnerable approaches 0. All later references to R_{DEG} and R_{BC} use the normalised form unless stated otherwise.

4.7 Implementation

The network is loaded from the PostGIS database into a NetworkX undirected graph object. Each row in `powergrid.nodes` becomes a node and each row in `powergrid.edges` becomes an edge in which edge direction is discarded. Where multiple edges exist between the same pair of nodes, only one edge is kept, consistent with the unweighted topological model of Section 4.3.

The robustness simulation follows the same loop for all three strategies. At each step, the next node to remove is selected based on the used strategy, the node and all its directly connected edges are removed, and the size of the largest connected component is recorded. This results in one (φ, S) data point per step. The strategies differ only in how the node is chosen. Random removal draws nodes uniformly at random; degree-based removal samples them with `random.choices`, using the current node degrees as weights and decrementing the degrees of the affected neighbours after each removal; and betweenness-based removal selects the node of highest betweenness centrality,

computed exactly with the Brandes algorithm as implemented in NetworkX and recomputed after every removal step. The repetition and averaging applied to the random and degree-based strategies, and the betweenness recomputation frequency, are those established in Section 4.4; all three strategies are run on every network, including the full European grid.

5 Network Analysis

This chapter presents the results that answer the three sub-questions of this thesis. Section 5.1 describes the setup of the experiments. Sections 5.2, 5.3 and 5.4 address Sub-question 1 by characterising the structural and topological properties of the grid: degree structure and bridges in Section 5.2, the concentration of shortest-path traffic in Section 5.3, and path-length behaviour in Section 5.4. Section 5.5 answers Sub-question 2 by simulating random and targeted node removals and measuring the resulting fragmentation. Section 5.6 answers Sub-question 3 by projecting the network-level vulnerability findings back onto the geographic map to locate the regions where structural weakness is concentrated.

5.1 Experimental Setup

All experiments are carried out at two levels: the full European transmission grid and one subgraph per country, each restricted to its domestic nodes and edges and reduced to its largest connected component (Section 4.2). Every simulation removes one node per step up to a maximum removal fraction of $\varphi_{\max} = 0.20$; this range covers the early removals in which targeted strategies produce the most fragmentation. Random removal is repeated 100 times and averaged. Degree-based removal is repeated five times, with the mean curve and one standard deviation across runs reported. Betweenness-based removal is deterministic and is recomputed after every step, a frequency fixed by a preliminary experiment on the Ireland subgraph (Figure 16 in Appendix A). All metrics and simulations follow the definitions of Chapter 4.

5.2 Network Topology

The extracted network forms a single connected component of $n = 6{,}795$ nodes and $m = 8{,}781$ edges spanning 36 countries, of which 25 are analysed individually. Table 1 reports the topological metrics of the full grid and of each country's largest connected component. The resulting network is sparse, locally tree-like. Its average degree is almost uniform across countries, but the properties beyond degree, the bridge fraction, degree assortativity, clustering, and path length, differ widely between them.

The first feature is sparsity. The full-grid density of 3.8×10^{-4} is far below the values typically reported for social or biological networks, and the average degree $\langle k \rangle = 2.58$ sits just above the minimum that any connected network can have: a connected graph on n nodes contains at least the $n - 1$ edges of a spanning tree, giving an average degree of at least $2(n-1)/n \approx 2$. With $m = 8{,}781$ edges against a spanning-tree minimum of 6,794, the grid carries only about 29% more edges than the sparsest structure that could hold it together. The maximum degree across the entire grid is only $k_{\max} = 11$, far short of what a scale-free network of this size would produce. Figure 3 shows the degree distribution: it is sharply peaked at $k = 2$ and decays quickly, with no heavy tail.

Table 1: Topological metrics for each country's LCC and for the full European grid: n = node count, m = edge count, $\langle k \rangle$ = average degree, $k_{\max}$ = maximum degree, $\langle C \rangle$ = average clustering, Br. = bridge fraction, r = degree assortativity, Int. = number of cross-border interconnectors touching the country (a full-grid property; these edges are excluded from the country subgraphs).

Country	n	m	$\langle k \rangle$	$k_{\max}$	$\langle C \rangle$	Br.	r	Int.
France	1,201	1,617	2.69	10	0.077	0.19	−0.11	33
Spain	1,075	1,317	2.45	9	0.022	0.33	+0.01	12
Germany	780	1,010	2.59	8	0.046	0.19	−0.01	40
Italy	590	752	2.55	8	0.047	0.18	+0.00	19
Great Britain	513	631	2.46	6	0.040	0.17	−0.10	9
Norway	288	334	2.32	6	0.016	0.24	+0.02	11
Ukraine	285	370	2.60	7	0.062	0.16	−0.12	14
Poland	246	338	2.75	9	0.065	0.17	+0.00	11
Sweden	209	261	2.50	8	0.057	0.27	−0.07	16
Switzerland	185	235	2.54	8	0.072	0.22	−0.02	28
Romania	169	201	2.38	6	0.005	0.28	−0.07	10
Bulgaria	153	181	2.37	7	0.031	0.23	+0.09	7
Portugal	153	214	2.80	8	0.101	0.22	−0.09	7
Austria	113	124	2.20	6	0.005	0.37	−0.01	17
Finland	105	127	2.42	8	0.020	0.15	−0.10	9
Belgium	71	81	2.28	6	0.053	0.33	−0.26	10
Czechia	67	90	2.69	10	0.049	0.23	−0.36	13
Netherlands	60	67	2.23	6	0.039	0.34	−0.44	9
Hungary	56	72	2.57	7	0.036	0.22	−0.37	17
Ireland	56	69	2.46	6	0.041	0.20	−0.12	2
Serbia	51	66	2.59	9	0.098	0.17	−0.15	13
Denmark	48	54	2.25	5	0.160	0.48	−0.09	9
Slovakia	47	55	2.34	6	0.043	0.25	−0.18	10
Greece	40	48	2.40	6	0.074	0.29	+0.02	6
Bosnia and Herzegovina	39	46	2.36	6	0.097	0.39	−0.18	13
Full EU grid	6,795	8,781	2.58	11	0.048	0.19	−0.03	215

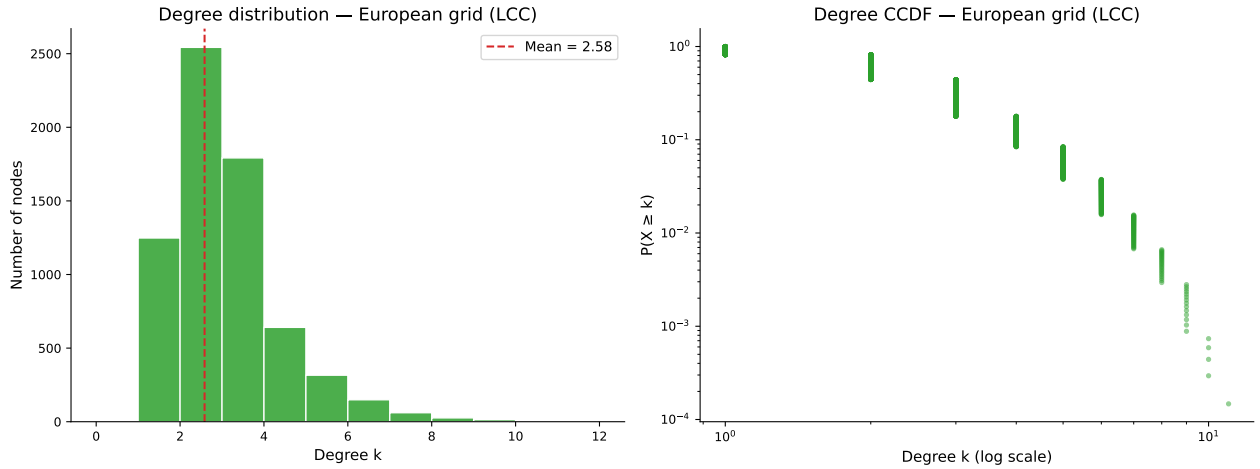

Figure 3: Degree distribution of the full European transmission grid LCC ($n = 6795$, $m = 8781$). Left: histogram of node degree k on linear axes. Right: complementary cumulative distribution function (CCDF) of k, with logarithmic horizontal and vertical axes.

The country-level distributions (Figure 17, Appendix A) share this low-degree, light-tailed shape, consistent with the finding that physical-infrastructure networks lack the hub-dominated topology of internet or social networks [RCVS07].

The second feature is local tree-likeness. The full-grid clustering coefficient is $\langle C \rangle = 0.048$, and 18.75% of all edges are bridges, single edges whose removal disconnects the network. Redundancy in the European grid is therefore structural rather than dense: alternative paths exist, but each individual node typically participates in only one or two of them. This scarcity of local redundancy is part of what, as we will see later, makes targeted removal so effective (Section 5.5).

Beyond the national grids, the full European network is tied together by a sparse set of cross-border interconnectors. Of the 8781 edges, 215 (2.4%) connect substations in different countries: 187 are so called alternating-current lines and 28 are HVDC links, the latter making up most of the 38 HVDC links in the dataset. These interconnectors are excluded from the country subgraphs by construction (Section 4.3) but retained in the full-grid analysis. The number of interconnectors per country varies widely, from 40 for Germany and 33 for France down to only 2 for the island grid of Ireland (Table 1).

The next feature is how the countries differ from one another. Density spans nearly two orders of magnitude, from 6.2×10^{-2} in Bosnia and Herzegovina ($n = 39$) to 2.2×10^{-3} in France ($n = 1201$), but this merely reflects the $n(n-1)$ scaling of the maximum edge count and says little about structure. Average degree is almost constant: across all 25 countries it stays within the narrow band 2.20–2.80. The structural variation lies elsewhere. Bridge fraction ranges widely, from 15% in Finland to 48% in Denmark, and degree assortativity ranges even more tellingly; Figure 4 shows both per country. Assortativity is the property that most clearly separates the national grids: it is close to zero or weakly negative for most countries, strongly negative in the heavily meshed grids of the Netherlands ($r = -0.44$), Hungary ($r = -0.37$) and Czechia ($r = -0.36$), and positive only in Bulgaria ($r = +0.09$) and Norway ($r = +0.02$). This matters because assortativity proves to be the single structural property that predicts how vulnerable a country is to a degree-based attack (Section 5.5): grids whose high-degree substations connect preferentially to other high-degree

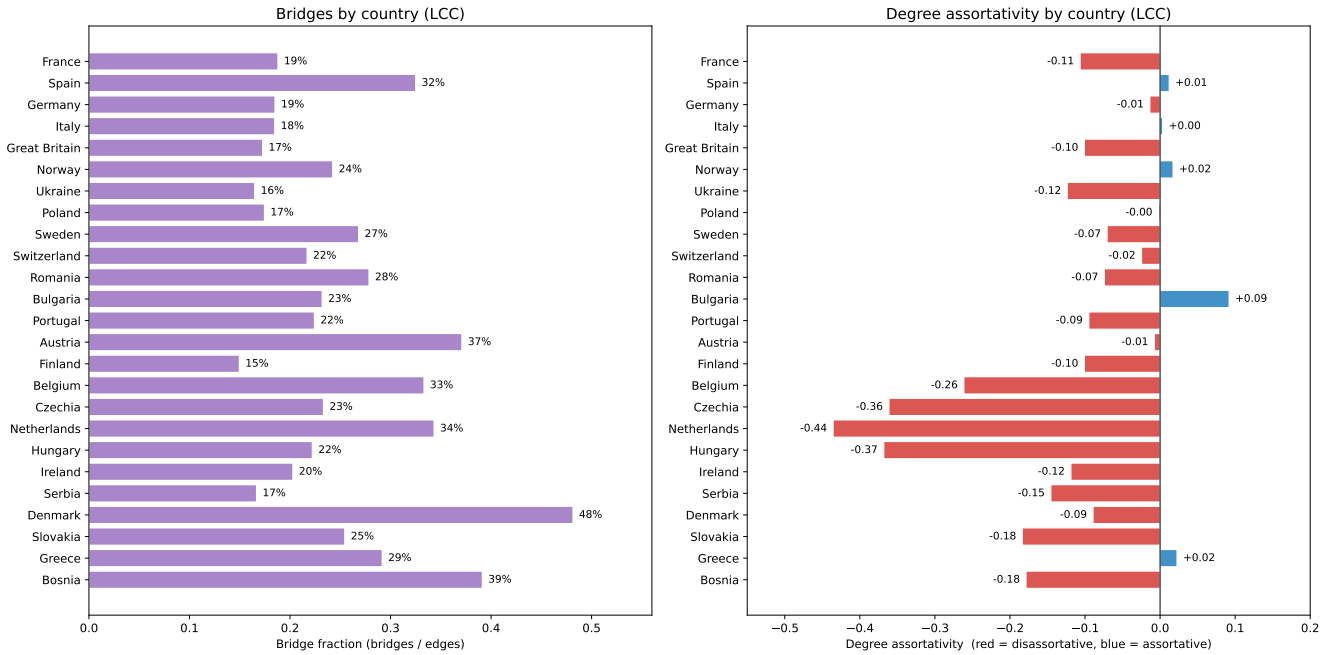

Figure 4: Country-level structural metrics for the largest connected component of each of the 25 countries, sorted from the largest to the smallest network. Left: bridge fraction (bridges divided by edges), with percentage labels. Right: degree assortativity r; red bars mark disassortative grids ($r < 0$), blue bars assortative grids ($r > 0$).

substations fragment fastest when those hubs are removed.

Norway stands apart on a second axis as well. Its modest node count combines with the longest country-level diameter ($D = 69$) and the largest size-normalised path length (1.15), values consistent with an elongated, chain-like topology that reappears in the robustness results of Section 5.5. Taken together, these results show that the European transmission grid is a sparse, low-degree, locally tree-like infrastructure network whose country-to-country differences arise in higher-order structure, assortativity, bridge density and path length, rather than in basic connectivity. This baseline topology frames the centrality and path analyses of Sections 5.3 and 5.4, which together complete the answer to Sub-question 1.

5.3 Centrality Analysis

Betweenness centrality measures how much of the network's shortest-path traffic passes through each node (Section 2.3). In the European grid this traffic is highly concentrated: the vast majority of substations carry almost no shortest-path load, while a small number act as critical intermediaries. Even the most central node has a normalised betweenness of only 0.29, and values fall steeply down the ranking: by rank 100 they have dropped below 0.06, and the median node carries almost no shortest-path traffic at all. No single node therefore dominates the grid in absolute terms. What separates the grid is the small group of nodes at the top of the ranking.

Table 2 lists the ten nodes with the highest betweenness in the full European grid. Eight lie in France and the other two just across the border in northern Spain; all ten operate at 400 kV. Their locations follow a corridor running from the Paris basin in the north-east, through central

and south-western France, down to the Spanish border. This confirms that shortest-path traffic across the pan-European network is routed through a small number of French substations, which together form the structural backbone of the continental grid.

Table 2: The ten nodes with the highest betweenness centrality in the full European grid. Coordinates are in decimal degrees (latitude, longitude); all ten substations operate at 400 kV. k = node degree, b = normalised betweenness centrality.

Rank	Country	Coordinates	k	b
1	France	(48.5, 3.9)	7	0.292
2	France	(48.4, 2.9)	8	0.203
3	France	(46.5, 1.6)	6	0.196
4	France	(45.1, −0.4)	7	0.193
5	France	(44.6, −0.6)	4	0.188
6	France	(44.0, −0.9)	3	0.187
7	France	(43.4, −1.4)	3	0.182
8	Spain	(43.2, −2.0)	3	0.180
9	Spain	(43.1, −2.3)	6	0.180
10	France	(47.1, 2.3)	6	0.178

Several of these nodes are barely connected: ranks five to eight have just three or four links each, yet they rank among the most central substations in Europe. Betweenness and degree therefore identify different critical nodes. A degree-based attacker, who targets the most-connected substations, would overlook these low-degree but high-betweenness bridges, whereas an attacker with full network knowledge would target them directly. This distinction is exactly what separates the two targeted-removal strategies of Section 5.5, and it shows why betweenness-based removal proves the more destructive of the two.

The same pattern appears within individual countries, though to varying degrees: in grids such as France, Spain and Germany most shortest-path traffic flows through just a handful of nodes, whereas in others it is spread far more evenly. This variation is quantified in Section 5.5, where the concentration of betweenness arises as the strongest structural predictor of a country's vulnerability to a betweenness-based attack. The path-length analysis in the following section completes the answers to Sub-question 1.

5.4 Shortest Path Analysis

Figure 5 reports, for each country, the average shortest path length $\langle \ell \rangle$ and its size-normalised form $\langle \ell \rangle / \sqrt{n}$ (Section 2.5). Path lengths are long: crossing the largest national grids takes on average 12.0 steps in France and 15.5 in Germany, with diameters reaching 43 in Italy. Judged against the two small-world conditions of Section 2.5, the grid meets the clustering condition but not the path condition. Its clustering is more than a hundred times that of a random graph of the same size ($\langle C \rangle \approx 0.048$, against $\langle k \rangle / n \approx 4 \times 10^{-4}$), the property earlier studies relied on when they described European grids as small-world [RCVS07, ELR18, HS21]. Its paths, however, are not short: crossing the full grid takes about 33 steps on average, nearly four times the ≈ 9 of a comparable random

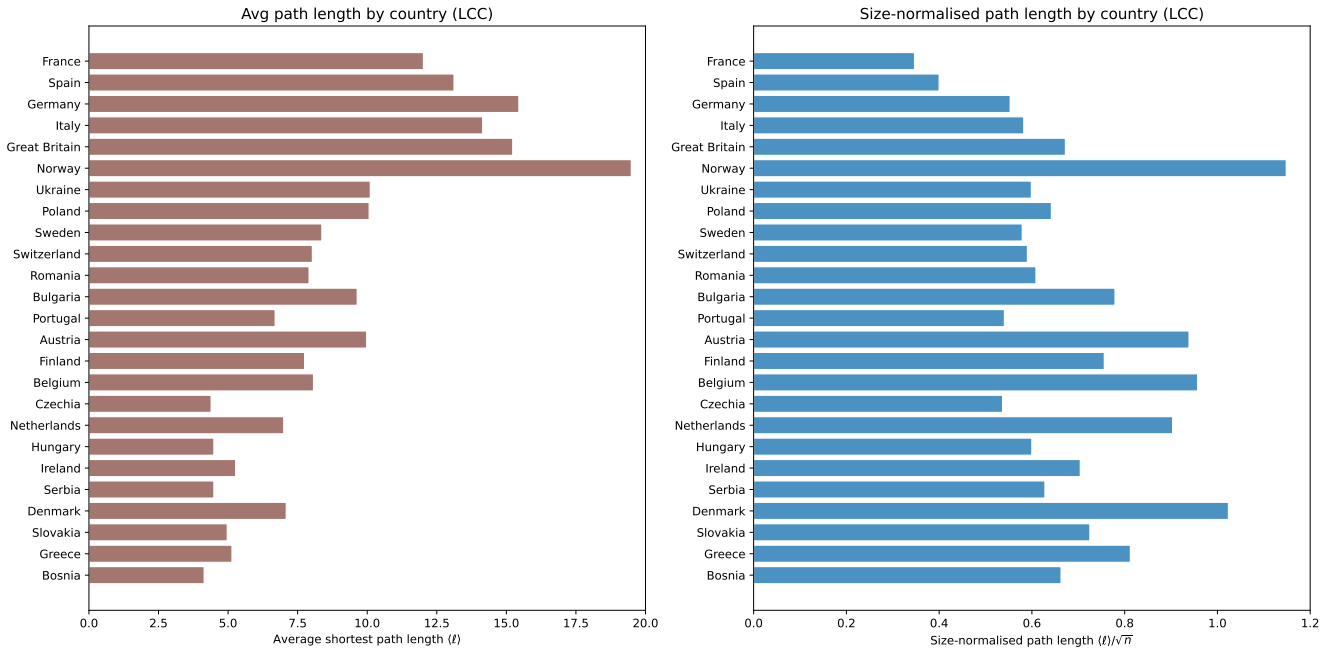

Figure 5: Path-length statistics for the largest connected component of each of the 25 countries. Left: average shortest path length $\langle \ell \rangle$. Right: size-normalised path length $\langle \ell \rangle / \sqrt{n}$.

graph. The grid is therefore not small-world in the strict sense, but a spatially embedded network in which geographic distance shows up as long routes through the grid.

The size-normalised ratio $\langle \ell \rangle / \sqrt{n}$ ranges from 0.35 in France to 1.15 in Norway. Norway is the clear outlier: despite having only 288 nodes it has both the longest size-normalised path length and a diameter of 69, meaning two substations can lie as much as 69 steps apart. This is the signature of the elongated, chain-like topology already visible in its bridge structure and assortativity (Section 5.2): the Norwegian grid stretches along the country's geography with few connections across its length. As the next section shows, this same structure makes Norway unusually fragile, fragmenting heavily even under random failure (Section 5.5).

Taken together, Sections 5.2 through 5.4 answer Sub-question 1. The European transmission grid is a sparse, low-degree, locally tree-like geographic network with long path lengths and shortest-path traffic concentrated in a small number of substations. It is not scale-free, and although its clustering is high, its long path lengths keep it from being small-world in the strict sense; its structural variation between countries lies not in basic connectivity but in higher-order structure: degree assortativity, bridge density and path length. These are the properties whose consequences for vulnerability will be examined in the remainder of this chapter.

5.5 Robustness under Simulated Failure

This section answers Sub-question 2 by simulating the progressive failure of the network and measuring how quickly it fragments. The three removal strategies of Section 4.4, random, degree-based and betweenness-based, are applied to the full European grid and to each of the 25 country subgraphs, and each country is summarised by the two vulnerability scores R_{DEG} and R_{BC} (Equations 8 and 9). The recomputation frequency used for the betweenness strategy was fixed by a preliminary

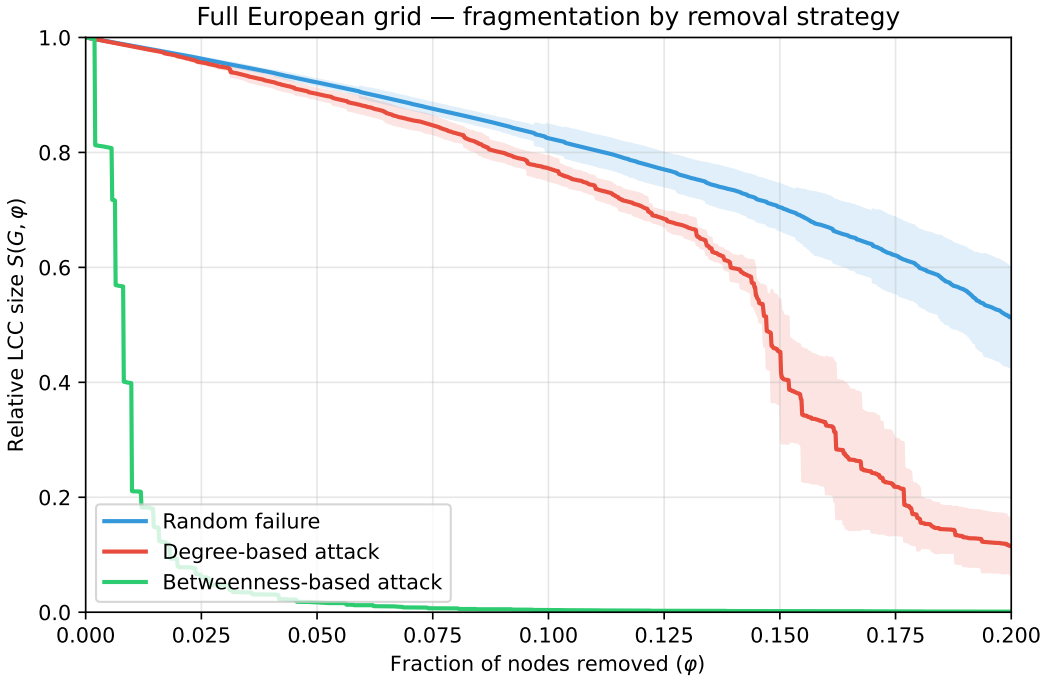

Figure 6: Fragmentation curves for the full European grid ($n = 6795$, $m = 8781$) under random, degree-based and betweenness-based node removal over $\varphi \in [0, 0.20]$. Shaded bands show one standard deviation across runs for the random and degree-based strategies.

experiment on the Ireland subgraph, reported in Appendix A.

Figure 6 shows the three fragmentation curves for the full European grid. The grid is highly fault-tolerant: under random failure the largest connected component still holds half its nodes after a fifth have been removed ($S = 0.51$ at $\varphi_{\max}$). Targeted removal is far more damaging, and the two strategies differ sharply in how. The betweenness-based attack is severe: removing barely one percent of nodes, the single most central substations, already breaks the grid into pieces ($S = 0.21$ at $\varphi = 0.01$), and by $\varphi = 0.05$ the network has essentially disintegrated ($S = 0.02$). The degree-based attack is slower and gentler in its early stages, tracking the random curve until around $\varphi = 0.13$, after which it collapses to $S = 0.12$. The gap between the curves is what the vulnerability score captures: the European grid withstands accidental failure well, but a knowledgeable attacker who targets its betweenness hubs can fragment it almost immediately.

Country-level comparison. The same three strategies applied to the country subgraphs reveal that national grids fragment in qualitatively different ways. Figure 7 shows four countries chosen to span this range. France is the clearest case of betweenness concentration: random failure leaves it almost intact ($S = 0.65$ at $\varphi_{\max}$), but betweenness-based removal collapses it to $S = 0.07$ by $\varphi = 0.05$ and to $S = 0.005$ by $\varphi_{\max}$, consistent with the small set of high-betweenness French substations identified in Section 5.3. Bulgaria shows the opposite emphasis: its widest gap is between random and *degree*-based removal ($S = 0.47$ versus $S = 0.22$ at $\varphi_{\max}$), the largest such gap of all 25 countries, indicating that its connectivity hangs on a few high-degree hubs rather than betweenness bottlenecks. Norway is fragile under every strategy: even random failure reduces it to $S = 0.33$,

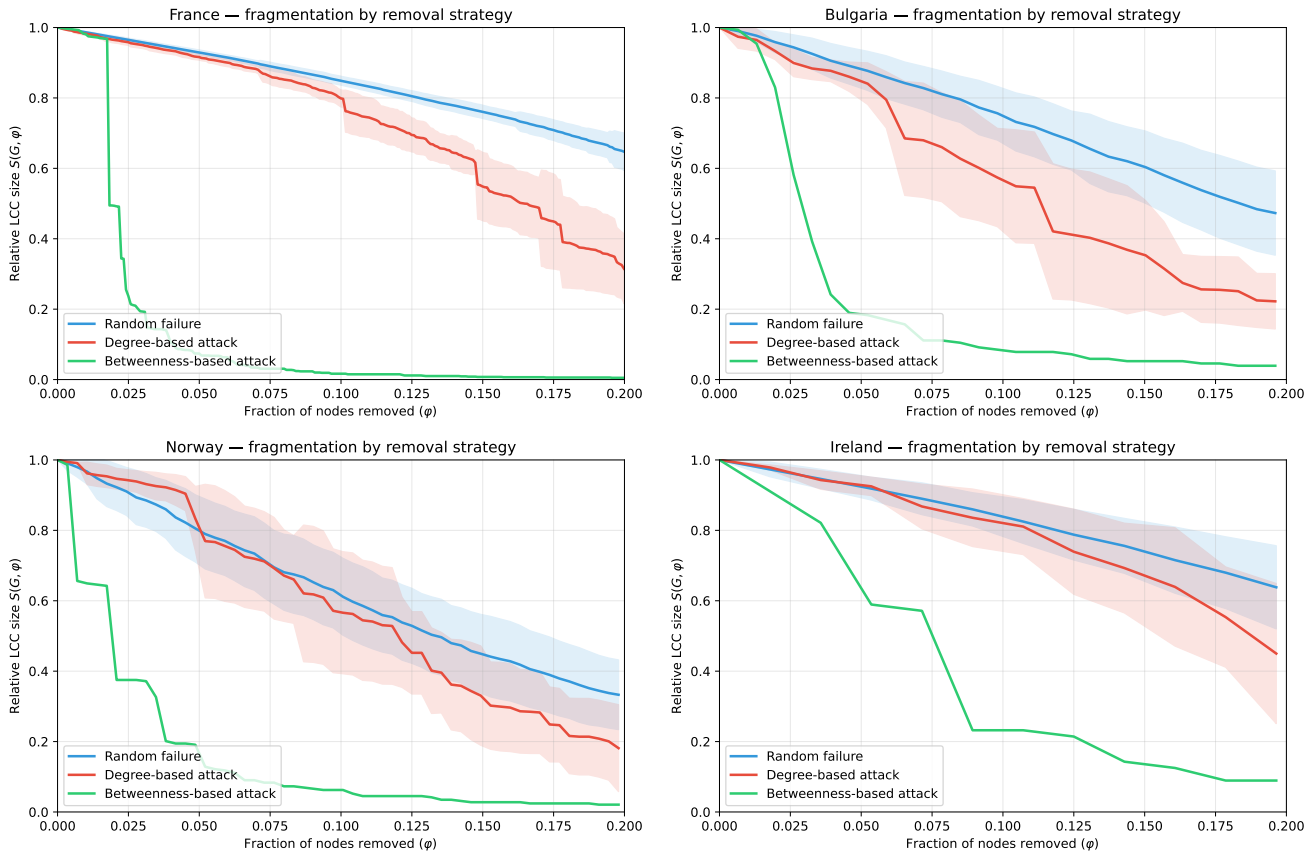

Figure 7: Fragmentation curves for France, Bulgaria, Norway and Ireland under the three removal strategies. Shaded bands show one standard deviation across runs for the random and degree-based strategies.

reflecting the elongated, chain-like topology already visible in its path lengths (Section 5.4). It fragments for a different reason than the others, having no dense core to lose: it is structurally thin throughout. Ireland is the most robust of the four, with random and degree-based curves staying close throughout ($S = 0.64$ and $S = 0.45$), meaning a degree-based attacker gains little over chance.

The vulnerability score R_{DEG} is only meaningful when the gap it measures is larger than the run-to-run noise of the stochastic strategies. Figure 8 makes this explicit for the two extreme countries and one near-zero case, drawing the standard-deviation bands around the random and degree-based curves. For Belgium the two bands overlap across the whole range; its degree-based curve even sits marginally above the random one ($S = 0.41$ versus $S = 0.40$ at $\varphi_{\max}$). Belgium's low $R_{\mathrm{DEG}} = 0.16$ therefore means that targeting adds no measurable damage over random failure, not that the grid is exceptionally strong. For Bulgaria and France the bands separate cleanly, so their high R_{DEG} and R_{BC} reflect a real structural effect rather than simulation noise. Across all four panels the betweenness curve drops fastest, confirming at the country level the pattern seen for the full grid.

Vulnerability ranking. Ranking the 25 countries by their normalised degree-based score (Figure 9) gives a spread of more than sixfold, from $R_{\mathrm{DEG}} = 0.16$ to $R_{\mathrm{DEG}} = 1.00$, confirming that the

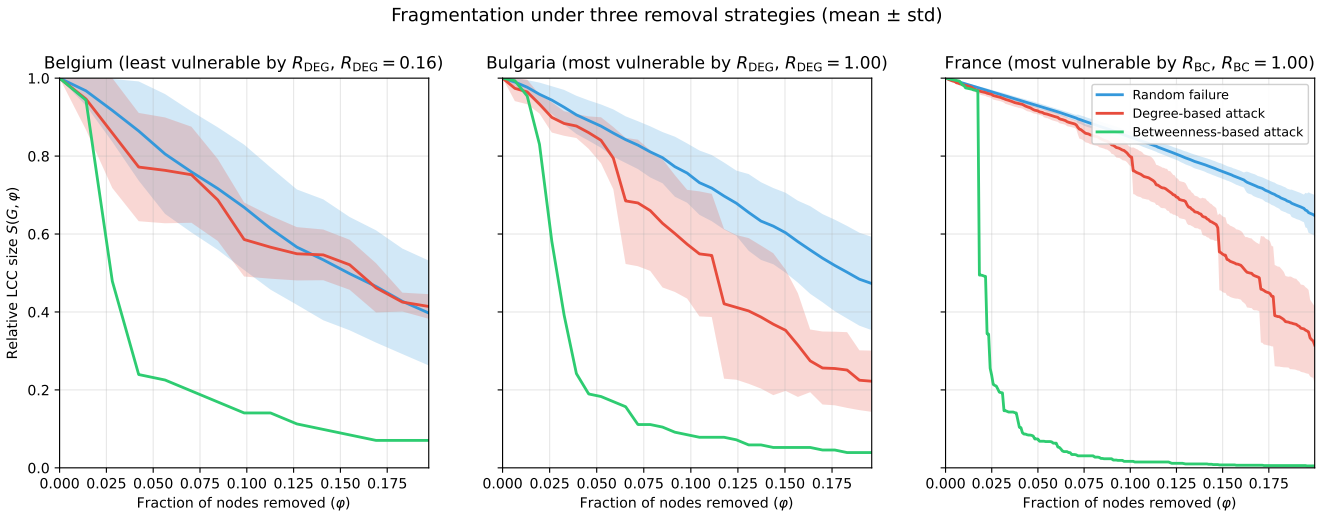

Figure 8: Fragmentation curves for Belgium, Bulgaria and France under random, degree-based and betweenness-based node removal. Shaded bands show one standard deviation across runs for the random and degree-based strategies.

structural differences of Section 5.2 translate into measurable differences in vulnerability. Bulgaria ($R_{\mathrm{DEG}} = 1.00$), Germany (0.78), Italy (0.77) and Spain (0.76) are the most exposed to a degree-based attack; Belgium (0.16), Poland (0.21) and Ireland (0.26) the least. The betweenness-based ranking is different: there France ($R_{\mathrm{BC}} = 1.00$), Spain (0.98) and Germany (0.96) lead, the large grids whose shortest-path traffic is funnelled through a handful of nodes. Figure 10 plots the two scores against each other. Almost every country lies above the diagonal, meaning a betweenness-based attack is at least as damaging as a degree-based one everywhere; the large western grids cluster in the high-vulnerability corner, while Bulgaria sits alone below the diagonal, vulnerable to degree-based removal but not to betweenness, because its load is spread evenly rather than concentrated in bottlenecks.

Structural correlates of vulnerability. To identify which structural properties drive these scores, each is correlated with R_{DEG} and R_{BC} across the 25 countries (Table 3). Degree-based vulnerability has only one significant predictor: degree assortativity ($r = +0.46$), confirming the link anticipated in Section 5.2, that grids whose high-degree substations connect preferentially to other high-degree substations fragment fastest when those hubs are removed. No other property correlates significantly with R_{DEG}. Betweenness-based vulnerability is far more strongly determined. It rises with the concentration of shortest-path traffic, measured by the betweenness Gini coefficient ($r = +0.72$), with average path length ($r = +0.70$) and with assortativity ($r = +0.51$), and falls with average node betweenness ($r = -0.88$). Together these describe large, sparse grids with long internal distances, in which a few nodes carry most of the routing, the structural condition that a betweenness-based attack exploits. Average clustering correlates with neither score.

Coverage and vulnerability. Because dataset coverage varies between countries (Section 4.2), a final check confirms that the ranking is not an artefact of missing data. Figure 11 plots each country's degree-based vulnerability against its coverage ratio. There is no systematic relationship:

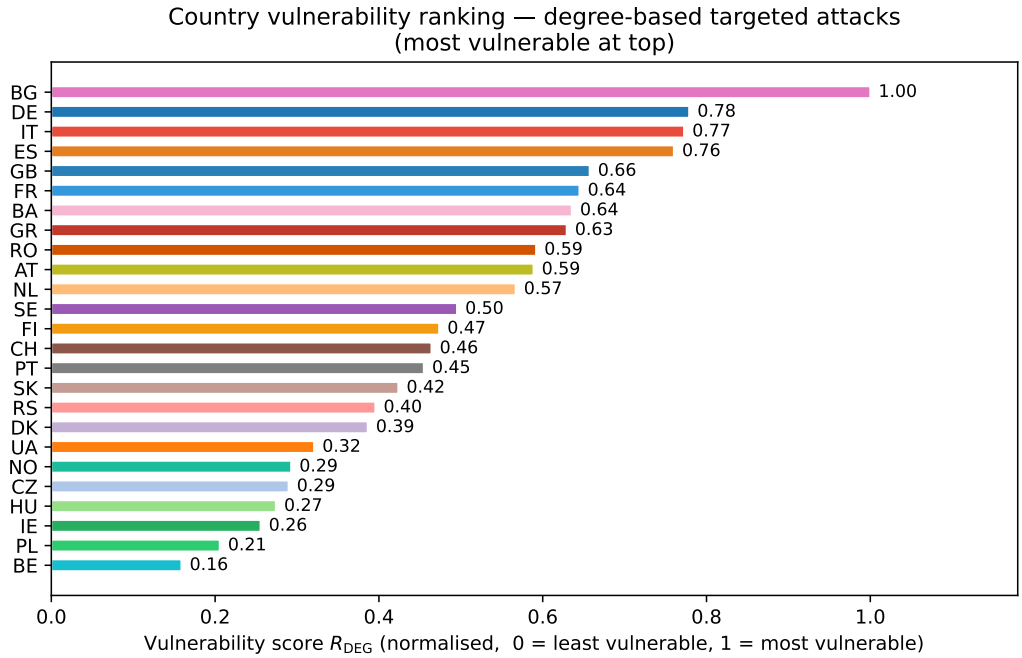

Figure 9: The 25 countries ranked by normalised degree-based vulnerability score R_{DEG} (Equation 10), most vulnerable at the top.

high- and low-vulnerability countries appear across the whole coverage range, so the differences in R_{DEG} reflect structure rather than how completely each grid is represented. The one country to read with caution is the Netherlands, whose coverage of 0.14 (Tier C) means its mid-ranking score likely understates its true vulnerability, since the missing infrastructure is probably the high-degree backbone that a targeted attack would exploit.

In answer to Sub-question 2, the European transmission grid is robust against random failure but noticeably vulnerable to targeted attack, and the two targeted strategies expose two different weaknesses. Degree-based vulnerability is governed by degree assortativity and is highest in Bulgaria, Germany, Italy and Spain; betweenness-based vulnerability is the more severe of the two and is

Table 3: Pearson correlation of seven structural properties with the normalised vulnerability scores R_{DEG} (degree-based) and R_{BC} (betweenness-based) across the 25 countries. Values marked * are significant at $p < 0.05$.

Structural property	R_{DEG}	R_{BC}
Average degree	-0.14	$+0.41^*$
Bridge fraction	$+0.04$	-0.45^*
Average clustering	-0.17	-0.19
Degree assortativity	$+0.46^*$	$+0.51^*$
Average betweenness	-0.32	-0.88^*
Average path length	$+0.32$	$+0.70^*$
Betweenness Gini	$+0.31$	$+0.72^*$

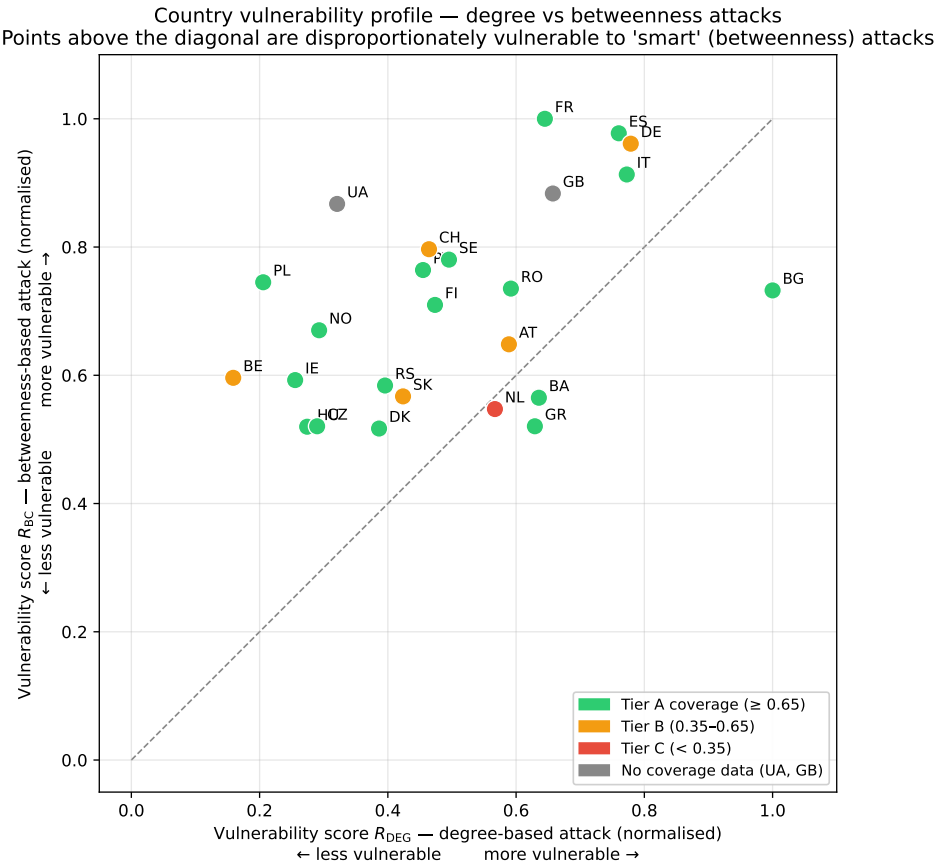

Figure 10: Each country's normalised vulnerability to a degree-based attack (R_{DEG}, horizontal axis) against its vulnerability to a betweenness-based attack (R_{BC}, vertical axis). Marker colour indicates dataset coverage tier; the diagonal is $y = x$. Points above the diagonal are more vulnerable to a betweenness-based attack than to a degree-based one.

governed by the concentration of shortest-path traffic, peaking in the large western grids of France, Spain and Germany. The next section maps these results back onto geography to locate where this vulnerability is physically concentrated.

5.6 Geographic Vulnerability

This section answers Sub-question 3 by projecting the network-level results of Section 5.5 back onto the map of Europe, to see where the structural vulnerability they measure is physically concentrated. A pan-European map first shows how the two vulnerability scores are distributed across the continent, after which three national case studies, France, Bulgaria and Norway, zoom in on the countries that sit at the extremes of the ranking and show what their vulnerability looks like on the ground.

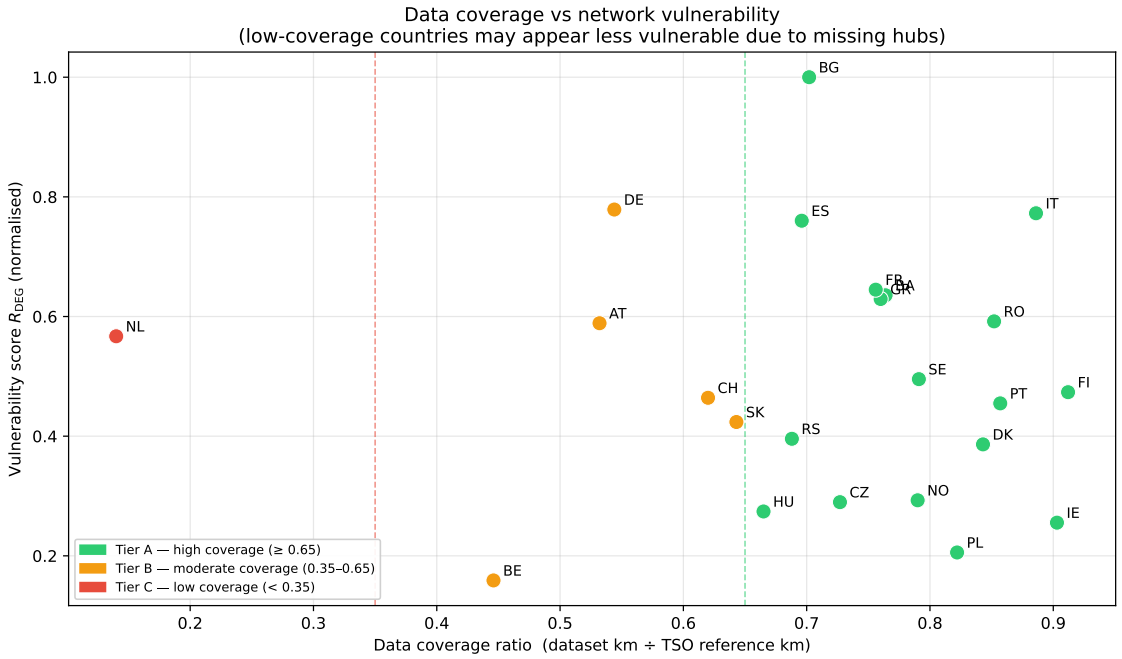

Figure 11: Normalised degree-based vulnerability score R_{DEG} against dataset coverage ratio (dataset circuit-kilometres divided by TSO reference circuit-kilometres). Marker colour indicates coverage tier; dashed lines mark the Tier A/B and Tier B/C thresholds.

Pan-European map

Figure 12 shades each country by its normalised vulnerability score, degree-based (R_{DEG}) on the left and betweenness-based (R_{BC}) on the right, with the cross-border interconnectors and the most central substations overlaid on the betweenness panel. The two panels show different geographies of weakness. Degree-based vulnerability is scattered: the darkest countries, Bulgaria, Germany, Italy and Spain, do not form a connected region. Betweenness-based vulnerability, by contrast, concentrates in a contiguous western band, France, Spain, Germany and Italy, the large interconnected grids through which most of the continent's shortest-path traffic is routed. The top-1% betweenness nodes ($n = 68$) cluster tightly inside this band, overwhelmingly in France, and they sit where the cross-border interconnectors are densest. Structural and geographic criticality therefore coincide: the substations that the network analysis flags as most central are also those that physically tie the western European grids together.

France

Figure 13 colours each French substation by its betweenness centrality on the full European grid. The high-centrality nodes are not spread evenly but trace a corridor running from the Paris basin south through central France to the Spanish border, the same backbone identified numerically in Section 5.3. This is the geographic counterpart of France's $R_{\mathrm{BC}} = 1.00$: removing this thin chain of substations is what collapses the French grid so quickly under a betweenness-based attack. Because these nodes also carry traffic between Iberia and the rest of the continent, their vulnerability is not France's alone.

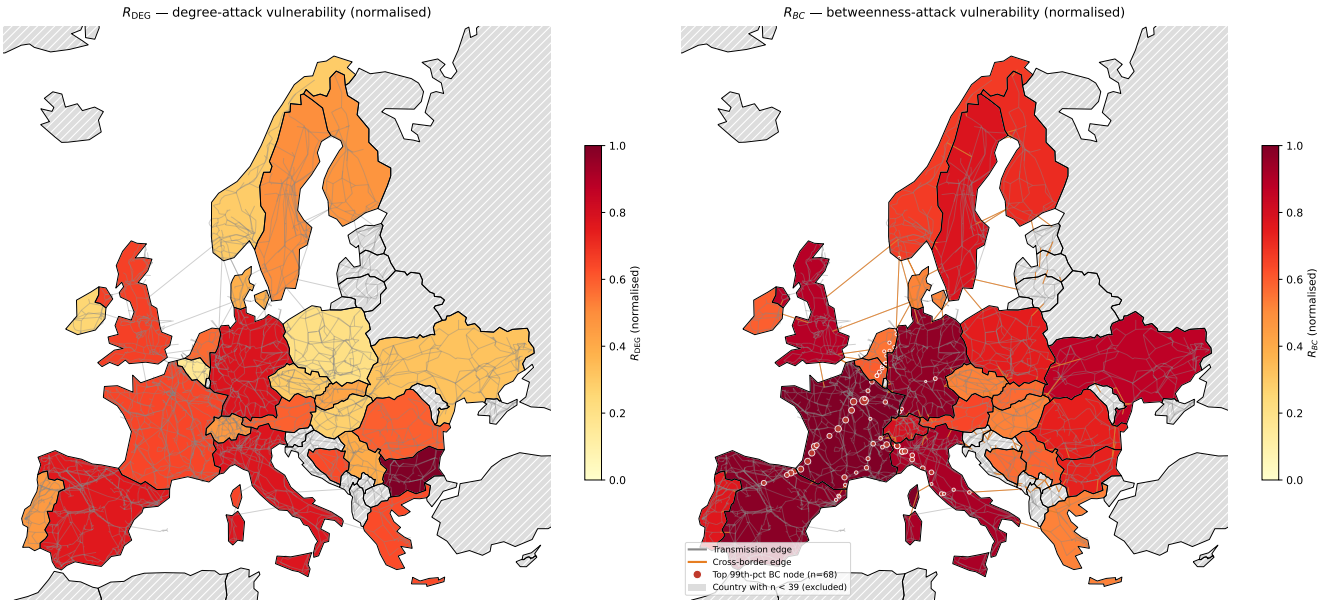

Figure 12: Pan-European vulnerability map. Left: countries shaded by normalised degree-based score R_{DEG}. Right: countries shaded by normalised betweenness-based score R_{BC}, with cross-border interconnectors in orange and the top-1% betweenness nodes ($n = 68$) as red points. Countries with a largest connected component smaller than 39 nodes are hatched and excluded.

Bulgaria

Bulgaria is the most vulnerable country to a degree-based attack ($R_{\mathrm{DEG}} = 1.00$) but only mid-ranked against betweenness. Figure 14 colours its substations by node degree, and the pattern explains the contrast: a few high-degree hubs stand out clearly, and their removal is what fragments the network, but the shortest-path load is spread across many nodes rather than funnelled through a single corridor. Bulgaria's weakness is therefore local, concentrated in identifiable hubs, rather than the distributed bottleneck structure seen in France.

Norway

Norway is the opposite case again: fragile under every strategy, yet without any dominant critical nodes. Figure 15 colours its substations by betweenness on the same scale as the France map. The whole grid renders pale, because Norway's maximum betweenness reaches only 66% of the value the French scale reaches. There is no concentrated bottleneck to attack; instead the elongated, chain-like topology (Section 5.4) makes the grid thin everywhere, so it fragments readily even under random failure. Norway's vulnerability is diffuse, a property of the whole network rather than of a few nodes.

Synthesis

In answer to Sub-question 3, structural vulnerability in the European transmission grid is concentrated rather than uniform. Betweenness-based vulnerability forms a contiguous western band, centred on the French corridor from the Paris basin to the Spanish border where the most central

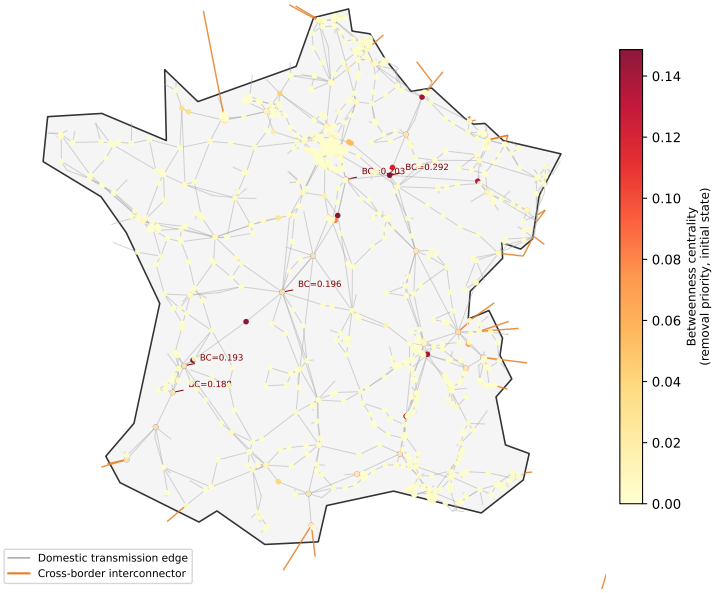

Figure 13: Mainland France, each substation coloured by betweenness centrality on the full European largest connected component. Darker red indicates higher betweenness centrality.

substations and the densest cross-border interconnectors coincide, while Bulgaria and Norway exemplify two further failure modes: localised high-degree hubs and whole-network thinness.

6 Discussion

This chapter steps back from the individual results and considers what they mean together. Section 6.1 interprets the findings of Chapter 5 in relation to the three sub-questions and draws them into a single picture of how the European grid is vulnerable. Section 6.2 discusses what this structural view makes it possible to understand or do. Section 6.3 sets out the limitations of the topological approach and the assumptions on which the results rest.

6.1 Interpretation of the Findings

The three sub-questions were answered separately, but their answers form one argument. The structural characterisation (Sub-question 1) found a grid that is sparse, low-degree and locally tree-like, neither scale-free nor small-world in the strict sense, whose variation between countries lies almost entirely in degree assortativity, bridge density and path length. Earlier studies of European grids report the same structural picture: light-tailed degree distributions and no dominant hubs [RCVS07, ELR18, HS21]. The failure simulations (Sub-question 2) then showed that these same properties are what govern vulnerability. Degree assortativity is the one significant predictor of vulnerability to a degree-based attack, while the concentration of shortest-path traffic, together with long path lengths, predicts vulnerability to a betweenness-based attack. Rosas-Casals et al. [RCVS07] classified grids as robust or fragile from the decay of the degree distribution alone; the results here add that how hubs connect to one another matters at least as much as how many

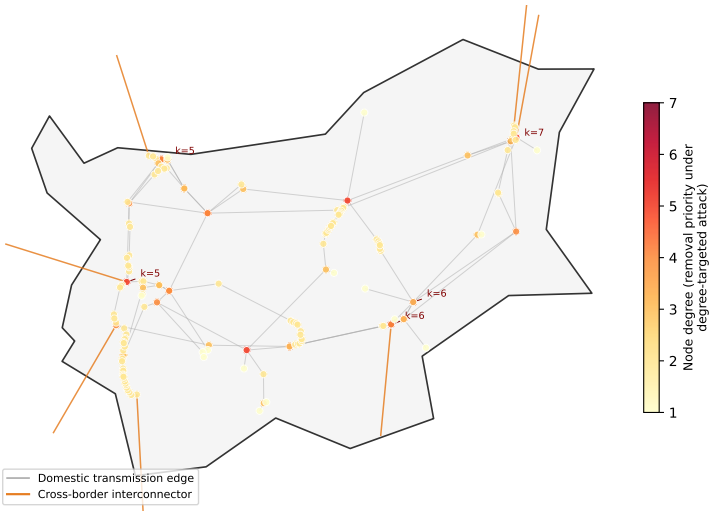

Figure 14: Bulgaria, each substation coloured by node degree on the country largest connected component. Darker red indicates higher degree.

there are. The structure described in the answer to the first sub-question is therefore not neutral; it is the cause of the vulnerability measured in the second sub-question.

The clearest interpretive result is that the grid does not have a single vulnerability but two, and that they do not coincide. A degree-based attacker, who can see only how many lines meet at each substation, and a betweenness-based attacker, who understands the whole network, fragment different countries. Bulgaria is the most exposed to the first but only mid-ranked against the second; France is the reverse, near the middle on degree but the most vulnerable country of all to a betweenness attack. The severity of the informed attack fits the foundational studies. In the North American and Italian grids, removing four percent of the highest-load nodes cost roughly sixty percent of connectivity [AAN04, CLM04]; the European grid likewise loses most of its connectivity within the first few percent of targeted removals. Hines et al. [HCSB10] reported a comparable disagreement between attack strategies, and the same pattern appears here: the two attackers identify different critical nodes. This means that "how vulnerable is this grid" has no single answer: it depends on what an attacker knows. The betweenness-based case is consistently the more destructive of the two, which matters because it is also the one that requires the most information to carry out.

Projecting these results onto geography (Sub-question 3) showed that vulnerability is concentrated rather than spread out. The most critical region is the western betweenness band, centred on the French backbone that runs from the Paris basin to the Spanish border, where the highest-centrality substations and the densest cross-border interconnectors coincide. That the nodes identified by a purely topological metric fall on the physical corridors tying the western European grids together suggests that the abstraction, despite its limitations, captures something real about where the system is fragile. Forsberg et al. [FTB23] found the highest-betweenness nodes of the Nordic grid clustered in a single region; the western band found in this study extends that observation to the scale of the continent. The ranking also agrees with the operational records studied by Rosas-Casals and Corominas-Murtra [RCCM09], in which France, the most betweenness-vulnerable country here, showed comparatively high interruption metrics. Hartmann and Cirunay [HC26], the closest

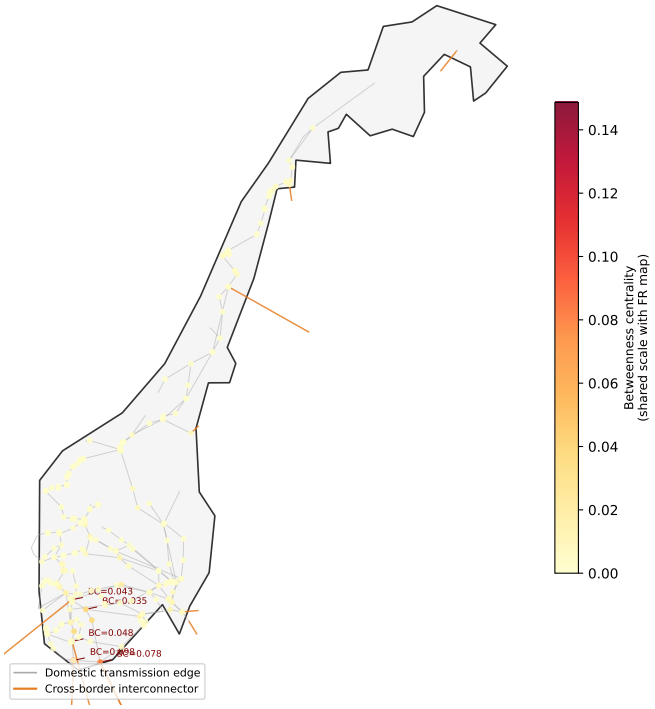

Figure 15: Norway, each substation coloured by betweenness centrality on the same colour scale as the France map (Figure 13). Norway's maximum betweenness reaches 66% of the value at which the French scale saturates, so the whole grid renders pale.

comparable study, apply a single removal strategy; the two scores used here show something that the single removal strategy cannot: a country's vulnerability rank depends on what the attacker knows.

6.2 Implications for Infrastructure Resilience

The practical value of this approach is that it locates structural weakness from topology alone, without the detailed electrical data that a power-flow study requires, and does so at a scale, 25 countries and the full continental grid, that such studies rarely reach. For an initial assessment of where a transmission system is most exposed, a structural screening of this type is cheap to compute and easy to update as the network changes, and it points to the substations and regions where a more detailed, physically grounded analysis would be most worthwhile.

The distinction between the two attack strategies has a direct reading for protection. The degree-based and betweenness-based results correspond to two different threat models, an adversary with only local knowledge and one with full knowledge of the network, and they identify different nodes. A defender therefore cannot treat "critical substations" as a single list: the nodes worth protecting against an informed attacker are the betweenness hubs of the French-Iberian corridor, which are not the same as the high-degree hubs that a less informed attacker would reach for. The maps also show that vulnerability is concentrated in a few regions rather than spread evenly across the grid. Protective effort can therefore be focused on those few corridors instead of being divided

across the whole network.

6.3 Limitations

These results rest on a deliberate simplification, and its limits should be stated. The analysis is purely topological: the grid is modelled as an unweighted graph with no representation of the physics that actually governs it. Line impedances, thermal capacities, generation and load are all absent, and power is assumed to take shortest paths, whereas in reality it spreads over all available lines according to their electrical properties. Real failures propagate through the redistribution of power flow and can cascade in ways that a static, flow-free model cannot reproduce. The nodes that matter most in the graph are not always the nodes whose failure would matter most in operation: Hines et al. [HCSB10] found only a mild correlation between topological rankings and the blackout sizes produced by power-flow simulations. The vulnerability scores in this thesis should therefore be read as structural indicators of where the network is fragile, not as predictions of how a real grid would fail.

The failure model itself is also a simplification. Nodes are removed one at a time and the only response recorded is the loss of connectivity; there is no cascading dynamics and no load redistribution beyond the shortest-path proxy that betweenness provides. The edge modelling introduces further assumptions: alternating-current lines, transformers and HVDC links are collapsed into a single type of undirected, unweighted edge (Section 4.3), which discards the differences in capacity and electrical role between them.

Finally, the results are only as complete as the data behind them. PyPSA-Eur is derived from OpenStreetMap, and its coverage varies by country (Section 4.2); for low-coverage countries such as the Netherlands a substantial part of the national grid is absent, so their scores must be read with caution. The coverage ratios used to flag this were themselves assembled from transmission-operator figures published in different years, and so give an approximate rather than an exact measure of completeness. The country-level results also exclude cross-border interconnectors by construction, so they describe each national grid in isolation rather than as part of the synchronous European system. None of these limitations undermines the structural comparison the thesis sets out to make, but each describes a boundary on how far its results can be pushed.

7 Conclusion

This thesis set out to answer how the European power transmission grid can be modelled as a network in order to identify its structurally vulnerable regions. The approach was to build an unweighted topological graph of the grid from the PyPSA-Eur dataset, characterise its structure, simulate the targeted removal of its nodes, and project the resulting vulnerability scores back onto the map of Europe. Taken together, the three sub-questions show that this purely structural model is enough to locate where the grid is most fragile, and that the regions it identifies are coherent rather than arbitrary.

The structural analysis established that the European grid is sparse, low-degree and locally tree-like, and that it is not scale-free; its clustering is high, but its long path lengths keep it from being small-world in the strict sense. What separates one national grid from another is not basic connectivity, which is almost uniform, but higher-order structure: degree assortativity,

bridge density and path length. The failure simulations then showed that these same properties determine vulnerability. The grid withstands random failure well but fragments quickly under targeted attack, and the two targeted strategies expose two distinct weaknesses: vulnerability to a degree-based attack is governed by degree assortativity, while vulnerability to the more destructive betweenness-based attack is governed by the concentration of shortest-path traffic and by long path lengths. Mapping these results onto geography showed that vulnerability is concentrated rather than spread out. The most structurally vulnerable region is the western betweenness band centred on the French backbone running from the Paris basin to the Spanish border, where the most central substations and the densest cross-border interconnectors coincide; high-degree countries such as Bulgaria, Germany, Italy and Spain are most exposed to a degree-based attack instead. A few sparse grids, Norway above all, are fragile under any failure, targeted or not. The fact that a topological metric points to the physical corridors that tie the western European grids together indicates that the model, despite its abstraction, captures a real feature of the system.

The main research question can therefore be answered as follows. The European transmission grid can be modelled as a topological network, and this model is enough to identify its vulnerable regions: it combines centrality measures, a node-removal simulation under three attack strategies, and a projection of the results back onto the map, and it needs no detailed electrical data to do so. These vulnerable regions are not spread evenly across the continent. They concentrate in a small number of countries and corridors, above all the French-Iberian betweenness backbone.

The most direct continuation of this work turns its question around. This thesis identifies vulnerability by removing nodes and measuring how the grid falls apart; the inverse problem is one of reinforcement: given the vulnerability scores, which nodes or edges could be added to improve resilience the fastest? This is a concrete network-optimisation question. Instead of only describing where the grid is weak, an analysis of this kind would point to where reinforcement helps most, that is, to the few new or rewired connections that lower a country's vulnerability score the furthest for the least added infrastructure. It would make the descriptive method developed here usable for planning.

A second direction follows from the central limitation of this work. Hines et al. [HCSB10] compared topological rankings with the blackout sizes from power-flow simulations for the North American grid and found only a mild correlation; for the European grid, no comparison of this kind exists on a current dataset. PyPSA-Eur makes one practical. The model assigns every line a standard type per voltage level, from which its impedance and thermal capacity follow, so the electrical parameters this thesis deliberately discards are already available in the dataset. A full flow-based vulnerability analysis, recomputing flows for every removal across every country, remains out of reach, but a single DC power-flow calculation is linear and computable at the size of the European network, which is enough to test whether the substations ranked as critical here remain critical when the physics is included. Cascading-failure simulations cost more, but they can be restricted to the regions this analysis flags as critical, above all the French–Iberian corridor. The points of disagreement would show what the topological abstraction misses.

Finally, the data and its scope leave room for refinement. Re-running the analysis on more complete or more recent snapshots would reduce the coverage caveats noted above, while weighting edges by capacity or voltage, or tracking how the structure changes as renewable generation expands, would move the model closer to the real system without giving up the scalability that makes the topological approach useful at European scale.

References

[AAN04] Réka Albert, István Albert, and Gary L. Nakarado. Structural vulnerability of the North American power grid. *Physical Review E*, 69(2):025103, 2004.

[AJB00] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, 2000.

[Bar16] Albert-László Barabási. *Network Science*. Cambridge University Press, 2016.

[BHS18] Thomas Brown, Jonas Hörsch, and David Schlachtberger. PyPSA: Python for Power System Analysis. *Journal of Open Research Software*, 6(1):4, 2018.

[BNX09] Ettore Bompard, Roberto Napoli, and Fei Xue. Analysis of structural vulnerabilities in power transmission grids. *International Journal of Critical Infrastructure Protection*, 2(1–2):5–12, 2009.

[BPD12] E. Bompard, E. Pons, and Di Wu. Extended Topological Metrics for the Analysis of Power Grid Vulnerability. *IEEE Systems Journal*, 6(3):481–487, 2012.

[CLM04] Paolo Crucitti, Vito Latora, and Massimo Marchiori. Model for cascading failures in complex networks. *Physical Review E*, 69(4):045104, 2004.

[CSSDS+15] Lucas Cuadra, Sancho Salcedo-Sanz, Javier Del Ser, Silvia Jiménez-Fernández, and Zong Geem. A Critical Review of Robustness in Power Grids Using Complex Networks Concepts. *Energies*, 8(9):9211–9265, 2015.

[ELR18] Rafael Espejo, Sara Lumbreras, and Andres Ramos. Analysis of transmission-power-grid topology and scalability, the European case study. *Physica A: Statistical Mechanics and its Applications*, 509:383–395, 2018.

[FTB23] Samuel Forsberg, Karin Thomas, and Mikael Bergkvist. 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, 2023.

[GKL17] N. Gündüz, S. Küfeoğlu, and M. Lehtonen. Impacts of natural disasters on swedish electric power policy: A case study. *Sustainability*, 9(2):230, 2017.

[HC26] Bálint Hartmann and M. T. Cirunay. Topology and fragility of European high-voltage networks: A cross-country comparative analysis. *Energy Reports*, 2026.

[HCSB10] Paul Hines, Eduardo Cotilla-Sanchez, and Seth Blumsack. Do topological models provide good information about electricity infrastructure vulnerability? *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 20(3):033122, 2010.

[HHSB18] Jonas Hörsch, Fabian Hofmann, David Schlachtberger, and Tom Brown. PyPSA-Eur: An open optimisation model of the European transmission system. *Energy Strategy Reviews*, 22:207–215, 2018.

[HS21] Bálint Hartmann and Viktória Sugár. Searching for small-world and scale-free behaviour in long-term historical data of a real-world power grid. *Scientific Reports*, 11(1):6575, 2021.

[MP24] Somnath Maity and Premananda Panigrahi. Vulnerability and Efficiency Assessment of Complex Power Grids Using Current-Flow Line Centralities. arXiv:2404.10005, 2024.

[PA13] Giuliano Andrea Pagani and Marco Aiello. The Power Grid as a complex network: A survey. *Physica A: Statistical Mechanics and its Applications*, 392(11):2688–2700, 2013.

[RCCM09] Martí Rosas-Casals and Bernat Corominas-Murtra. Assessing European power grid reliability by means of topological measures. *WIT Transactions on Ecology and the Environment*, 121:515–525, 2009.

[RCVS07] Martí Rosas-Casals, Sergi Valverde, and Ricard V. Solé. Topological vulnerability of the European power grid under errors and attacks. *International Journal of Bifurcation and Chaos*, 17(07):2465–2475, 2007.

[SK24] Julia E. Sullivan and Dmitriy Kamensky. Putin's power play: Russia's attacks on Ukraine's electric power infrastructure violate international law. *The Electricity Journal*, 37(2):107371, 2024.

[SRCCMV08] Ricard V. Solé, Martí Rosas-Casals, Bernat Corominas-Murtra, and Sergi Valverde. Robustness of the European power grids under intentional attack. *Physical Review E*, 77(2):026102, 2008.

[SWTL18] Benjamin Schäfer, Dirk Witthaut, Marc Timme, and Vito Latora. Dynamically induced cascading failures in power grids. *Nature Communications*, 9(1):1975, 2018.

[WKK+23] Wensheng Wang, Faezeh Karimi, Kaveh Khalilpour, David Green, and Manos Varvarigos. Robustness analysis of electricity networks against failure or attack: The case of the Australian National Electricity Market (NEM). *International Journal of Critical Infrastructure Protection*, 41:100600, 2023.

[WS98] Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, 1998.

[WZPS16] Fan Wenli, Liu Zhigang, Hu Ping, and Mei Shengwei. Cascading failure model in power grids using the complex network theory. *IET Generation, Transmission & Distribution*, 10(15):3940–3949, 2016.

A Supplementary Material

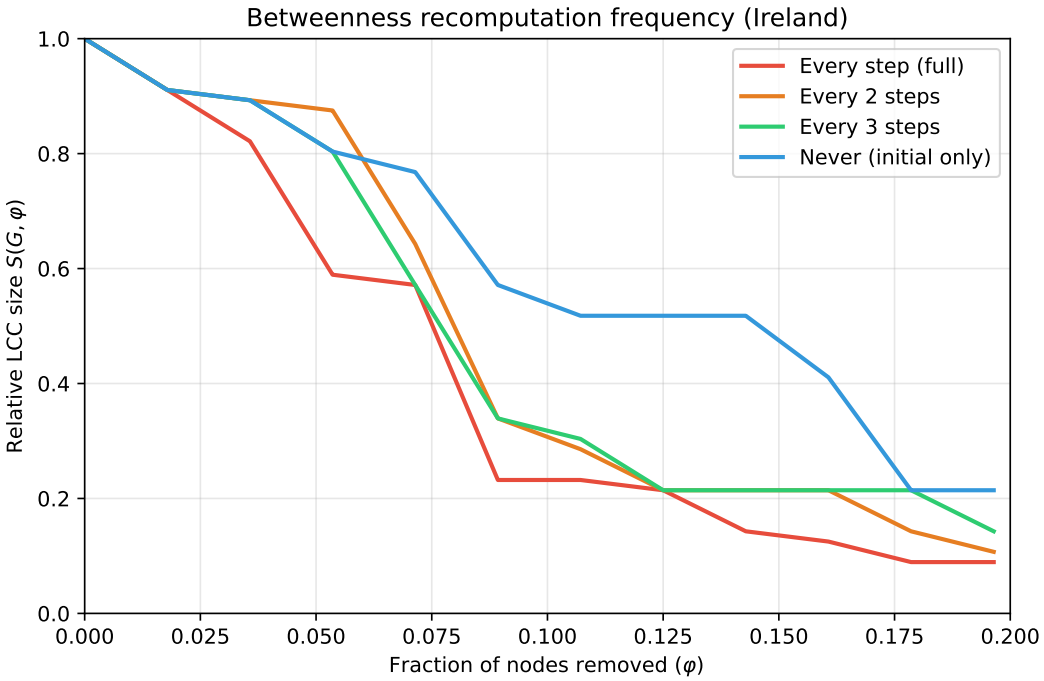

Figure 16: Fragmentation curves for the Ireland subgraph under betweenness-based removal with four recomputation frequencies: after every step, every two steps, every three steps, and computed once at the start.

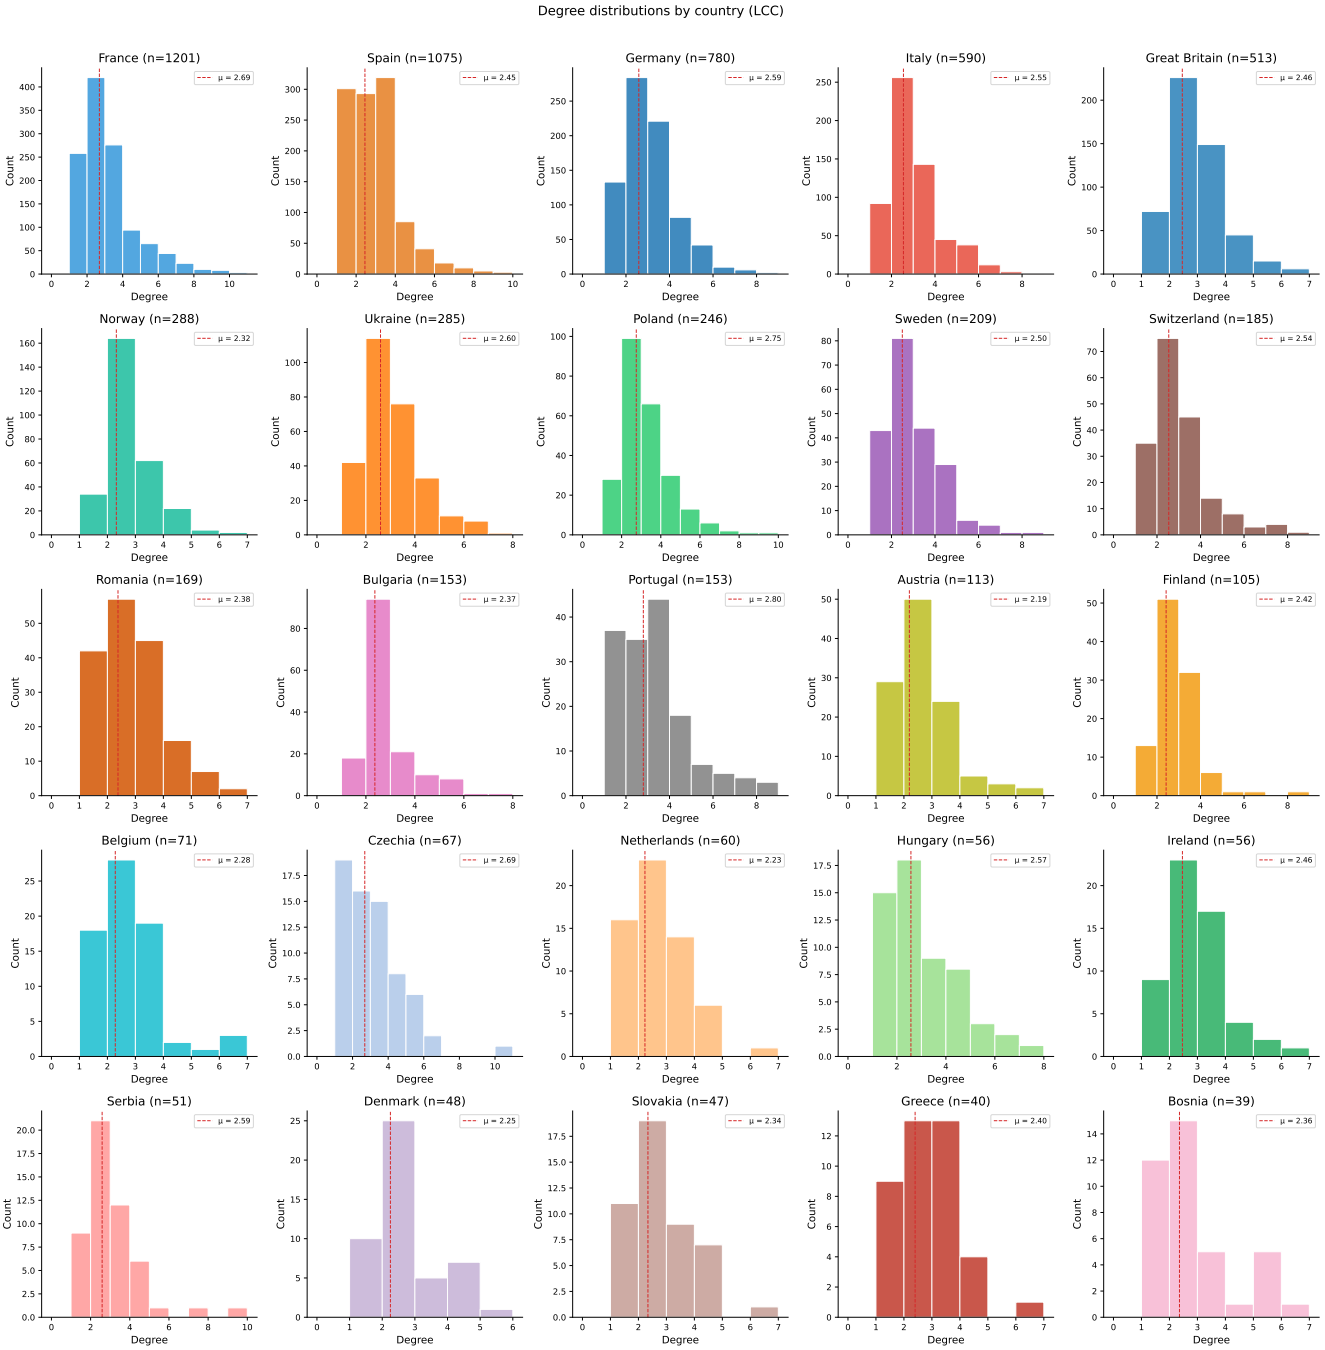

Figure 17: Degree distributions of the largest connected component of each of the 25 countries. Each shares the low-degree, light-tailed shape of the full grid in Figure 3.