Failure-Aware and Probabilistic Post-Disaster Restoration in Elastic Optical Networks using TVGs

Eonassis O. Santos*, Juliana de Santi†, Egemen Erbayat‡, Rujia Zou‡,
Sabidur Rahman§, Suresh Subramaniam‡, Gustavo B. Figueiredo*
*Federal University of Bahia, Brazil. Email:{eonassissantos, gustavobf}@ufba.br
†Federal University of Technology - Paraná, Brazil. Email: jsanti@utfpr.edu.br,
‡The George Washington University, USA. Email: {erbayat, rjzou, suresh}@gwu.edu
§AT&T Labs, San Ramon, California, USA. Email: k.m.sabidur.rahman@att.com

***Abstract*—Cascading failures in Elastic Optical Networks (EONs) severely threaten service availability during large-scale disasters, while traditional restoration strategies struggle to handle their dynamic and probabilistic nature. This work proposes a restoration approach that integrates probabilistic risk modeling into a Time-Varying Graph (TVG) framework. By associating failure probability, distance to disaster zones, and hop count with each network edge, composite cost functions guide risk-aware path selection. The method relies on a precomputed pool of k candidate paths, dynamically evaluating their feasibility and risk as the network degrades. Simulation results show that the proposed Time-Varying Graph Disaster Recovery with Probabilities (TVG-DRP) significantly improves restorability, availability, and resilience to re-affectation compared to three baselines: PRWR, PRPA, and TVG-DR.**

***Index Terms*—EONs, Cascading Failures, Disaster Recovery, Restoration, Time-Varying Graphs, Probabilistic Routing**

I. Introduction

Ensuring network resilience against large-scale failures and disasters has become a critical challenge in modern communication infrastructures. Such disruptive events can cause simultaneous link outages, severely impacting transmission in networks. To address these challenges, several restoration strategies for Elastic Optical Networks (EONs) have been proposed to mitigate the effects of disasters in optical networks.

Previous works often rely on static models, predefined protection resources, or reactive mechanisms based solely on topological criteria [1], [2]. Classical methods use path re-computation and spectrum reallocation, focusing on minimizing path length or reconfiguration delay [3], [4]. Also, several studies formulate the problem through Integer Linear Programming (ILP) or heuristics to address dynamic spectrum constraints. For instance, [5] formulates an ILP model for survivable EONs by integrating modulation-format awareness into protection schemes, improving resource efficiency and service resilience. A comprehensive overview of survivability strategies in EONs is given in [6].

While effective in isolated failure scenarios, these approaches do not capture the time-dependent evolution of failures nor incorporate probabilistic information about risk propagation. Consequently, they may incur high recomputation costs and are prone to selecting restoration paths that become quickly re-affected as the failure spreads.

To address this, DRAMA-PRO+ [7] incorporates mitigation-awareness for translucent EONs. The PRPA approach introduced in [2] proposes a restoration strategy that penalizes links near disaster zones to avoid paths likely to fail as disruptions propagate. Although these methods improve survivability, they largely depend on static penalty zones or introduce probabilistic risk only at relocation, without fully exploiting temporal models or precomputed paths to accelerate restoration.

These methods assume deterministic conditions and static topologies, thereby overlooking the progressive degradation and cascading effects that arise in large-scale disaster scenarios. Cascading failures represent a more complex failure model, where an initial disruption may disable core links and then propagate through overload, rerouting pressure, or loss of connectivity in surviving components. These failures unfold over time, affecting not just connectivity but also the predictability of service restoration.

This work adopts an *event-driven* perspective for modeling network evolution under large-scale disasters, where the network state is represented as a sequence of graph snapshots triggered by cascading failure events rather than fixed time intervals. Each failure event induces a structural change in the topology, and whenever an active path becomes infeasible, restoration is performed immediately by selecting an alternative route from a precomputed candidate pool. Consequently, the effectiveness of the proposed approach depends on the ordering and progression of failure events, rather than on a specific time span between snapshots. In scenarios where cascading failures lead to *progressive (monotonic) network degradation*, computing the candidate paths once on the initial graph G_0 is sufficient, since failed components are permanently removed and the relative ordering of the remaining feasible paths is preserved. The cascading failure process follows well-established probabilistic disaster models from the literature, which provide failure probabilities for each event and serve as inputs to the Time-Varying Graph (TVG), enabling the anticipated evolution of network states to be incorporated into restoration decisions. By integrating probabilistic failure awareness with event-driven TVG modeling and path stability, the proposed approach enables proactive and resilient restoration strategies that go beyond classical protection mechanisms designed for static or simultaneous failure scenarios.

In [8] we modeled the network under cascading failures using the TVG model, and proposed a strategy called *TVG-DR*. We have shown that TVGs under monotonic degradation

978-3-903176-78-2

preserve the distance-based path stability, allowing the use of a relative order of precomputed k-shortest paths during monotonic degradation.

In this paper, we propose the *TVG-DRP* (Time-Varying Graph Disaster Recovery with Probabilities). Its novelty compared to our earlier *TVG-DR* [8] lies in the integration of probabilistic risk metrics into the TVG framework. While TVG-DR relies exclusively on distance metrics, TVG-DRP extends this model by incorporating three risk-aware cost functions: failure probability, distance to disaster zones, and hop-based exposure. This probabilistic layer enables TVG-DRP to explicitly incorporate the link failure probabilities into the model and anticipate the likelihood of re-disruption, and to avoid paths with higher risk, rather than only relying on topological stability. Moreover, it operates by assigning dynamic risk-based weights to each edge in the network over time, reflecting both the likelihood of failure and the proximity to affected areas. Based on this enriched TVG model, we compute candidate paths per demand and select the one that minimizes a composite cost function defined by failure probability, distance to failure, and hop count. This strategy enables more informed and resilient restoration decisions in scenarios involving cascading failures.

To the best of our knowledge, TVG-DRP is the only approach in the literature that integrates cascading awareness, temporal modeling, probabilistic reasoning, and precomputed path pools, leading to measurable improvements in restorability, availability, and resilience. Simulation results show that TVG-DRP reduces non-restored requests by up to 40%, improves restorability by up to 9%, and increases availability by around 0.2%, when compared to existing approaches such as PRWR, PRPA, and TVG-DR.

The rest of this paper is organized as follows. Section II introduces the failure propagation and risk model, while Section III presents the network model and restoration constraints. Section IV describes the TVG-based modeling. Section V describes the probabilistic path stability concept. Section VI details the proposed TVG-DRP heuristics, and Section VII reports the evaluation setup and results. Finally, Section VIII concludes the paper and outlines future research directions.

II. Failure Propagation and Risk Model

Failures due to disasters in optical networks rarely occur as isolated events. Instead, the initial disruption often leads to a sequence of dependent failures over time, known as cascading failures [9]. Such disasters cause correlated, spatially localized failures that propagate over time, degrading service availability through direct damage and secondary effects such as overloads and cascading disruptions. Unlike static models, cascading failures demand adaptive, forward-looking restoration, selecting paths not only for current feasibility but also for their likelihood of remaining operational, to avoid re-disruptions and degraded service availability. Thus, robust restoration strategies must account for both the spatial and temporal evolution of these failures, requiring models that support continual re-evaluation and re-optimization.

To ensure comparability with prior studies, we adopt a probabilistic failure model based on [9] and extended in [2], [8], partitioning the network into three concentric zones around the disaster epicenter, as illustrated in Fig. 1: **Zone Z1** – Core disaster area, with high failure probability (p_1); **Zone Z2** – Surrounding area, with moderate failure probability (p_2); and **Zone Z3** – Peripheral area, with low failure probability (p_3), where $p_1 > p_2 > p3$.

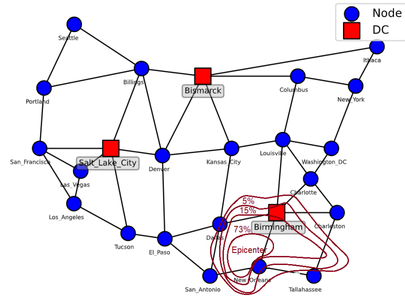

Fig. 1: Example the probabilistic failure model illustrating concentric zones Z1–Z3 around the epicenter with failure probabilities of 73%, 15%, and 5%, respectively [2], [9].

This zonal model reflects how proximity to the epicenter influences the likelihood of failure, guiding restoration weights and treating components located within or near Z_1 as high-risk regions to be avoided during path selection. These failures evolve as surviving components become overloaded or lose critical connectivity. At each time step t, a subset of network elements $F_t \subseteq V \cup E$, where V and E denote the node and link sets, fails based on its associated probability and zone.

To ensure reproducibility, the disaster scenario was modeled following the framework described in [9], where failures propagate spatially and temporally from a single epicenter. The reference topology corresponds to the 24-node, 43-link U.S. backbone used in [2], [8]. A disaster event starts at time t_0 with the network fully operational (G_0). Each newly triggered cascading stage produces successive snapshots G_t, which collectively form the TVG. At each stage t, the set of failed components $F_t \subset V \cup E$ is determined according to their spatial distance from the epicenter and the zone-specific failure probability p_i. Elements within Z_1 have the highest chance of failure and are removed first; failures then spread concentrically to Z_2 and Z_3, modeling a progressive degradation process. Ten independent scenarios were generated by varying the epicenter location to capture different spatial exposure conditions. This model emulates a large-scale disaster where failures occur gradually rather than simultaneously, producing monotonic degradation $\{G_0, G_1, ..., G_T\}$, as discussed later, that reflects the temporal evolution of the network.

III. Network Model and Restoration Objective

The EON is modeled as a graph $G = (V, E)$, where V is the set of optical nodes and E is the set of bidirectional fiber links. Each link $e \in E$ is associated with a set of available frequency slots S_e, capacity C_e, and slot bandwidth B_{slot}. A connection request is defined as $r_i = (s_i, d_i, bw_i)$, where s_i and d_i are the source and destination nodes, respectively, and bw_i is the requested bit rate. To accommodate bw_i, a contiguous and continuous set of slots $s \subseteq S_e$ must be reserved on each link e of the selected path, satisfying $|s| \cdot B_{\text{slot}} \geq bw_i$.

For each established request r_i, we precompute k candidate paths over the initial network, storing them in the pool $P_i^t = P_{i,1}, ..., P_{i,k}$ for later evaluation based on availability and risk.

As a disaster unfolds, established connections may be disrupted. At each time step t, the goal is to restore affected connection using the current operational topology. A valid restoration path P_i for connection r_i must satisfy the following constraints: $i)$ **Connectivity**: all nodes and links in P_i^t must be operational at time t; $ii)$ **Spectrum Feasibility**: there must exist a block of contiguous and continuous slots $s \subseteq S_e$ on each link $e \in P_i^t$ such that $|s| \cdot B_{\text{slot}} \geq bw_i$; and $iii)$ **Spectrum Continuity**: the selected slots must be the same on all links in the path.

The restoration algorithm selects the candidate path, reducing accumulated cost, with cost functions incorporating failure risk, proximity to failed components, and hop sensitivity, unlike traditional shortest-path strategies.

The overall objective is to minimize re-affectation risk, enhance restorability and availability, and ensure a feasible and sustainable service reestablishment under cascading failure conditions.

IV. Time-Varying Graph-Based Failure Modeling

To capture connectivity evolution under cascading failures, we model the network as a TVG, which represents progressive degradation and supports state-aware restoration. A TVG is defined as an ordered sequence of snapshots $\mathcal{G} = \{G_0, G_1, \ldots, G_T\}$, where each $G_t = (V_t, E_t)$ represents the state of the network at time t. The vertex set V_t and edge set E_t denote the active nodes and operational links at that instant, as affected by cascading failures.

Each edge $e \in E_t$ is characterized by its available slots S_e^t, failure probability $p_t(e)$, distance to the nearest failure $d_t(e)$, and hop indicator $ht(e)$, set to 1 for each traversed edge so that $\sum_{e\in P} ht(e)$ equals the path hop count.

This temporal modeling allows us to evaluate not only the feasibility of restoration paths at each time t, but also their stability and risk over time. In particular, it supports risk-aware restoration decisions that evolve with the disaster.

Monotonic Degradation Structure

In disaster-induced cascading scenarios, network connectivity tends to degrade progressively. Failures accumulate over time through the removal of nodes or links, but new elements are not added. This results in a *monotonic degradation* pattern:

Definition 1 (Monotonic Degradation TVG)**.** *A TVG $\mathcal{G} = \{G_0, G_1, ..., G_T\}$ exhibits monotonic degradation if for all $t > 0$, $G_t \subseteq G_{t-1}$ and therefore $G_t \subseteq G_0$.*

This property ensures that the structure of the network becomes increasingly constrained, but never augmented. As a consequence: $i)$. Paths that are feasible at $t = 0$ may become invalid in later G_t; $ii)$. The order of precomputed paths (based on cost or risk) in G_0 can remain valid for all G_t as long as they remain feasible; $iii)$. It enables the reuse of candidate path pools, reducing the computational overhead during restoration. This insight was formalized in our earlier work as the Path Stability (PS) property [8], which proves that under monotonic degradation, the relative cost order of the original k-shortest paths remains preserved among the surviving ones. This allows efficient restoration without recomputing paths from scratch under fixed deterministic hop weights.

By leveraging the monotonicity of TVGs, the restoration algorithm can make decisions based on precomputed candidates while continuously updating risk metrics and feasibility in response to failures.

V. Probabilistic Path Stability

In this paper, we discuss the Probabilistic Path Stability (PPS), where the restoration strategy explicitly accounts for the probability of failure of each link.

Let $\mathcal{G} = \{G_0, \ldots, G_T\}$ be a TVG with *monotonic* degradation, i.e., edges or nodes may fail over time and remain failed thereafter. For each edge e and stage t, let $p_t(e) \in [0, 1)$ denote the failure probability at stage t, and define the per-edge survival $s_t(e) = 1 - p_t(e)$.

In the deterministic setting, the PS of P through the TVG measures the persistence of feasibility (e.g., connectivity/resources) of P across stages; the induced ranking prefers paths that remain valid for longer horizons under the evolving failures. In the probabilistic setting, we define the survival of P over a horizon $H \geq 0$ starting at t as

$$\begin{aligned} S_{t,H}(P) &= \prod_{\tau=t}^{t+H} \prod_{e\in P} s_\tau(e) \\ &= \exp\big(-W_{t,H}(P)\big). \end{aligned} \tag{1}$$

where $W_{t,H}(P)$ is defined as the hazard of P over the horizon $H \geq 0$ and is given by

$$W_{t,H}(P) \triangleq \sum_{\tau=t}^{t+H} \sum_{e\in P} w_\tau(e),$$

$$w_\tau(e) \triangleq -\ln\big(s_\tau(e)\big) = -\ln\big(1 - p_\tau(e)\big).$$

Hence, for any fixed (t, H) and any two feasible paths P_a, P_b,

$$S_{t,H}(P_a) \geq S_{t,H}(P_b) \quad \Longleftrightarrow \quad W_{t,H}(P_a) \leq W_{t,H}(P_b).$$

Maximizing $S_{t,H}(P)$ is equivalent to minimizing $W_{t,H}(P)$.

Hence, PPS induces a ranking aligned with the PS objective when stability is defined as the probability that a path remains operational over the considered horizon. Suppose the PS ranks paths by the likelihood of remaining valid across stages (failures only remove feasibility). Then, under conditional failure independence across edges within each stage and the monotone-failure assumption across stages, the PS ranking coincides with the PPS ranking induced by $S_{t,H}(P)$; this is an approximation, since geographically correlated failures are partly captured by the zonal disaster model rather than by explicit inter-edge dependence.

We argue that if $\widehat{W}(P) = \sum_{\tau=0}^{H_0} \sum_{e\in E(P)} w_\tau(e)$ is a precomputed cumulative hazard on G_0 for horizon $H_0 \geq 0$ and $\preceq_{\widehat{W}}$ is the induced order. Then if $w_\tau(e)$ is non-decreasing in τ for all e (expanding disaster), then, across stages, the order $\preceq_{\widehat{W}}$ tends to be preserved among surviving candidates when re-evaluated with the instantaneous $W_{t,0}(P)$ (ties broken by $\widehat{W}$).

This analysis shows that the PPS model extends the deterministic PS concept by interpreting stability as the probability of a path remaining feasible over time. Within PPS, the survival $S_{t,H}(P)$ and the hazard $W_{t,H}(P)$ are dual representations of the same process, and under monotone failures, their induced rankings remain consistent across stages. This formulation provides the theoretical basis for the hazard-based cost functions analyzed in Section VII.

VI. Restoration Heuristics with Probabilistic Awareness

This section presents three restoration strategies that incorporate distinct risk metrics into TVG-based path selection. Each strategy defines a cost model that accounts for different risk dimensions, such as failure probability, distance to disaster zones, or hop count.

We consider different techniques to identify the best strategy in the scenario. Let $c(P_i^t)$ be the cost of a path P_i^t defined by one of the following models: $\boldsymbol{c_1}(P_i^t) = \sum_{e \in P_i^t} p_t(e)$ (risk only);$\boldsymbol{c_2}(P_i^t) = \sum_{e \in P_i^t} \left(d_t^{n(e)} + p_t(e)\right)$, and $\boldsymbol{c_3}(P_i^t) = \sum_{e \in P_i^t} \left(h_n^{n(e)} + p_t(e)\right)$, where $p_t(e)$ is the failure probability of edge e, $d_t(e)$ is the distance of edge e to failure regions, $h_t(e)$ equals 1 for each edge (so that $\sum_{e \in P_i} h_t(e)$ gives the hop count of path P_i), and $n(e)$ defines cost amplification sensitivity.

As the failure propagates from a known epicenter, the network degrades progressively across discrete time steps. The failure spreads through concentric zones Z_1, Z_2, Z_3 with decreasing disruption probabilities $p_1 > p_2 > p_3$ (Fig. 1) [9].

A. The TVG-DRP Algorithm

The proposed TVG-DRP (**T**ime-**V**arying **G**raph-**D**isaster **R**ecovery with **P**robabilities) algorithm (Alg. 1) is instantiated using one of the three risk models $\boldsymbol{c_1}$, $\boldsymbol{c_2}$, or $\boldsymbol{c_3}$.

Algorithm 1: TVG-DRP Restoration Algorithm.

Input: TVG $\mathcal{G}_\mathcal{T}$, set of disrupted requests D, path pools P_i

Output: Restored paths P_i^t for each disrupted connection d_i

1 Update G_t with failed components F_t
2 **for** *each disrupted connection* $d_i \in D$ **do**
3   Identify valid paths $P_i^{\text{valid}} \subseteq P_i$ in G_t
4   Evaluate $c_x(P)$ for each $P \in P_i^{\text{valid}}$
5   Select $P_i^t = \arg\min c_x(P)$
6   Allocate spectrum if possible; otherwise mark as blocked

In the proposed restoration mechanism (Alg. 1), for each disrupted connection $d_i = (s_i, d_i, bw_i)$ (Line 2), a feasible path P_i^t in G_t is sought (Line 3) that: i) connects the source s_i to the destination d_i via available nodes and links, ii) guarantees contiguous and continuous frequency slots on each link $e \in P_i^t$, and iii) provides sufficient bandwidth, i.e., $|s| \cdot B_{\text{slot}} \geq bw_i$, where $s \subseteq S_e^t$ represents the spectrum allocated on link e at time t. Subsequently (Line 4), the feasible precomputed candidate paths in the set $\mathcal{P}_i = \{P_{i,1}, \ldots, P_{i,K}\}$ are evaluated using distinct cost functions ($\boldsymbol{c_1}$, $\boldsymbol{c_2}$ and $\boldsymbol{c_3}$). Each strategy is tested separately by selecting the best valid path according to its respective cost function.

Leveraging the monotonic degradation property (Sec. IV) and PS (Sec. V), the relative cost ranking remains valid across time-varying snapshots. The algorithm then selects the lowest-cost feasible path according to the chosen risk model (Line 5), dynamically adapting the restoration to avoid high-risk zones and minimize re-disruption as the network deteriorates.

The cost function $\boldsymbol{c_1(P)}$ directly aggregates the instantaneous failure probabilities along the path, effectively approximating the first-order expansion of $W_{t,H}(P)$ for small $p(e)$. This model preserves the proportionality between cumulative risk and expected path hazard, but it ignores spatial correlations, potentially favoring paths that traverse high-density risk regions if their aggregate probability remains small.

In turn, $\boldsymbol{c_2(P)}$ introduces a geometric correction that penalizes edges closer to the epicenter of the disaster. This model captures the spatial gradient of risk while retaining the probabilistic weighting from Section V. In practice, it behaves as a spatially regularized hazard, discouraging route selections that may soon intersect the expanding failure region.

The expression $\boldsymbol{c_3(P)}$ adds a hop-count sensitivity term that penalizes longer paths, controlling for compounded risk due to the multiplicative nature of PPS (since each additional edge introduces an extra survival term $s_\tau(e)$). Hence, c_3 explicitly reflects the exponential growth of hazard with path length observed in $W_{t,H}(P)$.

Each strategy ($\boldsymbol{c_1}$,$\boldsymbol{c_2}$, $\boldsymbol{c_3}$) results in a distinct version of the algorithm, enabling comparative evaluation of risk-based heuristics.

Using a precomputed pool of k candidate paths entails a trade-off between restoration speed and solution space coverage. While larger k improves feasibility under extreme degradation, it also increases computational cost and latency. Since this work targets fast online restoration under cascading failures, it prioritizes timely recovery with a bounded candidate set over global recomputation. The parameter n controls the cost sensitivity to risk. Lower n favors shorter, moderately risky paths, whereas higher n penalizes longer or failure-prone routes. We evaluate c_1, c_2, and c_3 separately to isolate each component's effect; weighted combinations are possible, but left for future work to preserve interpretability and fit the paper scope.

Complexity Analysis

Let N be the number of disrupted connections, k the number of candidate paths per connection, and L the average path length. The computational complexity at each time step t is $O(N \cdot k \cdot L)$, as only precomputed paths are re-evaluated. This design enables efficient adaptation to dynamic network conditions with low computational overhead, though it constrains decisions to the quality of the initial candidate pool.

VII. Numerical Examples

We evaluated the TVG-DRP algorithm through simulations of cascading failures on the 24-node, 43-link EUA EON topology (Fig. 1) [2], [8], [9], where each link provides 400 frequency slots. Each simulation generates 100,000 connection requests following a Poisson arrival process. The requested bit rates are 10 Gbps (40%), 20 Gbps (30%), 40 Gbps (10%), 80 Gbps (10%), 100 Gbps (5%), and 200 Gbps (5%), thus reflecting heterogeneous traffic demands. Connection holding times follow an exponential distribution, and the offered network load varies between 550 and 850 Erlangs, measured as network-wide offered traffic (arrival rate × mean holding time), rather than per-slot-per-link occupancy.

For each source–destination pair, $k = 10$ candidate paths are precomputed using Yen's algorithm on the initial graph G_0. These paths are stored in a pool and reused throughout the simulation to reduce computational overhead and enable rapid recovery decisions. Ten failure scenarios are simulated based on the cascading failure model of Colman-Meixner et al. [9], which defines three concentric zones (Z1, Z2, Z3) around an epicenter with failure probabilities of 73%, 15%, and 5%, respectively. The failure model used was presented by [2], containing 10 failure scenarios[1].

Cascading events occur incrementally every 3,600 time units, progressively degrading the network and producing successive snapshots that form the TVG. This interval is chosen to be significantly larger than the average connection setup time and comparable to the mean holding time, allowing failures to propagate gradually rather than simultaneously. Upon disruption, the proposed TVG-DRP restoration algorithm is immediately triggered to reestablish the connection using the precomputed path pool and available spectrum resources.

During simulation, every connection affected by a failure at any stage of the cascading process triggered the execution of the proposed TVG-DRP restoration algorithm. The three variants of the proposed algorithm (TVG-DR**P** based on failure **probability** ($\boldsymbol{c_1(P)}$), TVG-DRP**D** based on **distance** ($\boldsymbol{c_2(P)}$), and TVG-DR**H** based on **hop** count ($\boldsymbol{c_3(P)}$)) are evaluated independently and compared against established baselines: PRPA [2] (α=0), which combines failure probability and hop count in its risk metric; PRWR [10], a restoration strategy that performs shortest-path rerouting without accounting for cascading failure risk; and TVG-DR [8], our previous model based solely on distance-based path stability. Each simulation was executed ten times, and the results are presented with 95% confidence intervals.

Fig. 2 shows the number of connections not restored after a disaster. The TVG-DRP variants consistently achieve the lowest values across all load levels, resulting from the precomputed pool of candidate paths and dynamic, risk-aware evaluation at each TVG snapshot. The distance-based TVG-DRPD (green) minimizes unrecovered connections by avoiding links near high-risk zones, while the hop-based TVG-DRPH (cyan) performs comparably under high loads by favoring shorter, less congested paths. In contrast, PRPA and PRWR, lacking temporal adaptability, struggle to maintain feasible alternatives as cascading failures progress, leading to more non-restored requests with increasing load.

Restorability, computed as the ratio between restored and disrupted requests, highlights the benefits of integrating probabilistic awareness into the TVG framework (Fig. 3). The TVG-

[1] https://github.com/GSRamalho/python-simple-anycast-wdm-simulator/tree/master/zones

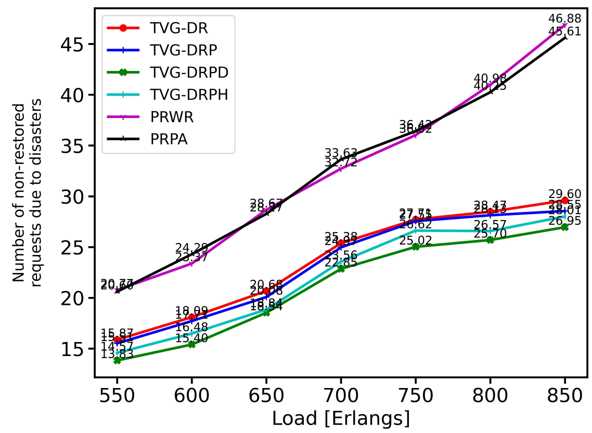

Fig. 2: Average number of non-restored requests (out of 100,000 generated requests) under cascading failure.

DRP variants consistently achieve higher restoration ratios across all load levels by dynamically re-evaluating the feasibility and risk of precomputed paths as the network degrades. The distance-based variant TVG-DRPD (green) delivers the best performance, reaching up to 9.89% higher restorability than PRWR at 850 Erlangs by penalizing links near the failure epicenter and favoring paths that remain operational throughout the cascading process. The hop-based TVG-DRPH (cyan) performs comparably by limiting exposure through shorter paths. In contrast, TVG-DR and the baseline strategies PRPA and PRWR lack adaptive risk awareness, preventing them from anticipating the evolving failure regions and resulting in earlier disconnections and lower restorability.

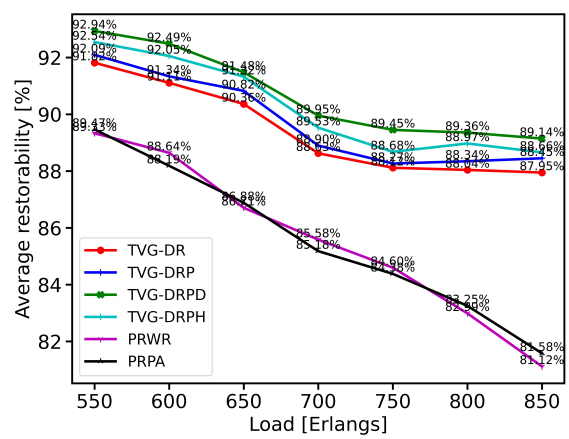

Fig. 3: Average network restorability.

Regarding Re-affected Requests (Fig. 4), the TVG-DRP variants achieve a clear reduction in secondary disconnections compared to baseline strategies. Their continual reassessment of probabilistic risk in each TVG snapshot penalizes links near the evolving disaster region, preventing restored paths from intersecting failure zones. The results achieved by the distance-based version TVG-DRPD (green) stem from its explicit incorporation of the spatial gradient of failure probability, which favors restoration routes that remain outside the dynamically expanding disaster region. The hop-based model TVG-DRPH (cyan) exhibits comparable performance, since shorter paths accumulate less compounded risk and traverse fewer potentially failing components. Both mechanisms effectively maintain restored flows in stable network areas, while PRWR, which relies solely on topological shortest paths, and PRPA, whose probabilistic weighting is static and applied only at restoration time, cannot anticipate the evolving nature of cascading failures. As a result, they often select routes that soon intersect expanding failure regions, leading to higher re-affection rates as load and network degradation intensify.

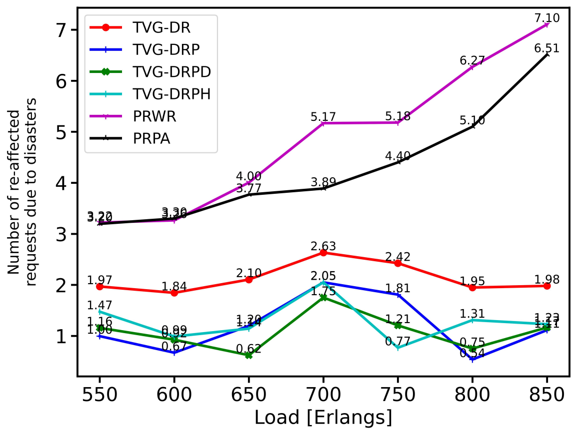

Fig. 4: Connections disrupted again (%) after restoration.

Due to higher restorability (Fig. 3) and fewer re-disconnections (Fig. 4), the TVG-DRP variants maintain higher service availability during cascading failures (Fig. 5). Risk-aware path selection avoids links with elevated failure probability or proximity to the epicenter, stabilizing active connections as the network degrades. The hop-based variant TVG-DRPH (cyan) achieves slightly higher availability under heavy load by favoring shorter paths with lower compounded risk, whereas TVG-DR and baseline strategies fail to anticipate progressive degradation, leading to earlier disconnections. Although the average availability gain is modest (≈0.2%), it remains significant in cascading scenarios, indicating that TVG-DRP primarily improves service stability rather than absolute uptime.

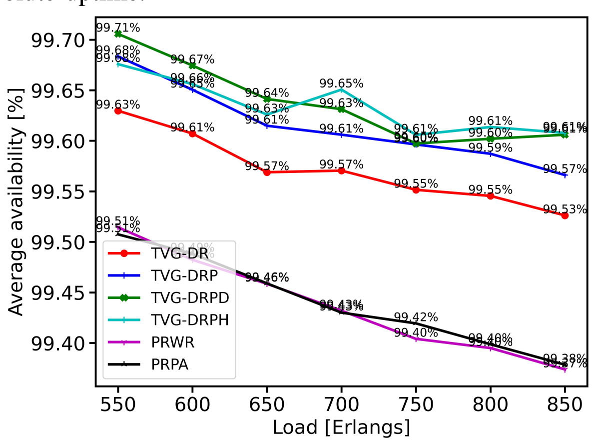

Fig. 5: Average network availability.

Similar blocking ratios across all algorithms and load levels indicate that TVG-DRP introduces no spectrum overhead; in all tested loads, the differences remained within the confidence intervals, so the detailed curves were omitted for space.

Overall, incorporating probabilistic reasoning into restoration heuristics significantly improves EON resilience under cascading failures. Each TVG-DRP variant offers complementary benefits and consistently outperforms static and risk-weighted baselines.

VIII. Conclusion

This paper introduced TVG-DRP, a family of restoration heuristics for EONs under cascading failures. By enriching a Time-Varying Graph (TVG) model with probabilistic metrics—failure probability, distance to disaster zones, and hop count—TVG-DRP enables adaptive, risk-aware restoration that evolves with network degradation. Simulation results across ten scenarios show that its three variants consistently outperform PRPA, PRWR, and TVG-DR, achieving up to 9% higher restorability, around 0.2% higher availability, and up to 30% fewer non-restored requests, while preserving computational efficiency. The key contribution lies in integrating probabilistic risk modeling with TVGs under monotonic degradation, leveraging path stability to enable fast, consistent online restoration across cascading stages without global recomputation. Future work will extend this framework with real-time failure prediction and SDN integration, as well as additional topologies, weighted combinations of c1–c3, physical-layer feasibility constraints, and traffic models with grooming or longer-lived circuit demands.

References

[1] R. Zou, H. Hasegawa, M. Jinno, and S. Subramaniam, "Drama+: Disaster management with mitigation awareness for translucent elastic optical networks," *IEEE Trans. Netw. Serv. Manag*, vol. 19, no. 3, pp. 2587–2599, 2022, http://doi.org/10.1109/TNSM.2022.3189362.

[2] G. S. Ramalho, K. O. Fonseca, D. F. Pigatto, C. Natalino, P. Monti, G. B. Figueiredo, and J. de Santi, "Cascading-failure-aware disaster recovery in optical cloud networks," in *Int. Conf. Opt. Netw. Des. Model. (ONDM)*, 2023, pp. 1–3, http://doi.org/10.23919/ONDM57372.2023.10144843.

[3] S. Li, R. Gu, G. Zhang, Y. Wang, Y. Wang, and Y. Ji, "Order aware service recovery algorithm in elastic optical network with multiple failures," in *Int. Conf. Netw. Netw. Appl. (NaNA)*, 2019, pp. 135–141, http://doi.org/10.1109/NaNA.2019.0003.

[4] J. P. Sterbenz *et al.*, "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Comput. Netw.*, vol. 54, no. 8, pp. 1245–1265, Jun. 2010, http://doi.org/10.1016/j.comnet.2010.03.005.

[5] K. D. R. Assis, R. C. Almeida, H. Waldman, A. F. Santos, M. S. Alencar, M. J. Reed, A. Hammad, and D. Simeonidou, "Sla formulation for squeezed protection in elastic optical networks considering the modulation format," *Journal of Optical Communications and Networking*, vol. 11, no. 5, pp. 202–212, 2019, http://dx.doi.org/10.1364/JOCN.11.000202.

[6] A. K. Sarma, S. K. Deka, and N. Sarma, "Survivable elastic optical network: A survey of failure scenarios and solutions," *Opt. Switch. Netw.*, vol. 58, p. 100823, 2025, https://doi.org/10.1016/j.osn.2025.100823.

[7] R. Zou, S.-C. Lin, M. Matsuura, H. Hasegawa, G. Figueiredo, and S. Subramaniam, "Drama-pro+: disaster recovery algorithm with mitigation awareness for protected translucent elastic optical networks," *J. Opt. Commun. Netw.*, vol. 17, no. 9, pp. D18–D29, 2025, http://doi.org/10.1364/JOCN.561155.

[8] E. O. Santos, L. F. Ferreira, J. de Santi, R. Zou, S. Subramaniam, and G. B. Figueiredo, "Post-disaster restoration based on path stability in elastic optical networks under cascading failures," in *IEEE Future Networks World Forum (FNWF)*. IEEE, 2025, to appear.

[9] C. Colman-Meixner, M. Tornatore, and B. Mukherjee, "Cloud-network disaster recovery against cascading failures," in *Proc. IEEE GLOBECOM*, Dec. 2015, http://doi.org/10.1109/glocom.2015.7417558.

[10] C. N. da Silva, L. Wosinska, S. Spadaro, J. C. W. A. Costa, C. R. L. Frances, and P. Monti, "Restoration in optical cloud networks with relocation and services differentiation," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 8, no. 2, pp. 100–111, Feb 2016, https://doi.org/10.1364/JOCN.8.000100.