

(RESEARCH ARTICLE)



A SELF-UPDATING FRAMEWORK FOR IDENTIFYING THE K-MOST VULNERABLE ENTITIES IN SMART GRIDS UNDER POWER-COMMUNICATION INTERDEPENDENCIES

Oyebode Abdul Salam ^{1,*}, Abdul Azeez ² and Abdul Rahman ³¹ Department of Fine and Applied ART, Ladoke Akintola University of Technology: Ogbomoso, Oyo State, Nigeria.² Department of Electrical and Electronics Engineering, Ladoke Akintola University of Technology: Ogbomoso, Oyo State, Nigeria.³ Department of Electrical Engineering, Faculty of Engineering, Ladoke Akintola University of Technology: Ogbomoso, Oyo State, Nigeria.

* Corresponding Author

ORCID Details

ABDUL SALAM OYEBODE: <https://orcid.org/0009-0009-9080-5462>ABDULAZEEZ OYEBODE: <https://orcid.org/0009-0003-9715-7142>ABDUL RAHMAN OYEBODE: <https://orcid.org/0009-0003-5077-2447>

Global Journal of Engineering and Technology Advances, 2026, 28(03), 197–211

Publication history: Received on 30 July 2026; revised on 10 September 2026; accepted on 12 September 2026

Article DOI: <https://doi.org/10.30574/gjeta.2026.28.3.0239>

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

Modern electricity networks depend on a communication layer for measurement, protection and control, and that communication layer draws its own power from the grid it supervises. This mutual dependence creates failure pathways that single-layer vulnerability analyses cannot capture, and it shifts continuously as topology, loading and control assignments evolve. This study develops and evaluates a self-updating framework that identifies the K entities whose simultaneous loss produces the largest interdependent damage. The framework couples a direct-current power-flow cascade simulator with a control-dependency propagation rule, and scores every node using an Interdependent Criticality Index that combines topological position, flow-based loading, interlayer support burden and dependency depth. Candidate sets are assembled by a lazy-greedy selection procedure, and a drift monitor based on Kendall's rank correlation and the Page-Hinkley test decides when the ranking must be revised. Revision is carried out incrementally over an affected set rather than by full recomputation. The framework was tested on four coupled systems built from the IEEE 14-, 39-, 118- and 300-bus test cases paired with synthetic communication topologies, across 500 operating snapshots per system and 1000 Monte Carlo replications. Measured against exhaustive ground truth, the framework achieved recall@5 between 0.85 and 0.94, compared with 0.38 to 0.52 for degree centrality and 0.66 to 0.80 for a static version of the same index. At K equal to 10 it produced 38.2 per cent demand not served on the coupled 118-bus system, exceeding betweenness-based selection by 10.9 percentage points, while incremental updating reduced update time from 41.6 s to 0.83 s.

Keywords: Smart grid vulnerability; Interdependent networks; Cascading failure; Critical node identification; Concept drift; Cyber-physical systems

1 INTRODUCTION

Electric power systems have ceased to be purely electromechanical assets. Their operation now rests on a supervisory layer whose entities depend on the power network at the same time as the power network depends on them, a coupling formalised by Roy and colleagues as a set of implicative interdependency relations spanning both layers [1]. That formalism has since been applied to the construction and maintenance of contingency lists for the joint system [2] and to the identification of the entities whose failure inflicts the greatest joint damage [3]. Wide-area measurement, remote switching, adaptive protection and market dispatch all depend on a supervisory control and data acquisition (SCADA) layer whose routers, remote terminal units and substation gateways are themselves electrically supplied by the network they govern [4,5]. The result is a bidirectional dependency: an electrical outage can silence the control equipment that would otherwise arrest it, and a communication outage can leave healthy plant unobservable and therefore effectively unavailable. Analyses that treat the power network and the communication network as separate objects systematically understate how far a disturbance can travel [6,7].

Interest in this problem accelerated after Buldyrev and colleagues showed that two mutually dependent networks fail through an abrupt first-order percolation transition rather than the graceful degradation seen in isolated networks [6]. Later work established that coupling strength, interlayer correlation and the availability of alternative supply paths all shape the transition point [8,9], and that mitigation effective in one layer may be counterproductive once the second is admitted [10]. Applied studies report the same qualitative behaviour in transmission systems paired with operational telecommunication assets [11,12]. Field evidence reinforces the concern: the 2015 attack on the Ukrainian distribution system combined remote breaker operation with denial of the channels needed to restore visibility, making the electrical and informational consequences inseparable [13,14].

Practical asset protection, however, rarely asks whether a system is vulnerable in the abstract. It asks which small set of assets deserves hardening, monitoring or redundant supply first. This is the K-most vulnerable entity problem, an interdependent analogue of the classical N-k security problem [15,16]. Roy and Sen formulated it on the joint power-communication network and solved it exactly by integer linear programming over the implicative interdependency model [3], and subsequently proposed regenerating the resulting K-contingency list as operating conditions evolve [2]. Two families of answers dominate. Topological ranking assigns each node a centrality score and selects the K highest [17,18]; such measures are inexpensive but correlate weakly with the electrical consequences of removal, since they ignore generation and demand placement, line ratings and transfer direction [19,20]. Optimisation-based ranking instead formulates the attacker's problem as a bilevel or mixed-integer program [15,21], which is electrically faithful but combinatorially expensive, and the expense grows sharply once a second layer is introduced.

A third limitation cuts across both families and has received less attention. Vulnerability rankings are almost always computed once, for a single reference operating state, and then treated as durable. Real systems do not oblige. Line outages for maintenance, seasonal load reshaping, distributed generation output, reconfiguration of communication routing tables and reassignment of control responsibility all change the interdependent structure on time scales of minutes to weeks. A ranking that was accurate at commissioning progressively decays, and the operator has no signal indicating when it has decayed enough to matter. In the machine learning literature this phenomenon is called concept drift, and a substantial body of technique exists for detecting and responding to it [22,23]; it has seldom been transferred to infrastructure vulnerability assessment.

The present study addresses these three gaps together. It proposes a self-updating framework in which the vulnerability ranking is treated as a maintained quantity rather than a one-off computation. The specific contributions are as follows. First, an Interdependent Criticality Index (ICI) is defined that fuses topological position, flow-based loading, interlayer support burden and dependency depth into a single node score applicable to both power and communication entities on a common scale. Second, the K-most vulnerable set is assembled using a lazy-greedy procedure that exploits the near-submodular structure of the interdependent damage function, giving near-exhaustive quality at a small fraction of exhaustive cost. Third, a drift monitor built from Kendall's rank correlation and the Page-Hinkley change detector determines when the ranking has degraded beyond an operator-specified tolerance, and an incremental updater then recomputes scores only over the affected subgraph. Fourth, the framework is evaluated on four coupled test systems against exhaustive ground truth, five ranking baselines, 500 operating snapshots per system and a full ablation of its components.

The remainder of this article is organised as follows. Section 2 describes the interdependent system model, the cascade simulator, the criticality index, the selection procedure, the self-updating mechanism and the experimental protocol. Section 3 reports identification accuracy, damage potential, computational scalability, drift-tracking behaviour, parameter sensitivity and component ablation, and discusses the operational implications and limitations of the findings. Section 4 concludes.

2 MATERIALS AND METHODS

2.1 Interdependent system representation

The coupled infrastructure at time t is represented as a two-layer graph

$$G(t) = (V_p \cup V_c, E_p \cup E_c \cup E\Phi) \quad (1)$$

where V_p is the set of power entities (buses, generators and loads), V_c is the set of communication entities (remote terminal units, substation gateways, routers and the control centre), E_p is the set of transmission lines, E_c is the set of communication links, and $E\Phi$ is the set of directed interlayer dependency arcs. Two dependency relations are distinguished. A supply arc runs from a power bus to a communication node that draws its electrical supply from that bus. A control arc runs from a communication node to the power entities whose measurement and switching functions it serves. The mapping Φ therefore encodes both relations, and the coupling ratio q is defined as the fraction of communication nodes that possess at least one supply arc to a power bus rather than to autonomous local generation.

Each power entity i carries an injection $P_i(t)$, a demand $D_i(t)$ and an operating status $s_i(t) \in \{0,1\}$. Each communication entity j carries a status $\sigma_j(t)$ and a service set Ψ_j listing the power entities it supervises. The notation used throughout is summarised in Table 1.

Table 1: Principal notation used in the framework.

Symbol	Definition
$G(t)$	Interdependent two-layer graph at operating snapshot t
$E\Phi$	Set of directed interlayer dependency arcs
q	Interlayer coupling ratio
α	Line tolerance parameter defining thermal headroom
$f\ell$	Active power flow on line ℓ
$C\ell$	Flow capacity of line ℓ
$ICl_i(t)$	Interdependent Criticality Index of entity i at snapshot t
S_k	Selected set of the K most vulnerable entities
τ	Kendall rank correlation between successive rankings
$\tau_{\min}$	Operator acceptance threshold on τ
λ	Page-Hinkley decision threshold
$A(t)$	Affected set requiring score recomputation

2.2 Cascading failure model

Electrical redistribution is modelled with the linearised direct-current power flow, which is standard in large-scale vulnerability studies because it retains the network effect of interest while remaining tractable over the very large number of scenarios required [24,25]. For a given injection vector P , the flow on line ℓ is

$$f\ell = \sum_i PTDF_{\ell i} \cdot P_i \quad (2)$$

where PTDF is the power transfer distribution factor matrix derived from the susceptance matrix of the surviving network. Line capacities are assigned proportionally to base-case flows through a tolerance parameter α ,

$$C\ell = (1 + \alpha) \cdot |f\ell^0| \quad (3)$$

following the load-capacity convention introduced by Motter and Lai and widely adopted since [26,27]. A line trips when its flow exceeds $C\ell$. After each tripping round the network is re-islanded, generation and demand are rebalanced within each surviving island subject to generator limits, and the flow solution is recomputed. The comparative behaviour of alternating-current and direct-current representations under failure has been examined in detail elsewhere; the direct-current model reproduces the ordering of severe events adequately for ranking purposes [28,29].

Communication propagation is superimposed on this electrical cascade. A communication node j remains operational only if it retains electrical supply through at least one live supply arc, retains a path to the control centre within the surviving communication layer, and has not been isolated by the loss of its own physical link set. Formally,

$$\sigma_j(t+1) = 1 \Leftarrow [\exists i \in \Phi^{-1}(j) : s_i(t) = 1] \wedge [j \leftrightarrow c_0 \text{ in } G_c(t)] \quad (4)$$

where c_0 denotes the control centre. When σ_j falls to zero, every power entity in the service set Ψ_j loses remote observability and control. Such entities are not de-energised, but they are frozen: no corrective switching, no automatic generation control participation and no adaptive protection setting change can be applied to them for the remainder of the cascade. This treatment reflects the operational reality that unobservable plant cannot be relied upon for corrective action, and it is the mechanism through which communication loss amplifies electrical damage. The two propagation processes are iterated in alternation until no further status change occurs.

The severity of a scenario is expressed as normalised demand not served,

$$\text{DNS}(S) = 100 \cdot [\sum_i D_i(0) - \sum_i D_i(\infty)] / \sum_i D_i(0) \quad (5)$$

where $D_i(\infty)$ is the demand served once the coupled cascade has settled. The interdependent damage function used for selection is $I(S) = E[\text{DNS}(S)]$, the expectation taken over the stochastic elements of the cascade.




Figure 1: Architecture of the self-updating vulnerability identification framework, showing the six functional modules and the closed feedback loop through which drift detection triggers incremental score revision.

2.3 The Interdependent Criticality Index

Exhaustive evaluation of $I(S)$ over all candidate sets is infeasible beyond very small systems, so a per-entity score is required to prune the search. The Interdependent Criticality Index of entity i at snapshot t is defined as a convex combination of four normalised components,

$$\text{ICI}_i(t) = w_1 \check{c}_i(t) + w_2 \mathfrak{l}_i(t) + w_3 \delta_i(t) + w_4 \dot{r}_i(t), \quad \sum w_k = 1 \quad (6)$$

The structural term $\check{c}_i$ is the betweenness centrality of i within its own layer, normalised to the unit interval and computed with the Brandes algorithm [30]. It captures the extent to which the entity lies on shortest transfer or routing paths.

The flow term $\mathfrak{l}_i$ measures electrical stress attributable to the entity. For a power bus it is the normalised sum of the loading ratios of incident lines,

$$\mathfrak{l}_i = (1 / |\ell(i)|) \sum_{\ell \in \ell(i)} |f_\ell| / C_\ell \quad (7)$$

For a communication node it is the mean of the same quantity taken over the power entities in its service set, which places both layers on a comparable scale.

The interlayer term δ_i quantifies the burden the entity carries across the coupling. For a power bus it is the fraction of communication nodes whose only supply arcs terminate at that bus; for a communication node it is the fraction of power entities in its service set that have no alternative supervisory node. This term is what distinguishes the index from single-layer measures, and it is the component that identifies entities of modest electrical significance whose loss nonetheless disables a disproportionate share of the control layer.

The dependency-depth term $\hat{r}_i$ is the expected number of alternating propagation rounds initiated by removing i alone, estimated from a pilot of 50 single-removal cascade simulations. Entities seeding long chains score highly even when their neighbourhood is unremarkable.

The weight vector w is initialised uniformly and then adapted online. After each snapshot in which ground-truth damage is observed, either from a real disturbance or from a scheduled simulation batch, the weights are revised by exponentially weighted gradient descent on the squared ranking error,

$$w_k \leftarrow w_k - \eta \cdot \partial / \partial w_k [\sum_i (ICI_i - \hat{I}_i)^2], \text{ followed by projection onto the simplex} \quad (8)$$

where $\hat{I}_i$ is the observed normalised damage attributable to i and η is a forgetting factor. This allows the index to specialise to the particular system rather than relying on weights transferred from the literature.




Figure 2: Two-layer representation of the coupled test system. Circles denote power buses, squares denote communication nodes, solid lines denote intra-layer edges and dashed red lines denote interlayer dependency arcs.

2.4 Selection of the K-most vulnerable set

Given the index, the K-most vulnerable set is the solution of

$$S_k^* = \arg \max \{ I(S) : S \subseteq V, |S| = K \} \quad (9)$$

The damage function is monotone and, on the coupled systems examined here, empirically near-submodular: the marginal contribution of an entity declines as the selected set grows, because damage regions increasingly overlap. Exact submodularity is not guaranteed in the presence of islanding, so the greedy solution is not accompanied by the classical $(1 - 1/e)$ guarantee [31]; nevertheless the structure is close enough that greedy selection with lazy evaluation performs strongly in practice, as Section 3 shows.

Selection proceeds as follows. Entities are ordered by ICI and the top 3K form a candidate pool, which reduces the number of cascade simulations by an order of magnitude relative to searching over V . The set is then built one element at a time, at each step adding the entity with the largest marginal gain

$$\Delta(v | S) = I(S \cup \{v\}) - I(S) \quad (10)$$

Marginal gains are held in a priority queue and re-evaluated lazily in the manner of the CELF procedure, so that an entity is re-simulated only when its stale gain still tops the queue [32]. Each evaluation of I uses 200 Monte Carlo cascade replications with randomised initial loading within ± 5 per cent of the reference schedule.

2.5 Self-updating mechanism

The framework treats the ranking as a maintained artefact. At every operating snapshot the current ICI vector is compared with the vector held in memory using Kendall's rank correlation coefficient [33],

$$\tau(t) = (n^c - n^d) / [\frac{1}{2} n (n - 1)] \quad (11)$$

where n^c and n^d are the counts of concordant and discordant pairs. A ranking is considered stale when $\tau(t)$ falls below the operator threshold $\tau_{\min}$, set to 0.80 in the experiments reported here.

Because τ fluctuates with ordinary operational noise, a threshold crossing alone is a poor trigger. A Page-Hinkley test is therefore run on the sequence of τ values [34]. Writing m_t for the cumulative deviation of τ below its running mean and M_t for its minimum to date, an update is triggered when

$$PH(t) = m_t - M_t > \lambda \quad (12)$$

with λ calibrated on a two-week burn-in period to yield at most one false trigger per 200 snapshots. The combination of a level threshold and a change detector suppresses reaction to transient excursions while responding promptly to sustained structural change.

When an update is triggered, the framework does not recompute the whole index. It identifies an affected set

$$A(t) = \{ v \in V : \text{dist}(v, \Delta E(t)) \leq h \} \cup \Phi_{\pm}(\Delta E(t)) \quad (13)$$

comprising all entities within h hops of a changed edge or status in either layer, together with the interlayer images of those entities. The hop radius h is set to 2, a value determined by preliminary testing to capture more than 95 per cent of the score movement while leaving the affected set small. Betweenness contributions are updated only over the affected set using the incremental scheme of Bergamini and colleagues [35], flow terms are recomputed from the updated PTDF matrix, and dependency-depth terms are re-estimated from a reduced pilot of 20 simulations restricted to affected seeds. Entities outside $A(t)$ retain their previous scores. The full loop, from layered input through scoring, selection, drift detection and incremental revision, is shown in Figure 1.

2.6 Test systems and experimental design

Four coupled systems were constructed. Power layers were taken from the IEEE 14-, 39-, 118- and 300-bus test cases distributed through the University of Washington archive and implemented in MATPOWER [36,37]. Communication layers were generated synthetically, since operational telecommunication topologies are rarely published: each communication layer was built as a Watts-Strogatz small-world graph with rewiring probability 0.15 and mean degree 4, a structure consistent with reported utility wide-area network designs [38]. Interlayer supply arcs were assigned by geographical proximity in the published bus coordinates, and control arcs were assigned so that each communication node supervises between two and six power entities. System characteristics appear in Table 2.

Table 2: Characteristics of the four coupled test systems.

System	Power nodes	Power lines	Comm. nodes	Comm. links	Interlayer arcs	Total V
Coupled IEEE 14	14	20	12	24	26	26
Coupled IEEE 39	39	46	30	60	69	69
Coupled IEEE 118	118	186	95	190	213	213
Coupled IEEE 300	300	411	240	480	540	540

For each system, 500 consecutive operating snapshots were generated to represent one simulated quarter of operation at approximately four-hour resolution. Snapshots incorporated diurnal and weekly demand shaping, scheduled line outages at a rate of 1.2 per 100 snapshots, stochastic renewable injection at four buses, and periodic reassignment of communication routing. Five permanent structural change events were injected at snapshots 95, 188, 262, 351 and 437 to test drift response.

Simulation parameters are listed in Table 3. All experiments were repeated over 1000 Monte Carlo replications and executed in Python 3.11 using NetworkX for graph operations and NumPy for linear algebra [39], on a workstation with an 8-core 3.6 GHz processor and 32 GB of memory. No graphics acceleration was used.

Table 3: Simulation parameters.

Parameter	Symbol	Value
Line tolerance	α	0.30 (0.10–0.80 in sensitivity study)
Interlayer coupling ratio	q	0.70 (0.20–1.00 in sensitivity study)
Kendall acceptance threshold	$\tau_{\min}$	0.80
Page-Hinkley threshold	λ	0.045
Affected-set hop radius	h	2
Forgetting factor	η	0.05
Candidate pool size	—	3K
Monte Carlo replications per evaluation	—	200
Replications per reported result	—	1000
Operating snapshots per system	—	500

2.7 Baselines and evaluation metrics

Five baselines were implemented. Random selection provides a lower reference. Degree centrality, PageRank and betweenness centrality represent the topological family and were computed on the union graph so that they had access to the same structural information as the proposed index [17,40]. Static ICI applies the full index but computes it once at snapshot zero and never revises it; comparison with this baseline isolates the contribution of the self-updating mechanism from that of the scoring rule. For the 14- and 39-bus systems, exhaustive enumeration of all sets of size five was tractable and provided ground truth.

Four metrics were recorded. Normalised demand not served, defined in Equation (5), measures damage potential. Recall@K is the proportion of the exhaustive ground-truth set recovered by a method,

$$\text{Recall@K} = |S_k \cap S_k^{GT}| / K \quad (14)$$

Kendall's τ against a freshly recomputed ranking measures ranking currency over time. Wall-clock update time measures computational cost. Differences between methods were assessed with the Wilcoxon signed-rank test on paired replications, with significance declared at $p < 0.01$ and Holm correction applied across the family of pairwise comparisons; the non-parametric test was chosen because DNS distributions are strongly right-skewed and heavy-tailed, a well-documented property of cascade severity [41].

3 RESULTS AND DISCUSSION

3.1 Damage potential of the identified sets

Figure 3 reports normalised demand not served on the coupled IEEE 118-bus system as the size of the removed set increases from one to twenty. All methods rise steeply at small K and flatten thereafter, the signature of the diminishing marginal returns that motivate greedy selection. The separation between methods is nevertheless substantial and persistent.




Figure 3: Normalised demand not served as a function of the number of removed entities on the coupled IEEE 118-bus system. Each point is the mean of 1000 Monte Carlo replications.

At K equal to 10, the proposed framework produced 38.2 per cent demand not served, against 31.6 per cent for the static version of the same index, 27.3 per cent for betweenness, 24.9 per cent for PageRank, 21.7 per cent for degree and 8.4 per cent for random selection. The margin over betweenness, the strongest purely topological competitor, is 10.9 percentage points, and the margin over the static index is 6.6 percentage points. All pairwise differences were significant at $p < 0.001$ under the Wilcoxon signed-rank test with Holm correction. Table 4 extends the comparison to all four systems.

Table 4: Normalised demand not served (per cent) at $K = 10$, averaged over 1000 replications. Standard deviations in parentheses.

Method	IEEE 14	IEEE 39	IEEE 118	IEEE 300
Random	11.2 (4.8)	9.6 (3.9)	8.4 (3.1)	7.1 (2.6)
Degree centrality	26.4 (5.1)	23.8 (4.4)	21.7 (3.8)	19.3 (3.4)
PageRank	29.1 (4.9)	26.5 (4.2)	24.9 (3.6)	22.4 (3.2)
Betweenness centrality	31.8 (4.7)	29.4 (4.0)	27.3 (3.5)	25.1 (3.1)
Static ICI	35.6 (4.2)	33.7 (3.6)	31.6 (3.2)	29.4 (2.9)
SU-ICI (proposed)	42.7 (3.8)	40.1 (3.3)	38.2 (2.9)	35.8 (2.7)

Two features of Table 4 merit comment. The advantage of the proposed framework is preserved across the range of system sizes examined, and absolute values decline gently with size in every row, as expected since larger systems offer more alternative transfer paths. The ordering of the baselines is also stable: flow-aware and interdependency-aware scoring consistently outperforms structure-only scoring, corroborating earlier findings that topological proxies alone are unreliable predictors of electrical consequence [19,20] and extending them to the two-layer setting.

Inspection of the selected sets clarifies the mechanism. On the coupled 118-bus system, the proposed framework selected on average 3.4 communication entities within its ten-element set, whereas betweenness selected 1.1 and degree selected 0.6. The communication entities chosen were rarely the highest-degree nodes in their layer; they were nodes that acted as sole supervisory parents for clusters of generation-adjacent buses. Their removal did not directly de-energise anything, but it froze corrective control precisely where the cascade needed to be arrested. This is the class of vulnerability that single-layer analysis is structurally incapable of detecting, and it accounts for most of the margin reported above.

3.2 Identification accuracy against ground truth

Damage potential shows that a method finds a damaging set, but not whether it finds the right one. Figure 4 reports recall@5 against exhaustively enumerated ground truth. Exhaustive enumeration was completed for the 14- and 39-bus systems; for the 118- and 300-bus systems a restricted ground truth was constructed by exhaustive enumeration over the union of all methods' top-40 candidates, which bounds the achievable recall from below.




Figure 4: Recall@5 against exhaustive ground truth for four coupled test systems. Bars show means over 1000 replications.

The proposed framework recovered between 0.85 and 0.94 of the ground-truth set, the static index between 0.66 and 0.80, betweenness between 0.49 and 0.64, and degree between 0.38 and 0.52. The gap between the proposed framework and the static index, which differ only in whether the ranking is maintained, averaged 0.19 across the four systems. That figure is the clearest single quantification of the cost of treating a vulnerability ranking as a fixed asset. Table 5 summarises the accuracy metrics.

Table 5: Identification accuracy and ranking currency across the four coupled systems, $K = 5$.

Method	Recall@5	Precision@5	Mean τ over 500 snapshots	Snapshots below $\tau_{\min}$
Degree centrality	0.44	0.44	0.51	382
PageRank	0.50	0.50	0.56	351
Betweenness centrality	0.57	0.57	0.59	334
Static ICI	0.73	0.73	0.62	298
SU-ICI (proposed)	0.90	0.90	0.93	11

3.3 Computational scalability

Figure 5 compares the wall-clock cost of a full ranking recomputation with that of an incremental update, over systems ranging from 26 to 2600 total nodes. The two largest points were obtained from synthetic coupled systems constructed by the same procedure in order to probe scaling behaviour beyond the standard test cases.




Figure 5 :Wall-clock time per ranking update for full recomputation and for incremental affected-set updating, on logarithmic axes. Annotations give the speed-up factor at each system size.

Full recomputation scales approximately as the cube of node count, dominated by the repeated betweenness and PTDF computations. Incremental updating scales close to quadratically, because the affected set grows far more slowly than the system. At 540 nodes, the size of the coupled 300-bus system, full recomputation required 41.6 s against 0.83 s for the incremental scheme, a speed-up of 50 times. At 2600 nodes the speed-up reached 81 times. Table 6 gives the underlying figures together with the mean affected-set size.

Table 6: Update cost and affected-set size by system scale.

Total nodes $ V $	Full recomputation (s)	Incremental update (s)	Speed-up	Mean $ A(t) $	$ A(t) / V $ (%)
26	0.21	0.012	18×	7.4	28.5
69	0.94	0.031	30×	13.1	19.0
213	6.80	0.142	48×	27.6	13.0
540	41.6	0.830	50×	51.3	9.5
1200	205.0	3.10	66×	88.7	7.4
2600	921.0	11.4	81×	152.4	5.9

These figures move the analysis from an offline planning activity to something that can run alongside the control room. An update completed in under a second on a 540-node system can be executed after every topology change, whereas a 42-second recomputation would in practice be scheduled nightly at best and would be stale for most of the operating day. The declining ratio of affected-set size to system size, from 28.5 per cent at 26 nodes to 5.9 per cent at 2600, indicates that the advantage widens rather than erodes as systems grow.

3.4 Drift tracking behaviour

Figure 6 traces Kendall's τ between each method's held ranking and a freshly recomputed reference over the 500-snapshot horizon of the coupled 118-bus system, with the five injected structural change events marked.




Figure 6: Ranking currency over 500 operating snapshots. Vertical dashed lines mark injected structural change events; the horizontal line marks the acceptance threshold $\tau_{\min} = 0.80$.

The static index decays monotonically, crossing the acceptance threshold at snapshot 128 and reaching 0.42 by the end of the horizon. Its ranking was therefore unacceptable for roughly three-quarters of the simulated quarter, and an operator relying on it would have been protecting the wrong assets for most of that period without any indication that this was so. The proposed framework held τ above 0.86 throughout, apart from brief excursions immediately following each injected event.

Those excursions are informative. After each structural change the framework's ranking degraded for a mean of 9.2 snapshots, roughly a day and a half of simulated time, before the monitor fired and the incremental update restored currency. This latency is the price of requiring a sustained signal before acting, and is tunable: reducing λ shortens it at the cost of more frequent updates. Across the horizon the monitor fired 11 times against 5 genuine events, a false-positive rate of 1.2 per 100 snapshots, consistent with the burn-in calibration target. Since a false positive costs under one second of computation, the asymmetry between an unnecessary update and a stale ranking is the central practical argument for the design.

3.5 Sensitivity to coupling and tolerance

The behaviour of interdependent systems is known to depend strongly on how tightly the layers are bound and on how much thermal headroom the electrical layer carries [8,9]. Figure 7 maps demand not served at K equal to 10 across the coupling ratio q and the tolerance parameter α .




Figure 7: Demand not served at $K = 10$ as a function of interlayer coupling ratio and line tolerance parameter, coupled IEEE 118-bus system. Contours are labelled in per cent.

Damage increases monotonically with coupling and decreases monotonically with tolerance, as expected, but the surface is not separable. In the tightly coupled, low-headroom corner (q above 0.8, α below 0.2) demand not served exceeded 45 per cent, while in the loosely coupled, generously rated corner it fell below 10 per cent. The gradient with respect to q is steepest at low α , which has a direct operational reading: where thermal margins are thin, reducing the dependence of communication equipment on grid-supplied power delivers disproportionate benefit. Conversely, in a

system with generous line ratings, investment in autonomous supply for communication assets buys comparatively little. This interaction has practical consequences for how hardening budgets should be split between the two layers, and it is invisible to any analysis that fixes one parameter while varying the other.

Coupling ratios at the upper end of the range examined are more representative of distribution-level installations than of transmission systems, where backup supply for critical communication assets is common practice. The framework's advantage over the baselines was preserved throughout the lower, more realistic half of the range.

3.6 Component ablation

Figure 8 reports the contribution of each component by removing it in turn from the full framework and re-running the coupled 118-bus experiment.




Figure 8: Ablation of framework components. Bars give demand not served at $K = 10$ (left axis); the dashed trace gives mean update time on a logarithmic right axis.

Removing the interlayer term from the index caused the largest loss of damage potential, from 38.2 to 27.4 per cent, confirming that the explicit representation of cross-layer support burden is the single most valuable element of the scoring rule. Removing the flow term cost 8.4 percentage points, and removing the drift monitor cost 6.3 percentage points. Removing lazy evaluation left damage potential essentially unchanged, at 37.9 per cent, but increased update time from 0.83 s to 6.42 s, a factor of 7.7. Lazy evaluation is therefore purely an efficiency device, as its theoretical basis implies, whereas the other three components carry accuracy consequences [32].

The ordering of these contributions supports the central claim of the study. The two components that most distinguish the framework from prior work, namely explicit interlayer scoring and maintained ranking currency, together account for 17.2 of the 38.2 percentage points of damage potential achieved. Neither is available in a conventional single-layer, single-shot vulnerability analysis.

3.7 Implications and limitations

Taken together, the results indicate that a vulnerability ranking should be treated as a maintained operational quantity rather than a planning-study output. The framework identifies asset sets that are both more damaging and more accurate than established alternatives, at a computational cost that permits revision after every topology change on systems of realistic size. For an operator this changes the character of the output: instead of a list produced annually, the ranking becomes a live indicator whose staleness is itself monitored.

The finding that communication entities occupy roughly a third of the identified vulnerable sets, while being almost invisible to topological ranking, has a direct bearing on protection investment. It suggests that hardening programmes allocated on the basis of electrical significance alone will systematically underfund the supervisory layer, and specifically will underfund the unglamorous single-parent supervisory nodes that carry no traffic distinction but hold sole control responsibility for consequential plant.

Several limitations qualify these conclusions. The direct-current power flow neglects reactive power and voltage collapse, which are implicated in a number of historical blackouts; alternating-current representations produce somewhat different failure sequences, although reported comparisons suggest the ordering of severe events is broadly

preserved [28,29]. The communication layers are synthetic, because operational topologies are not publicly available for security reasons, and their small-world construction may not reflect the more hierarchical designs used by some utilities. The control-loss model is binary, treating a supervised entity as either fully controllable or fully frozen, whereas real degradation is graded and protocol-dependent. The near-submodularity of the damage function is empirical rather than proved, so the greedy procedure carries no formal approximation guarantee under islanding. Finally, the drift experiments use injected structural events with known timing; performance against the more ambiguous drift patterns of real operational records remains to be established. Extension to alternating-current modelling, validation on utility communication topologies under non-disclosure, and incorporation of graded control degradation are the natural next steps.

4 CONCLUSION

This study set out to identify the K entities whose simultaneous loss inflicts the greatest damage on a smart grid when power and communication interdependencies are represented explicitly, and to keep that identification current as the system changes. The framework developed for the purpose combines an Interdependent Criticality Index spanning topological, flow-based, interlayer and dependency-depth information with lazy-greedy set selection, drift detection through Kendall's rank correlation and the Page-Hinkley test, and incremental score revision confined to an affected subgraph.

Across four coupled systems derived from standard test cases, the framework recovered 85 to 94 per cent of exhaustively determined ground-truth sets, against 38 to 52 per cent for degree centrality and 66 to 80 per cent for an otherwise identical index computed once and never revised. It produced 38.2 per cent demand not served at ten removals on the coupled 118-bus system, exceeding the strongest topological baseline by 10.9 percentage points. Incremental updating cut the cost of a revision from 41.6 s to 0.83 s at 540 nodes, and ranking currency was held above 0.86 over a horizon on which a static ranking decayed to 0.42.

Two findings carry the greatest practical weight. Communication entities constituted roughly a third of the identified vulnerable sets while being largely invisible to structural ranking measures, which implies that hardening decisions taken on electrical significance alone will misallocate protection. And maintaining ranking currency contributed almost as much to identification accuracy as the design of the scoring rule itself, which implies that the timing of vulnerability analysis deserves the attention normally reserved for its formulation. The computational profile of the framework makes such maintenance practical rather than aspirational, placing continuous interdependent vulnerability assessment within reach of the control room.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Acknowledgments

The authors thank the Department of Electrical and Electronics Engineering, University of Ilorin, for providing computing facilities used in the simulation study, and the anonymous reviewers whose comments improved the presentation of the drift-detection analysis. No external funding was received for this research.

Funding Source

This research received no external funding.

Disclosure of conflict of interest

The authors declare that they have no financial or non-financial conflicts of interest relating to the research reported in this article. No author holds a consultancy, patent application, equity position or other relationship with any organisation whose interests could be affected by the findings presented.

REFERENCES

- [1] Roy S, Chandrasekaran H, Pal A, Sen A. A new model to analyze power and communication system intra- and inter-dependencies. In: Proceedings of the 2020 IEEE Conference on Technologies for Sustainability (SusTech); 2020 Apr 23–26; Santa Ana, CA. Piscataway (NJ): IEEE; 2020. p. 1–8.
- [2] Roy S, Sen A. A self-updating K -contingency list for smart grid system. In: Proceedings of the 2021 IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC); 2021 Jan 27–30; Las Vegas, NV. Piscataway (NJ): IEEE; 2021. p. 341–6.

- [3] Roy S, Sen A. Identification of the K-most vulnerable entities in a smart grid system. In: Proceedings of the 2020 3rd International Conference on Advanced Communication Technologies and Networking (CommNet); 2020 Sep 4–6; Marrakech, Morocco. Piscataway (NJ): IEEE; 2020. p. 1–6.
- [4] Rinaldi SM, Peerenboom JP, Kelly TK. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Syst Mag.* 2001;21(6):11–25.
- [5] Sridhar S, Hahn A, Govindarasu M. Cyber-physical system security for the electric power grid. *Proc IEEE.* 2012;100(1):210–24.
- [6] Buldyrev SV, Parshani R, Paul G, Stanley HE, Havlin S. Catastrophic cascade of failures in interdependent networks. *Nature.* 2010;464(7291):1025–8.
- [7] Ouyang M. Review on modeling and simulation of interdependent critical infrastructure systems. *Reliab Eng Syst Saf.* 2014;121:43–60.
- [8] Parshani R, Buldyrev SV, Havlin S. Interdependent networks: reducing the coupling strength leads to a change from a first to second order percolation transition. *Phys Rev Lett.* 2010;105(4):048701.
- [9] Gao J, Buldyrev SV, Stanley HE, Havlin S. Networks formed from interdependent networks. *Nat Phys.* 2012;8(1):40–8.
- [10] Brummitt CD, D'Souza RM, Leicht EA. Suppressing cascades of load in interdependent networks. *Proc Natl Acad Sci USA.* 2012;109(12):E680–9.
- [11] Rosato V, Issacharoff L, Tiriticco F, Meloni S, De Porcellinis S, Setola R. Modelling interdependent infrastructures using interacting dynamical models. *Int J Crit Infrastruct.* 2008;4(1–2):63–79.
- [12] Johansson J, Hassel H. An approach for modelling interdependent infrastructures in the context of vulnerability analysis. *Reliab Eng Syst Saf.* 2010;95(12):1335–44.
- [13] Case DU. Analysis of the cyber attack on the Ukrainian power grid. Washington (DC): Electricity Information Sharing and Analysis Center; 2016.
- [14] Ten CW, Liu CC, Manimaran G. Vulnerability assessment of cybersecurity for SCADA systems. *IEEE Trans Power Syst.* 2008;23(4):1836–46.
- [15] Bienstock D, Verma A. The N-k problem in power grids: new models, formulations, and numerical experiments. *SIAM J Optim.* 2010;20(5):2352–69.
- [16] Salmeron J, Wood K, Baldick R. Analysis of electric grid security under terrorist threat. *IEEE Trans Power Syst.* 2004;19(2):905–12.
- [17] Freeman LC. A set of measures of centrality based on betweenness. *Sociometry.* 1977;40(1):35–41.
- [18] Holme P, Kim BJ, Yoon CN, Han SK. Attack vulnerability of complex networks. *Phys Rev E.* 2002;65(5):056109.
- [19] Hines P, Cotilla-Sanchez E, Blumsack S. Do topological models provide good information about electricity infrastructure vulnerability? *Chaos.* 2010;20(3):033122.
- [20] Cotilla-Sanchez E, Hines PDH, Barrows C, Blumsack S. Comparing the topological and electrical structure of the North American electric power infrastructure. *IEEE Syst J.* 2012;6(4):616–26.
- [21] Arianos S, Bompard E, Carbone A, Xue F. Power grid vulnerability: a complex network approach. *Chaos.* 2009;19(1):013119.
- [22] Gama J, Žliobaitė I, Bifet A, Pechenizkiy M, Bouchachia A. A survey on concept drift adaptation. *ACM Comput Surv.* 2014;46(4):1–37.
- [23] Bifet A, Gavaldà R. Learning from time-changing data with adaptive windowing. In: Proceedings of the SIAM International Conference on Data Mining; 2007 Apr 26–28; Minneapolis, MN. Philadelphia: SIAM; 2007. p. 443–8.
- [24] Soltan S, Mazaauric D, Zussman G. Analysis of failures in power grids. *IEEE Trans Control Netw Syst.* 2017;4(2):288–300.
- [25] Vaiman M, Bell K, Chen Y, Chowdhury B, Dobson I, Hines P, et al. Risk assessment of cascading outages: methodologies and challenges. *IEEE Trans Power Syst.* 2012;27(2):631–41.
- [26] Motter AE, Lai YC. Cascade-based attacks on complex networks. *Phys Rev E.* 2002;66(6):065102.
- [27] Crucitti P, Latora V, Marchiori M. Model for cascading failures in complex networks. *Phys Rev E.* 2004;69(4):045104.

- [28] Cetinay H, Soltan S, Kuipers FA, Zussman G, Van Mieghem P. Comparing the effects of failures in power grids under the AC and DC power flow models. *IEEE Trans Netw Sci Eng*. 2018;5(4):301–12.
- [29] Guo H, Zheng C, Iu HHC, Fernando T. A critical review of cascading failure analysis and modeling of power system. *Renew Sustain Energy Rev*. 2017;80:9–22.
- [30] Brandes U. A faster algorithm for betweenness centrality. *J Math Sociol*. 2001;25(2):163–77.
- [31] Nemhauser GL, Wolsey LA, Fisher ML. An analysis of approximations for maximizing submodular set functions—I. *Math Program*. 1978;14(1):265–94.
- [32] Leskovec J, Krause A, Guestrin C, Faloutsos C, VanBriesen J, Glance N. Cost-effective outbreak detection in networks. In: *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*; 2007 Aug 12–15; San Jose, CA. New York: ACM; 2007. p. 420–9.
- [33] Kendall MG. A new measure of rank correlation. *Biometrika*. 1938;30(1–2):81–93.
- [34] Page ES. Continuous inspection schemes. *Biometrika*. 1954;41(1–2):100–15.
- [35] Bergamini E, Meyerhenke H, Staudt CL. Approximating betweenness centrality in large evolving networks. In: *Proceedings of the 17th Workshop on Algorithm Engineering and Experiments*; 2015 Jan 5; San Diego, CA. Philadelphia: SIAM; 2015. p. 133–46.
- [36] Christie RD. Power systems test case archive. Seattle: University of Washington, Department of Electrical Engineering; 1999.
- [37] Zimmerman RD, Murillo-Sánchez CE, Thomas RJ. MATPOWER: steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Trans Power Syst*. 2011;26(1):12–9.
- [38] Watts DJ, Strogatz SH. Collective dynamics of small-world networks. *Nature*. 1998;393(6684):440–2.
- [39] Hagberg AA, Schult DA, Swart PJ. Exploring network structure, dynamics, and function using NetworkX. In: *Proceedings of the 7th Python in Science Conference*; 2008 Aug 19–24; Pasadena, CA. 2008. p. 11–5.
- [40] Page L, Brin S, Motwani R, Winograd T. The PageRank citation ranking: bringing order to the web. Stanford (CA): Stanford InfoLab; 1999. Technical Report 1999-66.
- [41] Dobson I, Carreras BA, Lynch VE, Newman DE. Complex systems analysis of series of blackouts: cascading failure, critical points, and self-organization. *Chaos*. 2007;17(2):026103.