Budgeted Link Shielding for Max-Min Path Availability under Correlated Failures: A Genetic Algorithm Approach

Nándor Bándi, Balázs Vass, Noémi Gaskó

***Abstract*—Regional disasters and other correlated hazards can simultaneously disable multiple links of a communication network, while reinforcing every vulnerable component is generally economically infeasible. We study the budgeted link-shielding problem, in which a set of links is selected for reinforcement subject to a cost constraint. Shielded links are assumed to remain operational under the modeled failures, and the objective is to maximize the minimum path availability over a prescribed set of protected endpoint pairs. The study goes beyond assuming network failures being independent, and supposes there are correlated failures of component sets. Since the number of possible shielding configurations is exponential and exact safest-path computation under arbitrary correlations is NP-hard, we evaluate each configuration by first constructing candidate paths from marginal link-failure probabilities and then computing the exact correlated-failure availability of the selected paths using inclusion–exclusion. We embed this evaluator into a binary genetic algorithm and investigate standard tournament selection and a connected-component-aware selection rule. A greedy risk-to-cost heuristic is used as a baseline. Experiments on real backbone-network topologies indicate that standard tournament selection provides the strongest overall performance, while the component-based preference offers no consistent advantage and the greedy method remains competitive in tightly constrained settings.**

***Index Terms*—correlated failures, genetic algorithm, link shielding, network availability, network resilience**

I. Introduction

Communication backbone networks are expected to provide high availability even when parts of the underlying infrastructure are damaged. Major disruptions, however, rarely affect network components independently: earthquakes, floods, construction accidents, and deliberate attacks may simultaneously disable several geographically co-located links. Such correlations can substantially change the availability of end-to-end paths and can make estimates based only on marginal link-failure probabilities overly optimistic [1]–[3].

Network operators can address failures through several complementary mechanisms [4]. Protection and restoration schemes provision alternative routes that carry traffic after a primary path becomes unavailable [5], while network-augmentation methods add links or capacity to improve connectivity [6]. A further option is *link shielding*, whereby selected physical links are reinforced before failures occur. Depending on the infrastructure, shielding may represent stronger cable protection, deeper burial, improved conduits, relocation, or another intervention that makes the protected link immune to the considered hazards. Previous studies have examined shielding vulnerable cuts, seismic-resistant cable laying and shielding, and minimum-cost shielding that guarantees connectivity under geographical failures [7]–[9]. These works establish shielding as a viable network-planning mechanism, but focus mainly on robustness or deterministic connectivity guarantees.

In this paper, we study an availability-oriented shielding problem. We are given an undirected network, a set of protected endpoint pairs, a shielding cost for every link, and a total shielding budget. The goal is to select a budget-feasible set of links to shield so as to maximize the minimum path availability among the protected pairs. The max-min objective prevents the optimization from improving already reliable connections while leaving one pair as a pronounced availability bottleneck. We adopt an idealized binary shielding model in which a shielded link cannot fail under the modeled hazards, and paths may be recomputed after selecting the shielding configuration.

To explore the exponentially large shielding search space, we develop a binary genetic-algorithm framework. An individual encodes the selected links as a bit vector, and its fitness is the minimum log-availability among the protected pairs. We examine standard tournament selection and a second selection rule that also favors shielding sets whose selected-edge subgraph has fewer connected components. The latter is motivated by the observation that a connected shielded subgraph spanning the protected terminals can provide fully shielded paths among them when the budget is sufficiently large. We also introduce a greedy baseline that repeatedly considers the currently least-available protected pair and selects an affordable link on its candidate path according to a risk-to-cost criterion.

Genetic algorithms have previously been applied to budgeted network-immunization problems [10], [11] and optimal network design [12]. Those studies, however, optimize contagion dynamics over protected nodes and network structure rather than end-to-end path availability under correlated physical link failures. In contrast, the present work applies

Nándor Bándi, Balázs Vass, and Noémi Gaskó are with Faculty of Mathematics and Computer Science, Babeș-Bolyai University of Cluj Napoca, Romania. Balázs Vass is also with Department of Telecommunications and Artificial Intelligence, Faculty of Electrical Engineering and Informatics, Budapest University of Technology and Economics, Műegyetem rkp. 3., H-1111 Budapest, Hungary. E-mails: {firstname.surname}@ubbcluj.ro.

This work was supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS-UEFISCDI, project no. PN-IV-P2-2.1-TE-2023-1977.

evolutionary search to the selection of physically shielded links under a probabilistic correlated-failure model.

The contributions of this paper are twofold. First, we formulate a budgeted link-shielding problem with a max-min path-availability objective under correlated link failures. Second, we develop two genetic-algorithm variants and a problem-specific greedy baseline.

The current experiments indicate that the genetic algorithm using standard tournament selection performs best overall. The connected-component-aware selection rule provides no consistent improvement, whereas the greedy method remains competitive under tight shielding budgets. These results suggest that directly exploiting the availability objective is more effective than imposing a general preference for topologically concentrated shielding sets.

The remainder of the paper is organized as follows. § II formulates the budgeted link-shielding problem and the correlated-failure availability model. § III presents the configuration evaluator, the two genetic-algorithm variants, and the greedy baseline. § IV describes the experimental setup and reports the numerical results. § V concludes the paper.

II. Problem Formulation

Let $G = (V, E)$ be a connected, undirected graph and $C \subseteq \{\{s,t\} \mid s \neq t,\ (s,t) \in V \times V\}$ the set of securely connected vertex pairs. Let $\texttt{cost} : E \to \mathbf{R}^+$ be the shielding (reinforcement) cost function defined on the edges, and let B be the total shielding budget. Let $E_B = \{E^* \subseteq E \mid \sum_{e \in E^*} \texttt{cost}(e) \leq B\}$ be the set of edge subsets whose cost does not exceed the budget.

Cumulative Failure Probability The cumulative failure probability of a given edge set S is the probability that every edge in S fails (and possibly other edges as well), and is denoted by CFP(S). Let CFP[G] be the set of all edge sets and cumulative failure probabilities whose CFP is strictly positive, in graph G.

$$\mathrm{CFP}[G] = \{(S, \mathrm{CFP}(S)) \mid \mathrm{CFP}(S) > 0\}$$

Safest paths Safely connected vertices communicate through paths that can have different rates of failure. Let $d_{s,t} \in G_{s,t}$ be a path in graph G connecting vertices s and t, and its associated availability probability (probability that no edge fails):

$$P(d_{s,t}) = P\Big(\bigcap\nolimits_{e \in d_{s,t}} \neg e\Big) = 1 - P\Big(\bigcup\nolimits_{e \in d_{s,t}} e\Big)$$
$$= 1 - \sum\nolimits_{S \neq \emptyset,\, S \subseteq d_{s,t}} (-1)^{|S|+1} \mathrm{CFP}(S).$$

Notion of availability The availability of a safely connected vertex pair can be defined as the probability of successful communication through the safest path connecting them. Let $A((s,t), \mathrm{CFP}[G])$ be the probability of availability between vertices s and t across the safest path having the failure probabilities CFP[G]. The availability between vertices s and t is defined as the availability of the safest path connecting them, formally this is defined as:

$$A((s,t), \mathrm{CFP}[G, E^*]) = \max(P(d_{s,t}) \mid d_{s,t} \in [G, E^*]_{s,t})$$

The shielded graph is denoted by $[G, E^*]$. Note that computing a safest path (and thus, computing the availability) between nodes s and t is NP-hard [13].

Problem Given an initial CFP[G], find the reinforced edge set $E^* \in E_B$ that stays within budget and maximizes the minimum availability probability of the safest paths between all securely connected vertex pairs. Formally, find:

$$\arg \max_{E^* \in E_B} \{\min \{A((s,t), \mathrm{CFP}[G, E^*]) \mid (s,t) \in C\}\}$$

The modelling of edge failures is expressive enough to model node failures, which can be done by considering all the adjacent edges as failed. This approach results in a more generic problem formulation.

Search space and cost There are $2^{|E|}$ possible shielding configurations. For each such configuration, there are $\sum_{s,t \in C} |G_{s,t}|$ possible paths. In order to explore the search space fully, $2^{|E|} \cdot \sum_{s,t \in C} |G_{s,t}|$ number of inclusion-exclusion formula evaluations are necessary.

Independent edge failures If edge failures are treated as independent, the inclusion-exclusion evaluations can be replaced by shortest path searches. Let $1 - \mathrm{CFP}(\{e\})$ be the probability that edge e does not fail (its availability), and let $d^*_{s,t}$ be the safest path. In the independent case it can be defined as:

$$d^*_{s,t} = \arg \max_{d_{s,t} \in [G, E^*]_{s,t}} P\Big(\bigcap\nolimits_{e \in d_{s,t}} \neg e\Big)$$
$$= \arg \max_{d_{s,t} \in [G, E^*]_{s,t}} \Big(\prod\nolimits_{e \in d_{s,t}} (1 - \mathrm{CFP}(e))\Big)$$

which is equivalent to the shortest path in the weighted graph:

$$\arg \min_{d_{s,t} \in [G, E^*]_{s,t}} \Big(\sum\nolimits_{e \in d_{s,t}} - \log(1 - \mathrm{CFP}(e))\Big)$$

Under independent edge failures, the marginal-weight candidate is an exact safest path. Under correlated failures, the candidate-path availability is a heuristic approximation of the exact safest-path availability: $A((s,t), [G, E^*]) \approx P(d^*_{s,t})$.

Reduced cost Assuming edge failures are not correlated, the number of inclusion-exclusion evaluations drops significantly. For a single shielding only $|C|$ number of shortest path calculations and inclusion-exclusion evaluations are necessary, $2^{|E|} \cdot |C|$ in total to cover every configuration.

Example 1: Figure 1 shows a small example graph. Having a shielding budget of 2 edges, the minimum availability was improved by 4 percentage points.

Used data structure The data structure used to represent a graph and its failure probabilities can be represented as the tuple $\mathrm{GCFP} = (E, \mathrm{CFP}[G])$. The *edges[GCFP]* function of such a structure represents the edge set E, and function *CFP[GCFP]* of such a structure represents the cumulative failure probability set $\mathrm{CFP}[G]$ associated with the structure. The *edges[cfp]* function represents the edge set S of an element *cfp* within the cumulative failure probability set $\mathrm{CFP}[G]$.

III. Proposed approach

The exponentially increasing size of the set of all possible shielding configurations prohibits the exhaustive search approach for large graphs. The optimization literature suggest a heuristics based approach to solving such problems. The combination of metaheuristics such as genetic algorithm and problem specific heuristics can yield acceptable results in a reasonable amount of time.

A. Evaluation of a Single Shielding Configuration

Algorithm 1 describes the evaluation of one shielding configuration. Assuming that the graph, the cumulative failure probabilities, the safely connected vertex pairs and the set of edges to shield E^* are given, the procedure defines the cumulative failure probabilities in light of the shielded edges and uses these to define the 'weights' of the graph. Next, it computes a marginal-probability shortest path for each protected endpoint pair and calculates the exact correlated-failure availability of each selected candidate path. As a last step the safest paths and availabilities (log) between vertex pairs are returned.

B. Genetic Algorithm

The genetic algorithm was the metaheuristic of choice due to its high effectiveness in exploring large search spaces. Each solution is encoded as a bitmask of the selected edges. A population is evolved until either the number of feasible solution evaluations limit, or the overall generations limit is reached. In realistic scenarios, the available budget enables the shielding of a small subset of edges.

Algorithm 2 describes the used genetic algorithm.

Selecting the right parents is important in order to be able to find solutions that are within the budget constraint. Parent selection variants employed include simple tournament selection (GA-1) and a tournament selection enhanced with a connected component based heuristic (GA-2). The simple tournament selection considers only the offspring with maximal fitness within a tournament, it does not enforce a preference regarding the connectivity of the shielded edges. Both genetic algorithm variants use the single point crossover strategy, bit-flip mutation

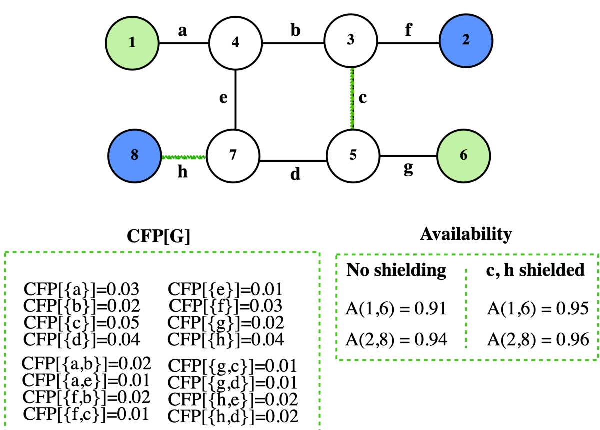

Fig. 1: Illustration of the effect of link shielding on availability. Before shielding the guaranteed availability was 0.91 and after, 0.95.

Algorithm 1 Evaluation of a shielding configuration

Require: **gcfp**: Cumulative failure probabilities and graph structure; **pairs**: List of (s,t) vertex pairs to protect; E^*: Set of edges to shield (reinforce);
Ensure: **paths**: Marginal-weight candidate path for each protected pair; **availabilities**: Log-availability of the selected candidate path for each protected pair

function EvaluateConfig(gcfp, pairs, E^*)
 ▷ Mask CFP values of shielded edges
 $\mathrm{CFP}[\mathrm{gcfp}'] \leftarrow \emptyset$
 for $\mathrm{cfp} \in \mathrm{CFP}[\mathrm{gcfp}]$ **do**
  if $E^* \cap \mathrm{edges}[\mathrm{cfp}] = \emptyset$ **then**
   $\mathrm{CFP}[\mathrm{gcfp}'] \leftarrow \mathrm{CFP}[\mathrm{gcfp}'] \cup \mathrm{cfp}$
  end if
 end for
 ▷ Build undirected weighted graph
 $G_w \leftarrow \emptyset$, $w \leftarrow \mathrm{Weights}(\mathrm{edges}[\mathrm{gcfp}], \mathrm{CFP}[\mathrm{gcfp}'])$
 for $e \in \mathrm{edges}[\mathrm{gcfp}]$ **do**
  $G_w \leftarrow G_w \cup (e, w_e)$
 end for
 ▷ Find shortest path for each connected pair
 paths $\leftarrow \{\}$, availabilities $\leftarrow \{\}$
 for $(s,t) \in$ pairs **do**
  $d^*_{s,t} \leftarrow \mathrm{Dijkstra}(G_w, s, t)$
  $\mathrm{paths}_{s,t} \leftarrow d^*_{s,t}$
  ▷ Calculate true availability via inclusion-exclusion formula
  $A_{s,t} \leftarrow \mathrm{Availability}(d^*_{s,t}, \mathrm{CFP}[\mathrm{gcfp}'])$
  $\mathrm{availabilities}_{s,t} \leftarrow log(A_{s,t})$
 end for
 return paths, availabilities
end function
function Weights(E, $\mathrm{CFP}[\cdot]$) ▷ Compute weight of each edge
 $w \leftarrow \emptyset$
 for $e \in E$ **do**
  $w_e \leftarrow -\log(1 - \mathrm{CFP}(\{e\}))$
 end for
 return w
end function
function Availability($d_{s,t}$, $\mathrm{CFP}[\cdot]$)
 $A \leftarrow 1$ ▷ Compute $P(d_{s,t}) = 1 - \sum_{S \neq \emptyset, S \subseteq d_{s,t}} (-1)^{|S|+1} \mathrm{CFP}(S)$
 for each non-empty edge subset $S \subseteq d_{s,t}$ **do**
  $A \leftarrow A - (-1)^{|S|+1} \cdot \mathrm{CFP}(S)$ ▷ Lookup in CFP table
 end for
 return A
end function

The heuristic based tournament selection considers the weighted sum of the relative fitness and number of connected components of the shielded edges. Algorithm 3 selects parents based on a composite score combining fitness (70%) and number of connected components in the shielded edge subgraph (30%). Given a large enough budget, a shielding configuration connecting all safely connected pairs (Steiner tree) is enough to ensure maximal availability. Therefore shielding solutions having less connected shielding components can be favorable to more disconnected solutions in this scenario.

Algorithm 2 Genetic Algorithm

Require: **gcfp**: Cumulative failure probabilities and graph structure; **pairs**: List of (s, e) vertex pairs to protect; $\boldsymbol{B}$: Shielding budget; $\boldsymbol{B_{evals}}$: Maximum number of feasible evaluations; $\texttt{cost}(\cdot)$: Shielding cost of each edge; $\boldsymbol{\mu}$: Population size; $\boldsymbol{p_m}$: Mutation rate; $\boldsymbol{p_c}$: Crossover rate; $\boldsymbol{\lambda_{elite}}$: Elite size;
Ensure: $\mathbf{x}^*$: Best solution found, $\mathbf{f}^*$: Log of availability of best solution (via shortest path approximation)

```
n ← |edges[gcfp]|
        ▷ Initialize random shielding configuration population with 20% shielding rate
population ← []
for i = 1 to μ do
    population_i ← random_bitstring(n, 0.2)
end for
x* ← 0^n, f* ← −∞                                   ▷ Best solution and fitness
evals_count ← 0, generation ← 0        ▷ Feasible evaluation and generation counters
while evals_count < B_evals and generation < 2 · B_evals do
    f ← []                                          ▷ Evaluate population
    for i = 1 to μ do
        if evals_count ≥ B_evals then
            break
        end if
        x ← population_i
        if Σ_{k=1}^{n} x_k · cost(edges[gcfp]_k) > B then
            f_i ← −∞
            continue
        end if
        shield ← {edges[gcfp]_j | x_j = 1}
        p, avail ← EvaluateConfig(gcfp, pairs, shield)
        f_i ← min(avail)
        evals_count ← evals_count + 1               ▷ Update global best
        if f_i > f* then
            x* ← x    f* ← f_i
        end if
    end for
    f̂ ← best_k(f, λ_elite)            ▷ Elitism, keep best individuals
    new_population ← {population_i | f_i ∈ f̂}
    while |new_population| < μ and evals_count < B_evals do
                                ▷ Generate offspring via genetic operators
        p_1, p_2 ← SelectParents(population, f)
        c_1, c_2 ← Crossover(p_1, p_2, p_c)
        c_1 ← Mutate(c_1, p_m)
        c_2 ← Mutate(c_2, p_m)
        new_population ← new_population ∪ c_1
        if |new_population| < μ then
            new_population ← new_population ∪ c_2
        end if
    end while
    population ← new_population
    generation ← generation +1
end while
    return x*, f*
```

Algorithm 3 Tournament Selection with Connected Component Weighting

Require: **gcfp**: Cumulative failure probabilities and graph structure; **population**: List of individuals; $\boldsymbol{f}$: Array of fitness values; $\boldsymbol{\nu}$: Tournament size; $\boldsymbol{\lambda}$: Number of parents to select
Ensure: parents: List of selected parent individuals

```
parents ← ∅
for i = 1 to λ do
    I ← []
    for j = 1 to ν do
        I_j ← random(1, |population|)
    end for                          ▷ Normalize fitness within tournament
    f_min ← min_{j∈I} f_j
    f_max ← max_{j∈I} f_j
    Δf ← f_max − f_min                   ▷ Calculate composite scores
    s ← []
    for each j ∈ I do
        if Δf < ε then
            f̃_j ← 0.5
        else
            f̃_j ← (f_j − f_min)/Δf
        end if
        c ← count_connected_components({edges(gcfp)_l | population_jl = 1})
        c_w ← 1/(c + 1)
        s_j ← 0.7 · f̃_j + 0.3 · c_w
    end for
    i* ← arg max_{j∈I} s_j
    parents ← parents ∪ population_{i*}
end for
    return parents
```

C. Greedy Algorithm

We use the greedy approach as a simple baseline. The unshielded edge that is considered the worst (bottleneck) is shielded iteratively until the budget limit is reached, as shown in Algorithm 4. This approach considers the ratio between edge weight and cost. It iteratively selects the edge that is cheapest to shield and has the lowest availability.

IV. Numerical experiments

Numerical experiments were carried out in order to assess the performance of the proposed approach. The Genetic Algorithm variants were tested on seven real graphs under various test settings. The greedy approach was used as a baseline.

Test settings Numerical experiments were performed with both a high and low objective function evaluation budget. In the low budget settings the total number of function evaluations was limited to 500, and in the high settings to 2500. In order to assess the performance in different vertex pair protection scenarios and budget limitations, three problem variants were tested. The first had only one vertex pair connected, in the second we sort all unordered endpoint pairs by decreasing great-circle distance and greedily select the first ten pairs whose endpoints have not previously been selected, and in the third every vertex was connected to every other vertex. The shielding budget allocated to each problem was calculated as half the shielding cost of the minimum approximated steiner tree (Kou's algorithm [14]) connecting all the protected vertices as terminal nodes. The cost of shielding an edge was set as the haversine distance connecting the nodes. The fully connected problem variants in the high computational budget setting were omitted due to computational limitations. Each experiment was repeated 20 times. The Genetic Algorithms had a population size of 50 in every test case, and the greedy approach was iterated until the budget limit was reached.

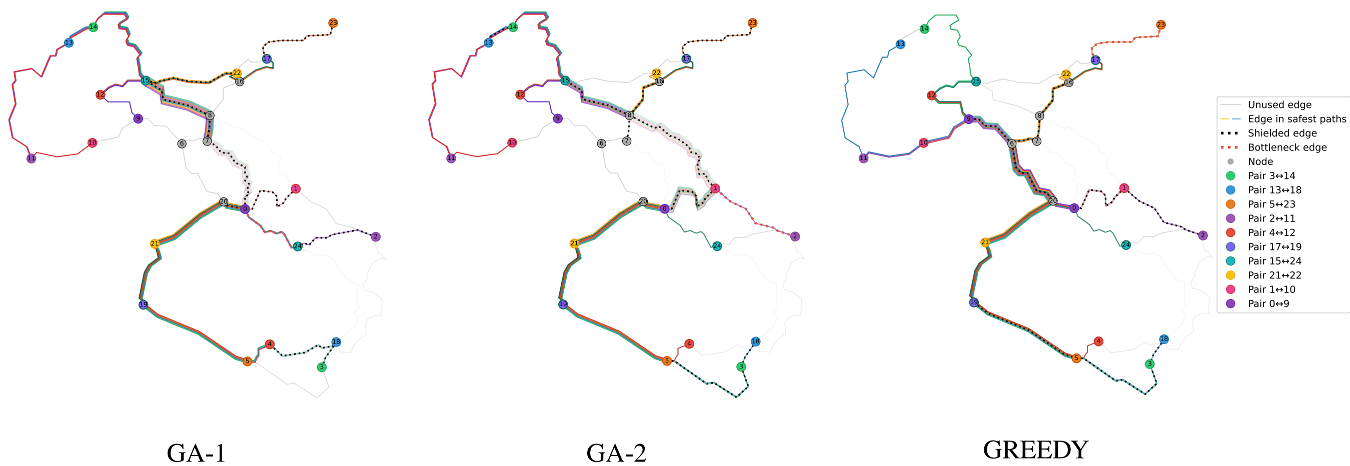

Fig. 2: Resulting shielding configurations for the interroute_j graph in the 10 pairs and 500 evaluations experimental setting. The safest paths between safely connected vertex pairs are highlighted. Shielded edges and the bottleneck edge (lowest availability) are also shown. The edges having a lower availability appear as more transparent.

Algorithm 4 Greedy Algorithm

Require: **gcfp**: Cumulative failure probabilities and graph structure; **pairs**: List of (s, e) vertex pairs to protect; $\boldsymbol{B}$: Budget constraint;
Ensure: $\mathbf{x}^*, \mathbf{f}^*$: greedy solution and availability
  function GREEDY(gcfp,pairs,B)
    $E^* \leftarrow \emptyset$
    $w \leftarrow$ Weights(edges(gcfp), CFP[gcfp])
    paths, avail $\leftarrow$ EvaluateConfig(**gcfp**, **pairs**, E^*)
    while $\sum_{e \in E^*} \texttt{cost}(e) < B$ **do**
        ▷ Consider pairs that need shielding only
      pairs$^* \leftarrow \{s, t \mid \exists\ e \in \text{paths}_{s,t} \text{ s.t } e \notin E^*\}$
      if pairs$^* = \emptyset$ **then**
        break
      end if ▷ Find pair with lowest availability
      p $\leftarrow$ arg $min_{s,t \in pairs^*}(\text{avail}_{s,t})$
        ▷ Find bottleneck edge
      $\text{e}_\text{b} \leftarrow$ arg $max_{e \in \text{paths}_p \setminus E^*}\left(\frac{w_e}{\texttt{cost}(e)}\right)$
      if $\texttt{cost}(e_b) + \sum_{e \in E^*} \texttt{cost}(e) > B$ **then**
        break
      end if
      $E^* \leftarrow E^* \cup \text{e}_\text{b}$ ▷ Shield bottleneck edge if within budget
      paths, avail $\leftarrow$ EvaluateConfig(**gcfp**, **pairs**, E^*)
    end while
    $\mathbf{f}^* \leftarrow min$(availaibilities)
    $\mathbf{x}^* \leftarrow$ *bitmask(E^*)*
    return $\mathbf{x}^*, \mathbf{f}^*$
  end function

A. Benchmarks and parameter tuning

The used real-world backbone networks are presented in Table I. Number of nodes, number of edges, graph density, the source of the graphs are detailed, and the source of the failure data is detailed.

The parameters of the Genetic Algorithm were tuned to the problem. The mutation rate was altered between 0.05 and 0.01, the crossover rate between 0.8 and 0.9, the elite size between 2 and 4 and the tournament size between 2 and 4, resulting in 16 Genetic Algorithm variants. These were tested on three graphs with 500 evaluations. The winning parameter configuration (highest minimum availability score using Wilcoxon rank-sum test) was that of the crossover rate being 0.9, the elite size being 2, a mutation rate of 0.05 and a tournament size of 4.

B. Results

Table II shows the results for the low computational budget test setting, and Table III shows the results for the 2500 function evaluation problem variants. Both tables show the average and standard deviation of the log of the guaranteed availability reached, for 20 independent runs. The best ranking results according to the Wilcoxon rank-sum test are highlighted in bold. The significance level was set to 0.05, and the test was one-sided. In case of an algorithm not being able to find a solution within the shielding budget, the average and standard deviation are denoted by *N/A*. The results point to the fact that both Genetic Algorithm approaches struggle to find a solution within a small budget. In most of the test cases for the 1 pair problem these approaches were not able to bring the solutions to be within the acceptable shielding budget. The greedy approach is more robust in these scenarios. In the cases where there was a higher shielding budget the Genetic Algorithm variants were able to outperform the greedy approach. This can be explained by the fact that a higher budget enables a larger set of possible solutions that they can reach. Both in the 500 and 2500 function evaluation cases the connected component count weighted tournament

TABLE I: Real-world problem inputs

Graph name	$\lvert V \rvert$	$\lvert E \rvert$	Density	Graph source	Failure data
22_optic_eu	22	45	0.19	[15]	[16], [3]
26_usanet	26	43	0.13	[15]	[16], [3]
28_optic_eu	28	41	0.10	[15]	[16], [3]
37_eu	37	57	0.08	[15]	[16], [3]
39_optic_north_american	39	61	0.08	[15]	[16], [3]
79_optic_nfsnet	79	108	0.03	[15]	[16], [3]
interroute_j	25	34	0.11	[16]	[16], [3]

TABLE II: Average and standard deviation of best fitness (log of guaranteed availability) per problem, graph, and algorithm for 500 function evaluations. Best ranking results according to the Wilcoxon rank-sum test are highlighted in bold.

Conn	Graph	GA-1		GA-2		GREEDY	
		Avg	Std	Avg	Std	Avg	Std
1 pair	22_optic_eu	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0001***	2.78e-20
	26_usanet	*N/A*	*N/A*	**-0.0016***	1.70e-03	**-0.0006***	0.00e+00
	28_optic_eu	-0.0005	7.23e-04	*N/A*	*N/A*	**-0.0003***	1.11e-19
	37_eu	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0037***	4.45e-19
	39_optic_n._am.	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0007***	0.00e+00
	79_optic_nfsnet	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0086***	1.78e-18
	interroute_j	**-0.0096***	5.33e-04	-0.0114	1.88e-03	-0.0171	3.56e-18
fully connected	22_optic_eu	**-0.0002***	6.06e-05	-0.0003	9.29e-05	-0.0003	5.56e-20
	26_usanet	**-0.0005***	6.67e-05	-0.0007	1.02e-04	-0.0006	1.11e-19
	28_optic_eu	-0.0001	1.83e-05	-0.0002	2.92e-05	**-0.0001***	0.00e+00
	37_eu	**-0.0003***	5.25e-05	-0.0004	6.19e-05	-0.0007	0.00e+00
	39_optic_n._am.	**-0.0006***	1.54e-04	-0.0009	4.44e-04	-0.0026	8.90e-19
	interroute_j	**-0.0286***	2.75e-03	-0.0338	3.52e-03	-0.0483	1.42e-17
10 pairs	22_optic_eu	**-0.0001***	4.04e-05	-0.0002	5.00e-05	-0.0002	2.78e-20
	26_usanet	**-0.0006***	8.41e-05	-0.0008	3.57e-04	-0.0006	0.00e+00
	28_optic_eu	-0.0001	1.53e-05	-0.0001	2.29e-05	**-0.0001***	1.39e-20
	37_eu	**-0.0002***	3.89e-05	-0.0004	2.06e-04	-0.0005	5.56e-20
	39_optic_n._am.	**-0.0007***	3.27e-04	-0.0015	8.64e-04	-0.0017	2.22e-19
	79_optic_nfsnet	**-0.0012***	7.32e-04	**-0.0018***	9.26e-04	-0.0137	0.00e+00
	interroute_j	**-0.0236***	1.89e-03	-0.0267	2.26e-03	-0.0298	7.12e-18

TABLE III: Average and standard deviation of best fitness (log of guaranteed availability), per problem, graph, and algorithm in case of 2500 function evaluations. Best ranking results according to the Wilcoxon rank-sum test highlighted in bold.

Conn	Graph	GA-1		GA-2		GREEDY	
		Avg	Std	Avg	Std	Avg	Std
1 pair	22_optic_eu	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0001***	2.78e-20
	26_usanet	*N/A*	*N/A*	**-0.0004***	2.21e-05	-0.0006	0.00e+00
	28_optic_eu	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0003***	1.11e-19
	37_eu	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0037***	4.45e-19
	39_optic_n._am.	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0007***	0.00e+00
	79_optic_nfsnet	*N/A*	*N/A*	*N/A*	*N/A*	**-0.0086***	1.78e-18
	interroute_j	**-0.0092***	5.96e-04	**-0.0093***	5.56e-04	-0.0171	3.56e-18
10 pairs	22_optic_eu	**-0.0001***	2.35e-05	**-0.0001***	2.49e-05	-0.0002	2.78e-20
	26_usanet	**-0.0005***	6.92e-05	**-0.0005***	8.92e-05	-0.0006	0.00e+00
	28_optic_eu	**-0.0001***	1.11e-05	**-0.0001***	1.63e-05	-0.0001	1.39e-20
	37_eu	**-0.0001***	2.75e-05	-0.0002	5.47e-05	-0.0005	5.56e-20
	39_optic_n._am.	**-0.0004***	4.14e-05	**-0.0003***	7.05e-05	-0.0017	2.22e-19
	79_optic_nfsnet	**-0.0004***	2.56e-04	*N/A*	*N/A*	-0.0137	0.00e+00
	interroute_j	**-0.0180***	2.15e-03	-0.0206	2.84e-03	-0.0298	7.12e-18

selection based heuristic proved to perform worse than the simple tournament selection, and similar to the simple greedy approach. This can be explained by the fact that selecting for solutions with a low connected component count can trap the algorithm into a local minimum. Figure 2 shows the safest paths and the best shielding configurations found by each approach in the case of interroute_j. The GA-1 found the best one, GA-2 the second best and the GREEDY ranked the third in this case. Considering all problems and the Wilcoxon-rank sum test with all methods, the Genetic Algorithm variant using the simple tournament selection ranked the first with a win rate of 72.5%, the greedy approach ranked second best with 25% win rate. GA-2 came in third and performed slightly worse than the greedy approach.

V. Conclusions

Communication networks are essential for modern society. An apparent failure can cause economic damage and hinder communication. We investigate the budgeted link-shielding problem, where the objective is to select a subset of network links for reinforcement while staying inside a given budget constraint. Reinforced (shielded) links are assumed to remain operational under the considered correlated network element failure scenarios. The goal is to maximize the minimum path availability among a specified set of protected source–destination (endpoint) pairs. To solve this problem we designed two genetic algorithm variants, and a baseline method for comparisons. Numerical experiments conducted on real-world problems show the effectiveness of the designed genetic algorithms.

Further work will address the study of other problem specific operators and the generalization of the proposed problem.

References

[1] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," *IEEE/ACM Trans. Netw.*, vol. 19, no. 6, pp. 1610–1623, 2011.

[2] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "The resilience of WDM networks to probabilistic geographical failures," *IEEE/ACM Trans. Netw.*, vol. 21, no. 5, 2013.

[3] B. Vass, J. Tapolcai, Z. Heszberger, J. Bíró, D. Hay, F. A. Kuipers, J. Oostenbrink, A. Valentini, and L. Rónyai, "Probabilistic shared risk link groups modelling correlated resource failures caused by disasters," *IEEE Journal on Selected Areas in Communications*, 2021.

[4] M. F. Habib, M. Tornatore, F. Dikbiyik, and B. Mukherjee, "Disaster survivability in optical communication networks," *Computer Communications*, vol. 36, no. 6, pp. 630–644, 2013.

[5] B. Todd and J. Doucette, "Fast efficient design of shared backup path protected networks using a multi-flow optimization model," *IEEE Transactions on Reliability*, vol. 60, no. 4, pp. 788–800, 2011.

[6] Z. Nutov, "Approximating node-connectivity augmentation problems," *Algorithmica*, vol. 63, no. 1, pp. 398–410, 2012.

[7] J. Zhang, E. Modiano, and D. Hay, "Enhancing network robustness via shielding," *IEEE/ACM Transactions on Networking*, 2017.

[8] Z. Wang, Q. Wang, M. Zukerman, and B. Moran, "A seismic resistant design algorithm for laying and shielding of optical fiber cables," *Journal of Lightwave Technology*, vol. 35, no. 14, pp. 3060–3074, 2017.

[9] B. Tao, M. Xiao, B. Khoussainov, and J. Peng, "Optimal shielding to guarantee region-based connectivity under geographical failures," in *IEEE INFOCOM 2022-IEEE conference on computer communications*. IEEE, 2022, pp. 1109–1118.

[10] A. Maulana, M. Kefalas, and M. T. Emmerich, "Immunization of networks using genetic algorithms and multiobjective metaheuristics," in *2017 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2017, pp. 1–8.

[11] R. F. Rodrigues, A. R. da Silva, V. da Fonseca Vieira, and C. R. Xavier, "Optimization of the choice of individuals to be immunized through the genetic algorithm in the sir model," in *International Conference on computational science and its applications*. Springer, 2018, pp. 62–75.

[12] B. Dengiz, F. Altiparmak, and A. Smith, "Local search genetic algorithm for optimal design of reliable networks," *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 3, pp. 179–188, 1997.

[13] Z. Tasnádi, B. Vass, and N. Gaskó, "Ant colony optimization algorithm for safest path computation in presence of correlated failures in backbone networks," in *ACM Genetic and Evolutionary Computation Conference (GECCO)*, Málaga, Spain, Jul. 2025.

[14] L. Kou, G. Markowsky, and L. Berman, "A fast algorithm for steiner trees," *Acta Informatica*, vol. 15, no. 2, pp. 141–145, Jun 1981.

[15] J. Tapolcai, L. Rónyai, B. Vass, and L. Gyimóthi, "List of shared risk link groups representing regional failures with limited size," in *IEEE INFOCOM*, Atlanta, USA, May 2017.

[16] A. Valentini, B. Vass, J. Oostenbrink, L. Csák, F. Kuipers, B. Pace, D. Hay, and J. Tapolcai, "Network resiliency against earthquakes," in *2019 11th International Workshop on Resilient Networks Design and Modeling (RNDM)*, Oct 2019, pp. 1–7.