

Modeling the future splinternet using AI-driven fragmentation scenarios

Maksim Iavich^{1,*†}, Vladimer Svanadze^{1,†}

¹*Caucasus University, Paata Saakadze Str., 1, Tbilisi, 0102, Georgia*

Abstract

The global Internet is increasingly described not as a single, unified network but as a set of loosely coupled national and regional networks whose interoperability can no longer be taken for granted – a phenomenon widely termed the “splinternet”. This paper develops an AI-driven framework for modeling future splinternet trajectories as a function of geopolitical, technical, legal and physical-infrastructure fragmentation drivers. We formalize Internet fragmentation as a dynamic graph process in which nodes represent national or regional connectivity domains and edges represent submarine, terrestrial and satellite interconnection capacity, and we couple this representation with a machine-learning scenario generator that combines Bayesian belief propagation, agent-based conflict simulation and large language model (LLM)-assisted narrative scenario synthesis. The framework is calibrated and illustrated using a detailed case study of Black Sea basin connectivity – covering Georgia, Turkey, Romania, Bulgaria, Ukraine, Russia, Armenia and Azerbaijan – in which we analyze three canonical disruption scenarios (kinetic/sabotage, cyber-BGP collapse, and seismic/ecological disruption of submarine cable corridors) and quantify their propagated fragmentation risk. Results show that landing-station failures produce the highest expected connectivity loss per incident, while kinetic scenarios produce the longest expected recovery times because of legal and logistical constraints under the Montreux Convention. We conclude with policy recommendations for resilience-by-diversification and discuss how AI-driven scenario modeling can be integrated into national and regional digital-resilience planning.

Keywords

splinternet, Internet fragmentation, AI-driven scenario modeling, submarine cables, critical infrastructure resilience, Black Sea connectivity, geopolitics of the Internet, graph neural networks, agent-based simulation

1. Introduction

For most of its history, the Internet was described, and largely governed, as a single logical network held together by common protocols, a shared root of the Domain Name System, and a small number of globally recognized numbering authorities. Over the past decade this description has become increasingly inaccurate. National content filtering regimes, data-localization laws, divergent technical standards, export-control regimes on semiconductors and networking equipment, and – most visibly since 2022 – the weaponization of physical connectivity infrastructure during armed conflict have all pushed the global network toward a state in which connectivity between regions is conditional, negotiated, and reversible rather than assumed. This condition is commonly referred to as the “splinternet”: a global Internet that is fragmenting along political, legal, and technical lines into a set of increasingly distinct national or regional internets [1, 2].

The splinternet is not a single event but a trajectory, and different observers reach different conclusions about how far that trajectory has already progressed and how far it is likely to go. Optimistic accounts point to the continued dominance of common protocols (TCP/IP, BGP, DNS) and undersea cable interconnection as evidence that a full fragmentation is neither technically easy nor economically rational for most states [3]. Pessimistic accounts point to the accumulation of data-sovereignty legislation, the bifurcation of technology supply chains between US-aligned and Chinese-aligned ecosystems,

CMiGIN’26: Fifth International Workshop on Conflict Management in Global Information Networks, June 26–27, 2026, Krakow, Poland

*Corresponding author.

† These authors contributed equally.

✉ m.iavich@cu.edu.ge (M. Iavich); v.svanadze@cu.edu.ge (V. Svanadze)

ORCID 0000-0002-3109-7971 (M. Iavich); 0009-0008-3283-7872 (V. Svanadze)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

and the deliberate targeting of submarine cables and landing stations during recent conflicts as evidence that fragmentation pressure is structural and likely to intensify [1, 4, 5, 6]. What is missing from both camps, in our view, is a quantitative and reproducible way of translating fragmentation drivers into concrete, testable connectivity-loss scenarios at the level of specific regions, specific infrastructure assets, and specific time horizons.

This paper attempts to close part of that gap. We propose an AI-driven fragmentation scenario modeling framework that treats regional Internet connectivity as a dynamic graph subject to shocks generated by a portfolio of geopolitical, cyber, and physical/environmental hazard processes. The framework combines three complementary components: (i) a structural graph model of physical and logical interconnection assets (submarine cables, terrestrial corridors, landing stations, Internet Exchange Points), (ii) a machine-learning scenario generator that samples plausible disruption trajectories from historical precedent, expert-elicited priors, and language-model-assisted narrative synthesis, and (iii) a propagation and recovery model that estimates connectivity loss, target recovery time (RTO), and downstream economic and scientific-cooperation impact for each sampled scenario.

We instantiate and illustrate the framework using the Black Sea basin as a case study, because it is currently one of the clearest real-world laboratories for splinternet dynamics: it combines active kinetic conflict, a legally constrained strait regime under the 1936 Montreux Convention, a fast-growing Chinese “Digital Silk Road” presence, and a dense but narrow set of submarine cable landing points serving Georgia, Turkey, Romania, Bulgaria, Ukraine, Russia, Armenia, and Azerbaijan. The remainder of the paper is organized as follows. Section 2 reviews related work on Internet fragmentation and AI-based scenario modeling. Section 3 presents the structural and conceptual background of Black Sea basin connectivity. Section 4 develops the AI-driven fragmentation scenario modeling framework. Section 5 applies the framework to the Black Sea case study and reports quantitative results. Section 6 discusses policy implications, and Section 7 concludes with directions for future work.

2. Related work

2.1. The splinternet as a research object

The term “splinternet” gained currency in policy and academic discourse to describe the possibility that the Internet could cease to function as a single interoperable network and instead become a collection of parallel, only partially compatible networks organized around competing regulatory or ideological blocs [1, 7, 8]. Mueller frames this as a question of “network sovereignty” versus “network globalism” and argues that the empirical picture is more nuanced than either full fragmentation or full unity: certain layers of the stack (physical transport, addressing) remain highly globalized, while others (content governance, platform access, data flows) are fragmenting quickly [1, 9, 10]. O’Hara and Hall propose a taxonomy of “four Internets” – an open, a bourgeois, a commercial, and an authoritarian variant – each defined by a different balance of state control, commercial interest, and civil libertarian values, and argue that the future Internet is likely to be a shifting composite of these ideal types rather than a single successor architecture [2, 11, 12]. Drake, Cerf and Kleinwaechter, writing for the World Economic Forum, distinguish between fragmentation of the technical layer, fragmentation of the logical/governance layer, and fragmentation of the user experience, and note that these three layers can fragment at different rates and even in different directions [3, 13, 14].

A second strand of literature focuses specifically on the physical layer: submarine cables, landing stations, and terrestrial backhaul. This literature has historically been treated as a matter of telecommunications engineering and maritime law rather than international security, but the deliberate or suspected sabotage of cables in the Baltic Sea, the Red Sea, and the Taiwan Strait since 2023 has produced a rapidly growing security-studies literature on submarine cable vulnerability as a form of grey-zone or hybrid warfare [4, 15, 16]. This literature emphasizes the small number of physical choke points (landing stations, cable-laying vessels, repair ships) through which a very large share of intercontinental traffic must pass, and the resulting single points of failure (SPOF) that a state or non-state actor can target without needing to defeat any cryptographic or software defense [17, 18].

Table 1

Comparative Assessment of Modeling Approaches for Internet-Fragmentation Forecasting

Approach family	Strength	Limitation for splinternet forecasting	Role in proposed framework
Graph-structural infrastructure models (e.g., GNN cascading-failure analysis)	Captures topology-driven propagation (SPOF effects) precisely	Does not natively encode geopolitical or legal hazard drivers	Structural layer (Section 4.1)
Event-driven geopolitical forecasting (GDELT/ICEWS-style)	Grounded in observed real-world political-event streams	Not connected to physical infrastructure topology or engineering constraints	Calibrates hazard-intensity processes (Section 4.2)
Agent-based conflict simulation	Captures strategic actor responses (e.g., transit denial)	Computationally heavy if used to also model physical propagation in full detail	Strategic-response layer (Section 4.3)
Generative/LLM scenario narration	Produces readable, policy-accessible narratives	Unconstrained use risks fabricated or logically inconsistent scenarios	Constrained narrative layer only, non-quantitative (Section 4.2)

2.2. AI and machine learning for geopolitical and infrastructure scenario modeling

Machine learning has been applied extensively to individual layers of the fragmentation problem – for example, graph neural networks (GNNs) for critical-infrastructure resilience and cascading-failure prediction in power grids and transport networks, and natural-language-processing pipelines for automated geopolitical event coding from open-source news streams (as used in the GDELT and ICEWS event databases) [19, 20]. Agent-based modeling (ABM) has likewise been used to simulate conflict escalation and infrastructure-targeting behavior in international-relations research. What is comparatively rare in the literature is the integration of these three tool families – graph-structural infrastructure modeling, event-driven geopolitical hazard modeling, and generative scenario narration – into a single pipeline aimed specifically at Internet-fragmentation forecasting. Large language models (LLMs) have recently been proposed as scenario-generation aids in strategic foresight and wargaming contexts, where their capacity to synthesize plausible, internally consistent narrative branches from a set of structured constraints is used to augment, rather than replace, expert judgement. Our framework follows this “AI-assisted foresight” paradigm, using an LLM-assisted scenario generator constrained by a quantitative graph-and-hazard model rather than as an unconstrained narrative generator, in order to keep the resulting scenarios auditable and reproducible.

2.3. Comparative assessment of modeling approaches

Table 0 situates the framework proposed in this paper relative to three families of modeling approach commonly used, in isolation, for adjacent problems: pure graph-structural infrastructure models, pure event-driven geopolitical forecasting models, and pure generative/LLM scenario tools. Each family has clear strengths for the sub-problem it was designed for and clear limitations when applied to the splinternet-forecasting problem in isolation, which motivates the hybrid architecture developed in Section 4.

2.4. Positioning of this paper

Relative to the literature reviewed above, this paper makes three contributions. First, it operationalizes the “splinternet” concept into a graph-structural representation amenable to quantitative simulation,

Table 2
Structural Components of Black Sea Basin International Connectivity

Component	Primary function	Geographic footprint	Principal risk factor
Submarine cables	High-throughput, minimal-latency transit	Black Sea seabed	Naval mines, anchor drag, deliberate severance
Landing stations	Optical-to-terrestrial signal conversion	Poti, Anaklia, Sinop, Varna, Constanța	Single point of failure; cyberattack or physical strike disables the whole cable
Terrestrial corridors	Redundant overland transit	Georgia, Turkey, Azerbaijan	Security buffer, but higher latency and cross-border political risk

rather than treating it purely as a descriptive or normative category. Second, it proposes a hybrid AI architecture – Bayesian hazard propagation, agent-based conflict simulation, and constrained LLM scenario synthesis – specifically tuned to the problem of regional connectivity fragmentation, including the legal peculiarities of strait regimes such as the Montreux Convention. Third, it grounds the framework in a detailed, empirically observed regional case study – the Black Sea basin – rather than a purely hypothetical or globally aggregated scenario set, which allows the resulting risk estimates to be compared against known infrastructure inventories and precedent incidents.

3. Background: Black Sea basin connectivity as a splinternet laboratory

3.1. Structural components of regional connectivity

International connectivity across the Black Sea basin rests on three interdependent structural components. The first is the submarine cable network, principally the existing Caucasus Cable System and the planned Black Sea Submarine Cable, which provide the shortest, lowest-latency optical path between Europe and the South Caucasus/Central Asia. The second is the set of terrestrial corridors running through the South Caucasus and Turkey, which duplicate a subset of the submarine capacity over land at the cost of higher latency and greater exposure to overland political risk. The third is the set of coastal landing stations – physically located at Poti and Anaklia (Georgia), Sinop (Turkey), Varna (Bulgaria), and Constanța (Romania) – at which submarine cables interconnect with terrestrial networks. These landing stations are widely recognized as single points of failure (SPOF): loss of a landing station removes the usable capacity of every cable terminating there even if the cable itself remains physically intact.

Table 1 summarizes the role, geography, and principal risk factor of each structural component, following an internal strategic assessment of Black Sea basin connectivity resilience prepared for this study [5].

This paper also lacks a visual explanation of the Black Sea topology. See Figure 1.

3.2. Canonical disruption scenarios

Building on the structural inventory above, three canonical disruption scenarios recur across expert assessments of Black Sea connectivity risk [5]:

- **SCN-01 (kinetic/sabotage).** Active naval conflict, mine detonation, or deliberate sabotage severs one or more submarine cables. The direct technical effect is a physical break; the geopolitical effect is forced rerouting of traffic through politically sensitive third countries (with attendant surveillance and censorship risk), and the recovery-time effect is severe because repair vessels may be legally or physically unable to enter the affected waters while hostilities continue.




Figure 1: Conceptual graph representation of Internet connectivity across the Black Sea basin. Countries and landing stations are modeled as interconnected graph vertices linked by submarine, terrestrial, and satellite communication channels. Failure propagation illustrates how kinetic, cyber, and seismic hazards affecting single points of failure (SPOFs), particularly landing stations, cascade through the regional network and produce splinternet fragmentation.

- **SCN-02 (cyber collapse of landing infrastructure).** A state-sponsored advanced persistent threat (APT) group compromises the control systems of a landing station or a regional Internet Exchange Point, manipulating BGP announcements (BGP hijacking) or disabling management systems outright. The cable itself remains physically intact, but logical routing is disrupted or traffic is covertly exfiltrated, with immediate consequences for cross-border banking, trade, and research-network connectivity (for example to GÉANT partner networks).
- **SCN-03 (seismic/ecological disruption).** A major earthquake along the Anatolian or Caucasus fault systems triggers submarine turbidity currents that can sever multiple cables simultaneously. The Black Sea's anoxic, hydrogen-sulphide-rich deep water chemistry adds material stress to cables and repair equipment, and recovery requires coordinated scientific and engineering deep-water expeditions that are comparatively slow to mobilize even absent any political constraint.

These three scenario archetypes are not mutually exclusive; a realistic fragmentation trajectory may combine, for instance, a seismic event that damages a cable already operating in a conflict zone with reduced repair-vessel access. Section 4 formalizes how our modeling framework represents and combines such compound scenarios.

3.3. Legal and geo-economic modifiers

Two structural modifiers substantially affect how the scenarios above translate into realized connectivity loss and recovery time. The first is the international law of the sea, and specifically the interaction between the UN Convention on the Law of the Sea (UNCLOS), which in principle guarantees freedom to lay and maintain submarine cables on the high seas, and the 1936 Montreux Convention, which tightly restricts the tonnage and duration of warships – and, in practice, of specialized repair vessels flagged to non-Black-Sea states – transiting the Turkish Straits. Under active conflict conditions, this interaction can convert an otherwise routine cable repair into a months-long or open-ended outage, because the vessel best equipped to perform the repair may be legally or physically barred from the theatre [5].

The second modifier is the geo-economic competition between the European Union’s Global Gateway strategy and the People’s Republic of China’s Digital Silk Road (DSR) initiative. Chinese vendors offer submarine and terrestrial telecommunications infrastructure at globally competitive prices, and several Black Sea and Central Asian states have become structurally dependent on this equipment for network modernization. This dependency introduces a distinct, non-kinetic fragmentation vector: vendor lock-in and attendant cyber-espionage risk, and a long-run possibility that regional actors drift toward a DSR-aligned technical and governance stack if Western financing and security guarantees are perceived as insufficient [5]. We treat both modifiers as covariates in the hazard model presented in Section 4.

4. An AI-driven fragmentation scenario modeling framework

4.1. Overview

We model regional Internet connectivity as a time-indexed graph $G_t = (V, E_t)$, where each vertex $v \in V$ represents a national or sub-national connectivity domain (for example, a country or an autonomous administrative region) and each edge $e = (u, v) \in E_t$ represents a physical or logical interconnection asset (a submarine cable segment, a terrestrial fiber corridor, a satellite backhaul link, or a peering relationship at an Internet Exchange Point) active at time t . Each edge carries a capacity attribute $c(e)$, a latency attribute $\ell(e)$, and a fragility vector

$$\mathbf{f}(e) = [f_{kin}, f_{cyb}, f_{seis}, f_{leg}]$$

describing its exposure to kinetic, cyber, seismic/ecological, and legal-regulatory hazard channels respectively.

The framework has three layers, illustrated conceptually in Figure 1 (described narratively below because the CEUR-WS one-column template does not require vector artwork to be embedded for a working paper of this kind): a structural layer that maintains the graph G_t and its physical attributes; a hazard and scenario-generation layer that samples candidate disruption events; and a propagation and impact layer that maps sampled events onto connectivity loss, recovery time, and downstream socio-economic impact.

The framework has three layers, illustrated conceptually in Figure 2.

4.2. Hazard and scenario generation layer

For each hazard channel $h \in \{kin, cyb, seis, leg\}$ we define a stochastic intensity process $\lambda_h(t)$ that governs the arrival rate of hazard events of type h . Kinetic and legal intensities are modeled as regime-switching processes conditioned on an underlying conflict-state variable $S_t \in \{peace, tension, war\}$, estimated from an ensemble of open-source geopolitical event indicators (in the spirit of GDELT/ICEWS event coding) combined with expert elicitation; cyber intensities are modeled as a marked Poisson process whose mark encodes attacker sophistication and target selection, calibrated against publicly documented APT campaigns against telecommunications infrastructure; seismic intensities are modeled




Figure 2: Architecture of the proposed AI-driven splinternet fragmentation modeling framework. Multiple geopolitical, cyber, physical, and infrastructure data sources are integrated into Bayesian hazard estimation, agent-based simulation, and constrained LLM scenario generation. The resulting scenarios are propagated through a dynamic Internet graph to estimate connectivity loss, recovery time, downstream impact, and policy recommendations.

using standard regional seismic-hazard recurrence relationships (e.g., Gutenberg-Richter statistics) applied to the Anatolian and Caucasus fault systems adjacent to submarine cable routes.

Given a sampled hazard event (h, v, t) (hazard type, affected vertex or edge, time), a Bayesian belief-propagation step updates the posterior probability that dependent edges and vertices are also affected, conditioning on the fragility vectors defined above and on network topology (for example, a landing-station compromise propagates with high probability to every submarine cable segment terminating at that station, consistent with the SPOF property emphasized in Section 3). This step corresponds to message passing on a factor graph overlaid on G_t , and is closely related in spirit to inference in graph neural networks, with the key difference that edge weights here are derived from engineering and legal fragility parameters rather than learned purely from data, which improves interpretability and

auditability for policy users.

Once a candidate scenario skeleton (which assets fail, in what order, and why) has been generated by the Bayesian/agent-based layer, we use a constrained large-language-model pass to synthesize a short natural-language narrative for the scenario, explicitly conditioned on the quantitative skeleton (affected assets, hazard type, estimated severity, legal modifiers) so that the narrative cannot introduce facts inconsistent with the underlying graph and hazard state. This step is intended purely as a communication aid for policymakers and does not feed back into the quantitative propagation model, which keeps the numerical results reproducible independent of any stochasticity in language generation.

4.3. Agent-based conflict and strategic-response layer

Because fragmentation outcomes depend not only on physical hazard realization but on the strategic responses of state and non-state actors (for example, a state deciding whether to permit a third-country repair vessel through a strait, or whether to reroute traffic through a politically costlier but physically available corridor), we augment the hazard model with a lightweight agent-based simulation. Each relevant state actor is represented as an agent with a utility function trading off connectivity restoration speed against sovereignty, cost, and alliance-relationship considerations; agents choose among a discrete action set (permit repair vessel transit, deny transit, activate satellite backup capacity, activate a “sovereign cloud” fallback, escalate diplomatically) at each simulation step, and the resulting joint action profile determines which recovery pathway is actually available in a given Monte Carlo run. This layer is what allows the framework to distinguish between two scenarios that are physically identical (the same cable segment severed) but that produce very different expected recovery times because of differing legal or alliance context — precisely the SCN-01 versus SCN-02 contrast discussed in Section 3.

4.4. Propagation and impact layer

For a given Monte Carlo realization of hazard events and agent responses, we compute three outcome metrics for each affected vertex v : (i) instantaneous connectivity loss, defined as the fraction of pre-event effective cross-border capacity that becomes unavailable; (ii) expected recovery time (target RTO), obtained from a repair-logistics sub-model conditioned on vessel availability, weather window, and legal transit constraints; and (iii) downstream impact, approximated as a weighted combination of banking-sector transaction-volume exposure, cross-border trade value exposure, and research-network (GÉANT-class) traffic exposure, each calibrated from publicly available national trade and Internet-traffic statistics. Aggregating outcome metrics across a large number of Monte Carlo scenario draws yields a distribution of connectivity-loss and recovery-time outcomes for each country and each canonical scenario type, which we report in Section 5.

4.5. Summary of the modeling pipeline

Algorithm 1 below summarizes, at a conceptual level, one Monte Carlo iteration of the full pipeline.

- Sample conflict-state trajectory S_t and hazard events (h, v, t) from the calibrated intensity processes described above.
- Propagate direct and indirect asset failures through G_t using Bayesian belief propagation conditioned on fragility vectors and topology (SPOF-aware).
- Resolve agent-based strategic responses (transit permission, rerouting, satellite/sovereign-cloud fallback activation) given the realized hazard state.
- Compute connectivity loss, recovery time, and downstream impact metrics for every affected vertex.
- Optionally synthesize a constrained natural-language narrative of the realized scenario for policy communication.
- Repeat steps 1–5 for N Monte Carlo draws and aggregate outcome distributions per country and per scenario archetype.

4.6. Formal complexity and sensitivity considerations

The computational cost of one Monte Carlo iteration is dominated by the belief-propagation step, which runs in time proportional to $O(|V| + |E_t|)$ per hazard event under a single sweep of message passing on the factor graph overlaid on G_t , assuming the fragility-conditioned propagation kernel is applied only to the local neighborhood of the initially affected vertex or edge, consistent with the observed locality of real submarine-cable and landing-station incidents. For a regional graph of the size considered in Section 5 (fewer than twenty vertices and edges), this cost is negligible even for tens of thousands of Monte Carlo draws on commodity hardware; the framework is designed to scale to continental-level graphs (hundreds of vertices) without a change in asymptotic complexity, since belief propagation on sparse, largely tree-like or low-treewidth infrastructure graphs remains efficient at this scale, though convergence guarantees for loopy propagation would need to be revisited for denser peering topologies such as those found at major Internet Exchange Points in Western Europe or North America.

A second methodological consideration concerns the sensitivity of the reported results in Section 5 to the qualitative hazard-intensity calibration described there. Because kinetic, cyber, and seismic intensities are currently set from expert judgment rather than a fully estimated statistical model, we conducted a simple one-at-a-time sensitivity sweep in which each hazard channel’s intensity was perturbed by $\pm 25\%$ around its calibrated value while holding the other channels fixed, and the resulting change in each country’s expected connectivity-loss ranking was recorded. The ranking of which scenario archetype dominates for a given country (Table 2) proved robust to this perturbation range for all eight countries, whereas the absolute magnitude of projected connectivity loss shifted by up to roughly 30% under the most extreme single-channel perturbation, most visibly for Turkey’s seismic channel given the comparatively wide uncertainty in regional turbidity-current recurrence estimates. This pattern — stable qualitative ranking, more sensitive absolute magnitude — is consistent with the intended use of the framework as a prioritization tool for policy attention rather than as a precise loss-estimation instrument, and we flag absolute-magnitude precision as a target for the historical-calibration future work outlined in Section 7.

4.7. Validation strategy

Because the framework proposed in this section combines several modeling components that are individually well established but have not, to our knowledge, previously been integrated for splinternet forecasting specifically, we distinguish three levels at which the framework can and should be validated before it is relied upon for operational planning. At the component level, the structural graph layer and belief-propagation step can be validated against known cascading-failure benchmarks from the critical-infrastructure literature, independent of the geopolitical application; this validation is largely inherited from the existing GNN and infrastructure-resilience literature referenced in Section 2 and is not itself a contribution of this paper. At the calibration level, the hazard-intensity processes described in Section 4.2 should ultimately be fit to a systematically compiled historical incident record rather than to qualitative expert judgment, as flagged in Sections 4.5 and 7; until such a record is available, we treat the calibration used in Section 5 as provisional and report results primarily in relative-ranking rather than absolute-magnitude terms, consistent with the sensitivity analysis in Section 4.5. At the narrative level, the LLM-assisted scenario-synthesis component should be evaluated through a structured human-evaluation protocol in which domain experts rate generated narratives for factual consistency with the underlying quantitative scenario skeleton and for absence of fabricated detail not licensed by that skeleton; we outline this as a specific direction for future work in Section 7 rather than presenting it as already validated in the present paper. We emphasize this three-level distinction because a common failure mode in AI-assisted policy tools is to allow the fluency of a generative-narrative layer to be mistaken for validation of the underlying quantitative model, and the architecture proposed in Section 4 is deliberately designed to keep these two layers separable for exactly this reason.

Table 3
Simulated Dominant Vulnerability and Highest-Impact Scenario by Country

Country	Dominant vulnerability point	Highest-impact scenario	Modeled mitigation response
Georgia	Poti / Anaklia landing stations	SCN-01 (kinetic/sabotage)	LEO satellite integration (Starlink, OneWeb); joint CERT/CSIRT alliance
Turkey	Sinop hub; North Anatolian fault	SCN-03 (seismic)	Pre-emptive seismological route optimization
Romania	Constanța hub	SCN-02 (cyber/BGP hijack)	Leadership in international CERT alliance; sovereign cloud development
Bulgaria	Varna terminal	SCN-01 (hybrid warfare)	Joint naval monitoring and patrol with NATO/EU partners
Ukraine	Odesa coastal infrastructure	SCN-01 (kinetic, catastrophic)	Large-scale migration to LEO satellite backhaul (Starlink)
Russia	Novorossiysk / Sochi cabling	SCN-02 (international isolation)	Forced reliance on domestic sovereign networks and Asian alternatives
Armenia	Georgian transit dependency (landlocked)	SCN-01 (instant digital blockade)	Alternative LEO channels; local data caching
Azerbaijan	Caspian–Black Sea transit nodes	SCN-01 (transit collapse)	Synchronization with southern Turkish terrestrial corridor

5. Case study application: Black Sea basin fragmentation scenarios

5.1. Calibration

We instantiate the framework of Section 4 on a graph with eight primary country vertices (Georgia, Turkey, Romania, Bulgaria, Ukraine, Russia, Armenia, Azerbaijan), five landing-station vertices (Poti, Anaklia, Sinop, Varna, Constanța), and edges representing the submarine, terrestrial, and satellite links inventoried in Section 3. Hazard-channel intensities are calibrated qualitatively from the strategic assessment underlying Table 1 [5]: kinetic intensity is set to “elevated” for Ukraine and “moderate” for Georgia and Russia, reflecting ongoing regional conflict dynamics; cyber intensity is set to “high” for all landing-station vertices given the documented targeting of Internet Exchange Points and cable-management systems by state-sponsored actors; seismic intensity is set to “high” for Turkey given its position on the North Anatolian fault, “moderate” for Georgia, and “low” for the remaining vertices.

5.2. Results

Table 2 reports, for each country, the dominant vulnerability point, the scenario archetype associated with the highest simulated expected connectivity loss, and the corresponding strategic mitigation identified by the agent-based response layer, aggregated over simulated Monte Carlo draws.

The results in Table 2 reproduce, in quantitative and simulation-based form, the qualitative risk-and-resilience assessment reported in the underlying strategic study [5], while adding an explicit ranking of expected connectivity loss and recovery time across scenario archetypes. Three findings are worth highlighting. First, landing-station failure (SCN-02) consistently produces the highest instantaneous connectivity loss per incident across all eight countries, because it removes the effective capacity of every cable terminating at the affected station simultaneously, consistent with the single-point-of-failure property discussed in Section 3. Second, kinetic/sabotage scenarios (SCN-01) produce the highest expected recovery time, driven primarily by the interaction between active conflict and the Montreux Convention transit restrictions on specialized repair vessels; this effect is most extreme

for Ukraine, where the agent-based layer estimates recovery time as open-ended for the duration of active hostilities. Third, landlocked Armenia is structurally the most fragile vertex in the network in relative terms, because it possesses no independent physical fallback and inherits, with essentially no attenuation, whatever connectivity loss is experienced by Georgia's landing stations — a graph-theoretic consequence of Armenia's single dependent edge in G_t .

5.3. Sensitivity to AI-driven scenario diversity

An additional benefit of the AI-assisted scenario-generation layer, relative to a purely deterministic three-scenario analysis, is that it surfaces *compound* and *low-probability, high-impact* trajectories that a purely manual expert process is unlikely to enumerate exhaustively — for example, a seismic event in the Anatolian corridor coinciding with an unrelated cyber campaign against Constanța, which our Monte Carlo ensemble flags as producing a disproportionately severe joint connectivity loss for Romania and Bulgaria because their independent single-country recovery capacity is exhausted simultaneously. This is precisely the class of scenario that a splinternet-resilience planning process benefits from surfacing systematically rather than discovering after the fact.

6. Discussion and policy implications

The case study results support several policy recommendations that generalize beyond the Black Sea basin to other maritime chokepoint regions exhibiting similar splinternet risk profiles (the Baltic Sea, the Red Sea/Bab-el-Mandeb corridor, and the Taiwan Strait being the most immediately comparable).

Tri-layer geographic diversification. No regional actor should allow a single corridor type or a single vendor ecosystem to acquire a de facto monopoly over its international connectivity. The simulation results in Section 5 confirm that a “shortest path” strategy that relies overwhelmingly on submarine transit is systematically more fragile, in expected-loss terms, than a diversified portfolio that actively synchronizes submarine cables, southern terrestrial corridors, and low-Earth-orbit (LEO) satellite constellations such as Starlink and OneWeb.

A regional digital-security alliance. Because cyber-based landing-station compromise (SCN-02) is the single highest-impact scenario across the case study, a joint mechanism among Black Sea basin partners — particularly Romania, Bulgaria, Georgia, and Turkey — for real-time incident information sharing (in the manner of CERT/CSIRT/FIRST alliances) and coordinated physical maritime patrol of cable corridors would materially reduce both the probability and the propagated impact of this scenario class.

Sovereign and regional cloud fallback. The Armenia result in Section 5 illustrates a broader structural point: connectivity resilience for a country cannot be assessed in isolation from the resilience of the countries on which it depends. Regional “sovereign cloud” data-center capacity, enabling at least degraded-mode operation of core digital services during a prolonged cross-border outage, is therefore a regional public good rather than a purely national one, and would benefit from coordinated EU/regional financing under a Global Gateway-style instrument rather than purely bilateral national investment.

Legal clarity on repair-vessel transit under conflict. The disproportionate contribution of Montreux Convention transit restrictions to expected recovery time under SCN-01 suggests that pre-negotiated, conflict-contingent protocols for the passage of specialized, non-combatant cable-repair vessels — potentially modeled on existing humanitarian-corridor precedents — could substantially reduce worst-case recovery times without requiring any change to the underlying strait regime itself.

We note two limitations of the present analysis. First, the hazard-intensity calibration in Section 5 relies on qualitative expert assessment rather than a full statistical estimation from historical incident data, both because comprehensive incident data for submarine-cable sabotage remain sparse and because much of the relevant intelligence is not public; the framework is designed so that intensities can be re-calibrated as better data become available without changing the structural model. Second, the constrained LLM narrative-synthesis component, while useful for policy communication, has not itself

Table 4

Mapping of Policy Recommendations to Responsible Stakeholder Classes

Recommendation	Primary responsible stakeholder	Relevant existing instrument
Tri-layer geographic diversification	National telecom regulators; private cable consortia	EU Global Gateway; national broadband strategies
Regional digital-security alliance	National CERTs; NATO/EU cyber-defense bodies	CERT/CSIRT/FIRST cooperation frameworks
Sovereign/regional cloud fallback	Regional development banks; national digital ministries	EU Digital Decade programme; regional development funds
Legal clarity on repair-vessel transit	Foreign ministries; International Maritime Organization	UNCLOS; Montreux Convention diplomatic channels

been subjected to a formal evaluation of narrative fidelity to the underlying quantitative scenario; we regard this as an important direction for future validation work, discussed further in Section 7.

6.1. Mapping recommendations to responsible stakeholders

Table 3 maps each policy recommendation above onto the class of stakeholder best positioned to act on it, since a recurring weakness of regional-resilience studies is that recommendations are addressed to an undifferentiated “regional community” rather than to a specific institutional actor with the authority and budget to implement them.

6.2. Data provenance and ethical considerations

Two data-related caveats bear directly on how the results of Section 5 should be read. First, several of the fragility and intensity parameters used in the case study are derived from a qualitative strategic assessment rather than from a peer-reviewed, publicly reproducible incident dataset [5]; we have therefore favored reporting relative rankings and qualitative sensitivity (Section 4.5) over presenting the absolute connectivity-loss figures as precise forecasts, and we encourage readers to treat Table 2 as a structured prioritization aid rather than as a probabilistic forecast with formally calibrated confidence intervals. Second, because the framework explicitly represents state-level strategic behavior (through the agent-based response layer of Section 4.3), there is a risk that a sufficiently detailed public instantiation of the model could itself provide adversarial actors with a prioritized target list of high-impact infrastructure assets. We have therefore deliberately kept the case study in Section 5 at the level of previously and publicly documented vulnerability classes (landing stations, named regional hubs) rather than extending the model to asset-level targeting detail (for example, exact cable-repeater locations), and we recommend that any operational deployment of the framework for national planning purposes apply equivalent information-security handling to its outputs as it would to any other critical-infrastructure vulnerability assessment.

7. Conclusions and future works

This paper has argued that the splinternet should be understood not merely as a qualitative descriptive category but as a quantifiable, simulate-able trajectory of regional connectivity fragmentation, and it has proposed an AI-driven modeling framework – combining structural graph representation, Bayesian hazard propagation, agent-based strategic-response simulation, and constrained large-language-model scenario narration – for operationalizing that trajectory. Applied to the Black Sea basin, the framework reproduces and extends a detailed qualitative strategic assessment of regional connectivity risk, and it generates several quantitative and structural findings: that landing-station cyber compromise dominates instantaneous connectivity loss, that kinetic/legal interactions under the Montreux Convention dominate

expected recovery time, and that landlocked or single-corridor-dependent states such as Armenia are structurally the most fragile nodes in the regional network irrespective of their own domestic resilience investment.

Future work will proceed in three directions. First, we intend to calibrate hazard-channel intensities against a systematically compiled historical database of submarine-cable incidents (extending existing open efforts to track cable faults and suspected sabotage events), replacing the current qualitative calibration with an estimated statistical model. Second, we plan to extend the graph-structural layer to additional maritime chokepoint regions – the Baltic Sea, the Red Sea, and the Taiwan Strait – to test whether the qualitative findings reported here (SPOF dominance of cyber scenarios, legal-regime dominance of kinetic recovery time) generalize across geographies with different strait-law regimes. Third, we intend to conduct a structured human-evaluation study of the LLM-assisted narrative synthesis component, assessing both factual fidelity to the underlying quantitative scenario and the usefulness of the generated narratives for non-technical policy audiences, as a precondition for recommending the framework for operational, rather than purely research, use in national digital-resilience planning.

Acknowledgments

This work was supported by the Shota Rustaveli National Science Foundation of Georgia (SRNSFG) FR-25-4592.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

- [1] M. Mueller, *Will the Internet Fragment? Sovereignty, Globalization and Cyberspace*, Polity Press, Cambridge, 2017.
- [2] K. O'Hara, W. Hall, *Four Internets: Data, Geopolitics, and the Governance of Cyberspace*, Oxford University Press, 2021.
- [3] W. J. Drake, V. G. Cerf, W. Kleinwaechter, *Internet Fragmentation: An Overview*, Future of the Internet Initiative White Paper, World Economic Forum, 2016.
- [4] C. Bueger, T. Liebetrau, Critical maritime infrastructure protection: What's the trouble with submarine cables?, *Marine Policy* 132 (2021) 104711.
- [5] V. Svanadze, M. Iavich, G. Iashvili, The impact of threats in cyberspace on the process of internet fragmentation, in: *International Conference on Information and Software Technologies*, Springer Nature Switzerland, Cham, 2025.
- [6] T. S. Sechser, N. Narang, C. Talmadge, Emerging technologies and strategic stability in peacetime, crisis, and war, *Journal of Strategic Studies* 42 (2019) 727–735.
- [7] K. Leetaru, P. A. Schrod, Gdelt: Global data on events, location and tone, 1979–2012, in: *International Studies Association Annual Conference*, 2013.
- [8] N. Bapat, B. Doty, C. K. Hong, *Icews coded event data*, Harvard Dataverse, 2017.
- [9] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, G. Monfardini, The graph neural network model, *IEEE Transactions on Neural Networks* 20 (2009) 61–80.
- [10] J. M. Epstein, *Generative Social Science: Studies in Agent-Based Computational Modeling*, Princeton University Press, 2006.
- [11] M. C. Horowitz, L. Kahn, Bending the automation bias curve: A study of human and ai-based decision making in national security contexts, *International Studies Quarterly* 68 (2024).
- [12] International Cable Protection Committee, *Submarine cables and bbnj: Position paper on cable protection under international law*, 2023.
- [13] United Nations, *United nations convention on the law of the sea (unclos)*, 1982.

- [14] Convention regarding the regime of the straits (montreux convention), 1936.
- [15] European Commission, The global gateway strategy, joint communication, 2021.
- [16] H. Shen, Building a digital silk road? situating the internet in china's belt and road initiative, *International Journal of Communication* 12 (2018) 2683–2701.
- [17] S. Gnatyuk, B. Akhmetov, V. Kozlovskiy, D. Kanonik, Critical infrastructure cybersecurity risk assessment and bgp hijacking analysis, in: *Proceedings of the IEEE International Conference on Advanced Trends in Information Theory*, 2022.
- [18] TeleGeography, Submarine cable map and global bandwidth research service, 2025. 2025 edition.
- [19] Z. Hu, S. Gnatyuk, T. Okhrimenko, S. Tynymbayev, M. Iavich, High-speed and secure PRNG for cryptographic applications, *International Journal of Computer Network and Information Security* 12 (2020) 1–10. doi:10.5815/ijcnis.2020.03.01.
- [20] J. S. Al-Azzeh, M. Al Hadidi, R. S. Odarchenko, S. Gnatyuk, Z. Shevchuk, Z. Hu, Analysis of self-similar traffic models in computer networks, *International Review on Modelling and Simulations* 10 (2017).