Enhancing Network Robustness via Shielding

Jianan Zhang *, Eytan Modiano * and David Hay †
*Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA, USA
†School of Engineering and Computer Science, Hebrew University, Jerusalem, Israel

***Abstract*—We consider shielding critical links to guarantee network connectivity under geographical and general failure models. We develop a mixed integer linear program (MILP) to obtain the minimum cost shielding to guarantee the connectivity of a single SD pair under a general failure model, and exploit geometric properties to decompose the shielding problem under a geographical failure model. We extend our MILP formulation to guarantee the connectivity of the entire network, and use Benders decomposition to significantly reduce the running time by exploiting its partial separable structure. We also apply simulated annealing to solve larger network problems to obtain near-optimal solutions in much shorter time. Finally, we extend the algorithms to guarantee partial network connectivity, and observe significant reduction in shielding cost, especially when the failure region is small. For example, when the failure region radius is 60 miles, we observe as much as 75% reduction in shielding cost by relaxing the connectivity requirement to 95% on a major US infrastructure network.**

I. Introduction

Communication networks are subject to natural disasters and attacks, such as hurricanes, earthquakes, electromagnetic pulse attacks [1]. Network failures may result in tremendous financial loss and hinder effective recovery to the affected regions. Therefore, it is important for network designers to guarantee that the network can withstand failures that may result from disasters or attacks.

Several metrics measure the performance of the network. The most fundamental requirement is connectivity, without which it is impossible to support any application that requires communication through the network. Another metric, important for quality of service guarantee, is the maximum amount of traffic carried by the network. In case of network failures, one cannot expect the network to support the same amount of traffic as before the failure. However, low priority applications such as movies can be throttled to give higher priority to critical applications in case of network failures. In this paper, we focus on guaranteeing network connectivity, and assume that networks are able to use limited resources to support critical applications using service differentiation.

Previous research considers geographical failures [2], [3] and general failures [4], [5] to assess the robustness of the network. Geographical failure models capture the effects of natural disasters and physical attacks; e.g., all links in the failure region are destroyed. Under the general failure model, an arbitrary set of links may fail, i.e., each failure may affect a specified set of links, whose number and location is determined by the nature of the failure.

A common approach to design robust networks is through redundancy and backup routes (see [6], [7] for a survey of protection techniques for optical networks). An alternative approach, which we consider in this paper, is through shielding of critical links. Shielded network infrastructure can survive disasters and attacks. Previous research suggests strengthening cables to resist physical attacks [7], and upgrading or covering vulnerable components to resist electromagnetic pulse attacks [8]. More robust optical fibers and cables are being developed to improve network reliability [9], [10]. Recently, Google announced to reinforce undersea cables to resist attacks and avoid frequent repairs [11].

Due to the cost of shielding, it may not be economical to shield the entire network. Instead, critical parts of the network can be identified and shielded to guarantee network robustness. Previous work identifies critical parts to shield in order to achieve certain performance objectives in various applications [12]–[17]. The authors in [12] design optimal topologies, given different levels of shielding cost, link construction cost, and utility of network connectivity, under the assumption of uniform costs for all links. In [13], the authors formulate a road network retrofit problem, and use a two stage stochastic programming approach to decide which roads to retrofit to minimize the average performance loss incurred by a disaster. Fortifying facilities to minimize the transportation cost and path length is also considered in [14], [16], [17]. However, no previous research is devoted to identifying the minimum cost shielding to guarantee network connectivity. The shielding problem we consider shares similarities with fixed charge problems [18], [19], where the fixed costs of using network resources make the selection of resources difficult.

In this paper, we aim to design robust networks by shielding critical parts of the network. We determine the minimum cost shielding to guarantee that the network remains connected after a failure, under both geographical and general failure models. First, we consider a single source-destination (SD) pair in the network and determine the minimum cost shielding to guarantee its connectivity. We develop a mixed integer linear program (MILP) to formulate the shielding problem in the case of general failures, and identify properties of optimal shielding in the case of geographical failures to decompose the shielding problem to multiple subproblems, each of which determines

This work was supported by NSF grant CNS-1017800, DTRA grants HDTRA-09-1-0050, HDTRA1-14-1-0058, and by a grant from the U.S.-Israel Binational Science Foundation and the Israeli Centers of Research Excellence (I-CORE) program (Center No. 4/11).

978-1-4799-7795-6/15/$31.00 ©2015 IEEE

the optimal shielding for a disjoint set of links. Then, we extend the MILP formulation to consider guaranteeing the connectivity of the entire network. By identifying the partial separable structures of the MILP, we apply the Benders decomposition technique [20] to reduce the running time and solve network shielding problems of realistic size. A heuristic based on simulated annealing further reduces the running time significantly while achieving good results. Moreover, we observe that shielding cost can be significantly reduced if the connectivity requirements are slightly relaxed.

II. Failure models and network shielding

Geographical failure models can be used to model real world disasters and attacks [2], [3], [21]. In this paper, we consider the disk failure model, which captures the effect of electromagnetic pulse attacks. A disk failure with a given radius may occur anywhere in the network, and all the links intersecting the disk region are affected. Multiple failures can be represented by one failure which dominates them, where a failure dominates another if it affects all the links affected by the other failure. The number of dominating failures is polynomial in the number of links and can be efficiently obtained by computational geometry techniques [2], [3].

In addition, we consider a general failure model that represents the failures of shared risk groups [4], [22]. Instead of being limited to be within a geographical region, the set of failed links can be arbitrary, possibly restricted by the nature of the attack or disaster. Under this model, the possible failures and links affected by each failure are described explicitly.

We aim to shield links by using the minimum cost to guarantee network connectivity after any single failure event. For simplicity we assume that shielded links do not fail. We first consider guaranteeing the connectivity of a single SD pair, and later extend to the connectivity of the entire network. Finally we relax the connectivity requirement for the entire network to allow for partial connectivity, which requires much less shielding. Throughout this paper all nodes are assumed to be reliable. Straightforward extensions of the proposed approaches can be applied to node shielding as well.

III. Guaranteeing connectivity of a single SD pair

We start by considering the shielding problem in order to guarantee the connectivity of a single SD pair. It suffices to shield links to guarantee that a path will exist after any failure event. Clearly, if only one link can fail at a time, the links in the minimum-cut-equals-1-set need to be shielded. The minimum-cut-equals-1-set is the set of links among which any single link failure disconnects the SD pair. Any other single link failure will not disconnect the SD pair and need not be shielded. However, if a failure event affects several links and disconnects the SD pair, not all the affected links need to be shielded in order to guarantee a path between the SD pair after the failure event. We aim to determine the links that need to be shielded with minimum cost to guarantee the connectivity of the SD pair after any failure under both the general and geographical failure models.

A. Shielding under the general failure model

Under the general failure model a failure is specified by the set of failed links. The network is represented by a graph $G=(V,E)$. Each failure z affects a set of links $E^{(z)}$. The objective is to shield a set of links E^* with minimum shielding cost, to guarantee that s and d are connected through $G=(V,E'^{(z)})$ for all z, where $E'^{(z)}=(E\setminus E^{(z)})\cup E^*$.

The optimal shielding problem under the general failure model can be modeled by a MILP. Let $t_{ij}^{(z)}$ indicate whether or not failure z affects link (i,j). Failure z affects a set of links $E^{(z)}=\{(i,j)|t_{ij}^{(z)}=1\}$. Each link (i,j) has shielding cost c_{ij}. Both $t_{ij}^{(z)}$ and c_{ij} are problem parameters. The decision variables are $x_{ij}^{(z)}$ and h_{ij}, which represent the amount of flow carried on link (i,j) after failure z and whether or not to shield link (i,j), respectively. The set of shielded links is $E^*=\{(i,j)|h_{ij}=1\}$. Since the links are undirected, (i,j) is the same link as (j,i), in which case $c_{ij}=c_{ji}$ and $h_{ij}=h_{ji}$. The minimum shielding cost to resist any possible failure is given by the following MILP.

$$\min \quad \sum_{(i,j)\in E} c_{ij}h_{ij}/2 \tag{1}$$

$$\text{s.t.} \quad \sum_{\{j|(i,j)\in E\}} x_{ij}^{(z)} - \sum_{\{j|(j,i)\in E\}} x_{ji}^{(z)} = \begin{cases} 1, & \text{if } i=s \\ -1, & \text{if } i=d \\ 0, & \text{otherwise} \end{cases} \quad \forall z \tag{2}$$

$$x_{ij}^{(z)} - h_{ij} \le 1 - t_{ij}^{(z)} \qquad \forall (i,j)\in E, z \tag{3}$$

$$h_{ij} - h_{ji} = 0 \qquad \forall (i,j)\in E$$

$$x_{ij}^{(z)} \ge 0 \qquad \forall (i,j)\in E, z$$

$$h_{ij} = \{0,1\} \qquad \forall (i,j)\in E$$

Since we consider a connectivity problem, only unit flow need to be carried from s to d, which is guaranteed by the flow constraints (2). Constraints (3) guarantee that in case failure z occurs and affects link (i,j) ($t_{ij}^{(z)}=1$), unit flow can be carried on link (i,j) only if it is shielded ($h_{ij}=1$). If link (i,j) is not affected by failure z ($t_{ij}^{(z)}=0$), it can carry unit flow regardless of shielding in case of failure z. The factor of 1/2 in the objective accounts for the fact that each shielded link is counted twice ($h_{ij}=h_{ji}=1$).

The above algorithm can be applied to obtain the optimal shielding under the general failure model. For example, in Fig. 1, which represents the topology of the XO communication backbone network and consists of 60 nodes and 71 links [23], we consider the failure model where all the links incident to any two nodes are affected by a failure, and the number of failures is $\binom{60}{2}$ (i.e., this model allows for link failures incident to up to 2 nodes). The cost of shielding each link is represented by the length of the link (in latitude/longitude degree unit). Given the SD pair Seattle-Miami, the optimal shielding obtained by the algorithm is represented by the thick links with total cost 45.98.

B. Shielding under the geographical failure model

We model geographical failures as disks with a given radius (i.e., all links intersected by the disk region fail). Given a

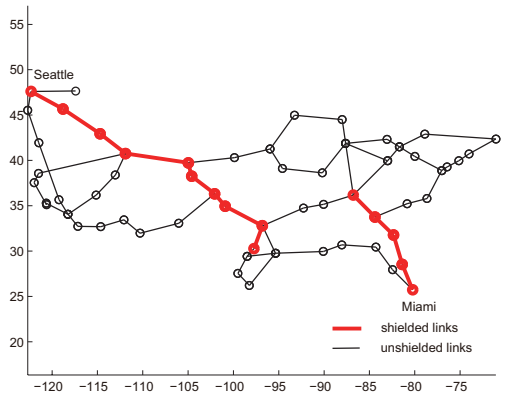

Fig. 1. Optimal shielding under the failure model where all the links incident to any two nodes are affected by a failure.

network topology and an SD pair, it is necessary to identify the set of geographical failure regions, each of which disconnects the given SD pair. We call such regions bottleneck regions. Finding the bottleneck regions can be accomplished by checking whether the disk failure disconnects the SD pair [24]. Since the number of dominating failure regions is polynomial in the number of links [2], [3], this task can be done in polynomial time. In the following we exploit properties of bottlenecks to decompose the shielding problem to several subproblems, each of which consists one or more bottlenecks and can be solved independently.

In order to guarantee a path between the given SD pair, for each bottleneck a shielded path that starts and ends outside the failure region must exist. Otherwise, there would be no path going through this bottleneck after the failure occurs and the SD pair is disconnected. If the shielded path is part of a path between the SD pair, such a shielded path is sufficient to guarantee that the SD pair is connected after a disk failure occurs at this bottleneck.

We start by describing a simple algorithm that can be used to find a shielded path to guarantee the connectivity of an SD pair after a failure occurs at a bottleneck. We illustrate the algorithm by using the example in Fig. 2.

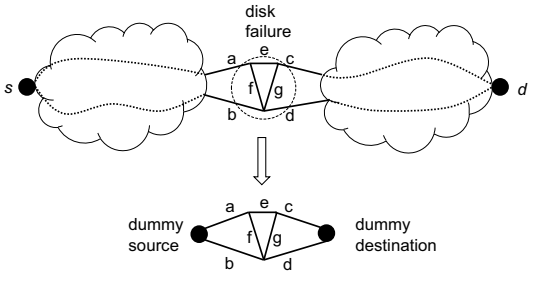

Fig. 2. A single bottleneck.

Based on previous analysis, we have the following claim.

Algorithm 1 Bottleneck Shielding Algorithm

1) Among the links intersected by the bottleneck region, find the links that are connected to the source node without going through any link in the bottleneck (links a, b in Fig. 2). A link is connected to the source if a path exists between the source and one end node of the link. These links have to cross the boundary of the failure region. Among the two end nodes of each link, there is one node outside the failure region. Merge these nodes to form a dummy source, as shown in Fig. 2.
2) Find the links that are connected to the destination without going through any link in the bottleneck (links c, d in Fig. 2), and merge their ends outside the failure region to form a dummy destination.
3) Among all the links intersected by the bottleneck, shield a path between the dummy SD pair. This path will survive the failure affecting this bottleneck.

Claim 1. *A shielded path between the dummy SD pair is necessary and sufficient to guarantee the original SD pair connectivity after a disk failure occurs at this bottleneck.*

If there is only one bottleneck between an SD pair, we only need to shield a "shortest path" between the dummy SD pair, where the "length" of each link represents its shielding cost. However, generally there may be multiple bottlenecks. Recall that each bottleneck includes a set of links which can fail simultaneously and whose failure disconnects the SD pair. In order to guarantee the connectivity of the SD pair in case of any disk failure, it is necessary to shield a path through every bottleneck. If the bottlenecks are disjoint and do not share common links (Fig. 3), shielding the shortest path between each dummy SD pair is optimal.

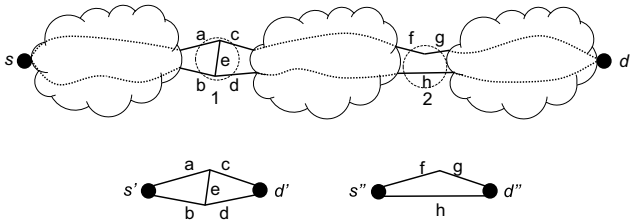

Fig. 3. Non-overlapping bottlenecks.

However, different bottlenecks may overlap and share common links (Fig. 4). Shielding the links in one bottleneck may affect the shielding for another bottleneck. For example, between the first dummy SD pair (s', d'), shielding link c also leads to a shielded link c between the second dummy SD pair (s'', d''). Thus it is necessary to consider all the overlapping bottlenecks jointly.

Nevertheless, if a set of overlapping bottlenecks do not share common links with another set of overlapping bottlenecks, these two sets can be considered separately, because shielding decisions for one set do not affect the shielding decisions for the other in order to shield a path in each bottleneck.

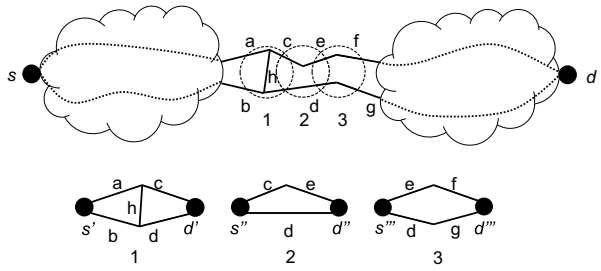

Fig. 4. Overlapping bottlenecks.

The optimal shielding for a set of overlapping bottlenecks is given by MILP (1) which includes constraints only associated with the failures in this set. In short, under the geographical failure model, instead of considering all the failures at once, the problem can be decomposed to multiple smaller MILPs (one per overlapping bottlenecks set) that can be solved more efficiently.

We illustrate the algorithm using the network in Fig. 5, which is the same as in Fig. 1, where now a failure is any disk with radius 2° (about 120 miles). Given the SD pair Seattle-Miami, there are 4 bottleneck regions represented by dashed circles, and all the links intersected by the dashed circles are candidate links among which shielding decisions are made. In each of the two disjoint bottlenecks, a shortest path is shielded, represented by the thick links with costs 5.12 and 5.88, respectively, while the overlapping bottleneck has shielding cost 10.00, yielding to a total shielding cost 21.00.

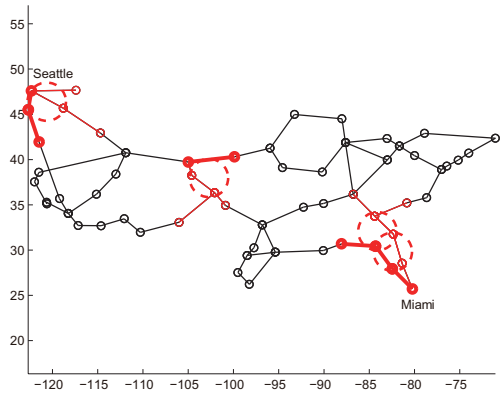

Fig. 5. Bottlenecks and shielded links given disk failure radius 2°.

IV. Guaranteeing connectivity of the entire network

In backbone networks, where nodes represent routers, it is important to guarantee that all the nodes are connected. Since links are shared by many SD pairs, shielding links to guarantee the connectivity of one SD pair may benefit another SD pair. The union of the optimal shielding for each SD pair may not be the optimal shielding for the network. In fact, all the SD pairs must be considered jointly in order to determine the optimal shielding.

A. Shielding under the geographical failure model - the cases of huge and tiny failures

We start by considering two special cases of geographical failures. In the first case, the failure region is huge and contains all the links in the network. Unshielded links are destroyed by a failure event. In this case, in order to keep all the nodes connected, one must shield at least a minimum spanning tree, where the weight of each link is its shielding cost.

In the second case, the failure region is tiny, which affects either a single link, or the links incident to one node. If the network has tree structure, every link's failure would disconnect the network. Therefore all the links in the tree structure have to be shielded. On the other hand, if nodes form a cycle, removing the links incident to one node does not affect the connectivity of the other nodes (recall that we only consider one failure at a time). Therefore, it suffices to guarantee that each node is incident to at least one shielded link, which connects this node with the remaining nodes after any failure. For a cycle, the optimal set of links to shield is its minimum edge cover. Minimum edge cover of a graph is a set of edges of minimum weight such that every node in the graph is incident to at least one edge in the set. The calculation of minimum edge cover takes polynomial time and can be obtained by calculating the maximum matching in a transformed graph [25]. The detailed algorithm can be found in Chapter 2 of [26].

B. Shielding under the general failure model

Finally we consider optimal shielding to guarantee that the entire network is connected under the general failure model. The network is represented by a graph $G=(V,E)$. Each failure z affects a set of links $E^{(z)}$. Our objective is to shield a set of links E^* using minimum shielding cost, to guarantee that $G=(V,E'^{(z)})$ is connected for all z, where $E'^{(z)}=(E\setminus E^{(z)})\cup E^*$.

The MILP formulation for this problem is similar to the formulation for the single SD pair connectivity problem, except that the constraints guarantee the connectivity of the entire network instead of a single SD pair. The variables and parameters have the same meaning as in MILP (1).

$$\min \quad \sum_{(i,j)\in E} c_{ij}h_{ij}/2 \tag{4}$$

$$\text{s.t.} \quad \sum_{\{j|(i,j)\in E\}} x_{ij}^{(z)sd} - \sum_{\{j|(j,i)\in E\}} x_{ji}^{(z)sd} = \begin{cases} 1, & \text{if } i=s \\ -1, & \text{if } i=d \\ 0, & \text{otherwise} \end{cases} \quad \forall z,s,d$$

$$\begin{aligned} & x_{ij}^{(z)sd} - h_{ij} \le 1 - t_{ij}^{(z)} && \forall (i,j)\in E, z, s, d \\ & h_{ij} - h_{ji} = 0 && \forall (i,j)\in E \\ & x_{ij}^{(z)sd} \ge 0 && \forall (i,j)\in E, z, s, d \\ & h_{ij} = \{0,1\} && \forall (i,j)\in E \end{aligned} \tag{5}$$

In MILP (4), for each SD pair and failure scenario, there is a flow variable for each link. The number of variables is

very large since there are many possible failure scenarios. It is difficult to directly solve MILP (4) for large problem instances. However, the flow variables after one failure couple with the flow variables after another failure only through the decision variables h in (5). Given h, it is easy to determine whether there are feasible flows between all the SD pairs after each failure, by only considering the flow variables and constraints related to each failure. Benders decomposition can be applied to problems with such partial separable structure.

1) Benders decomposition: Benders decomposition accelerates the calculation of an optimization problem with partial separable structure and many constraints, and has been applied to resilient network design [19], [20]. Instead of considering all the constraints at once, it first solves a relaxed problem that has only a few constraints, and then check whether there are any violated constraints. If there is none, the solution is optimal. Otherwise, a violated constraint is added to the relaxed problem and the problem is solved again. The relaxed problem is called the master problem, and the violated constraints are identified by solving subproblems.

The MILP (4) can be reformulated as follows. It starts with a master problem with constraints only on h.

$$\begin{aligned} \min \quad & \sum_{(i,j)\in E} c_{ij}h_{ij}/2 \\ \text{s.t.} \quad & h_{ij}-h_{ji}=0 \quad \forall (i,j)\in E \\ & h_{ij}=\{0,1\} \quad \forall (i,j)\in E \end{aligned}$$

After obtaining h, check whether there are violated constraints by solving subproblems, each corresponding to checking whether the network is connected after each failure. If the linear program (LP) (6) is feasible and has optimal value 0, the network is connected after failure z. If it is infeasible, the associated constraint has been violated.

$$\min \quad 0 \tag{6}$$

$$\text{s.t.} \sum_{\{j|(i,j)\in E\}} x_{ij}^{(z)sd} - \sum_{\{j|(j,i)\in E\}} x_{ji}^{(z)sd} = \begin{cases} 1, \text{ if } i=s \\ -1, \text{ if } i=d \\ 0, \text{ otherwise} \end{cases} \quad \forall s,d \tag{7}$$

$$x_{ij}^{(z)sd} - h_{ij} \le 1 - t_{ij}^{(z)} \qquad \forall (i,j)\in E, s, d \tag{8}$$

$$x_{ij}^{(z)sd} \ge 0 \qquad \forall (i,j)\in E, s, d$$

It is more efficient to add the violated constraint by considering the dual of LP (6). The dual is represented by LP (9), where dual variables p and r corresponds to primal constraints (7) and (8), respectively. If LP (9) is unbounded, the constraint $\sum_{sd}[p_d^{*(z)sd} - p_s^{*(z)sd} - \sum_{(i,j)\in E}(1-t_{ij}^{(z)}+h_{ij})r_{ij}^{*(z)sd}] \le 0$ is added to the master problem, where $(p_d^{*(z)sd}, p_s^{*(z)sd}, r_{ij}^{*(z)sd})$ is an extreme ray of LP (9) which led to its unboundedness. The new constraints avoid unbounded costs along extreme rays, to guarantee that LP (9) is bounded and that LP (6) is feasible.

$$\max \quad \sum_{sd}[p_d^{(z)sd} - p_s^{(z)sd} - \sum_{(i,j)\in E}(1-t_{ij}^{(z)}+h_{ij})r_{ij}^{(z)sd}] \tag{9}$$

$$\begin{aligned} \text{s.t.} \quad & p_j^{(z)sd} - p_i^{(z)sd} - r_{ij}^{(z)sd} \le 0 \quad \forall (i,j)\in E, s, d \\ & r_{ij}^{(z)sd} \ge 0 \qquad \forall (i,j)\in E, s, d \end{aligned}$$

While for the original problem, the MILP has a large number of variables and constraints, with Benders decomposition it is possible to solve each subproblem using an LP, and the size of each subproblem is small. In Benders decomposition, checking whether a subproblem is bounded corresponds to checking whether the network is connected after one failure in our problem. Instead of checking only one failure after obtaining a new shielding decision in each iteration, multiple failures can be checked. This is particularly helpful in our problem since different failures affect different links, and the links that need to be shielded are likely to be different. The number of violated constraints added before resolving the master problem provides a tradeoff between the number of master iterations and the running time of each iteration. In our numerical evaluations, the number of constraints that we added was equal to the number of nodes in the network, and we observed more than 50% running time saving compared with the standard Benders decomposition algorithm.

2) Simulated annealing: Finally, we developed a heuristic based on simulated annealing [27], [28] to solve the problems faster. Simulated annealing is a method to search for globally optimal solutions for nonconvex optimization problems. It starts at an initial state, and then aims to find a neighbor state, preferably a state with smaller cost. If such a neighbor state with smaller cost is found, the current state is replaced with the neighbor state. Otherwise, if the neighbor state has larger cost, the current state is replaced with the neighbor state with some small probability. Simulated annealing avoids being stuck in a local minimum without continuing further searches. The probability to replace the current state with a neighbor state that has higher cost depends on the difference of the costs of the two states. The higher the cost of the neighbor state, the less likely to enter this state.

Let S be the set of shielded links and S^c be the set of unshielded links in the current state. Initially, all the links are shielded. In order to find a neighbor state which differs from the current state in only one or two shielded links, one of the following three operations is performed:

1) Randomly remove one link from S.
2) Randomly remove one link from S, and randomly shield one more link from S^c.
3) Randomly shield one more link from S^c.

In these operations, the probability of removing a link is proportional to the shielding cost of the link. Thus, links having larger shielding costs are more likely to be removed. The probability of adding a link to the shielded set is proportional to the multiplicative inverse of its shielding cost, so that links with small shielding costs are more likely to be added.

Since the objective is to find a neighbor state with smaller shielding cost, the operations are done sequentially during the first few iterations of simulated annealing. For example, after one shielded link is removed (operation 1), if the current shielding is feasible, a neighbor state with smaller shielding cost is identified. If any removal of shielded link leads to an infeasible shielding, one more shielded link is added after the removal (operation 2) in search of a feasible shielding state.

If neither works, one more shielded link is added without removing any shielded link (operation 3), in order to retain a new feasible shielding state.

After finding the neighbor (new) state, next determine whether to replace the current state with the neighbor state. If the neighbor state has smaller shielding cost than the current state, it replaces the current state. On the other hand, if the neighbor state has larger shielding cost compared with the current state, it replaces the current state with probability $p = exp(-\delta/T)$, where δ is the difference between the shielding costs of the two states and T is the temperature. It is possible to enter a state with higher shielding cost and avoid being stuck at a local minimum. If the neighbor state is rejected (with probability $1-p$), perform the last operation (2 or 3) from the current state and try again.

During the first fewer iterations, T is large so that it is easy to enter a state that has larger shielding cost to explore more possible states. As the number of iterations increases, T decreases to make it less likely to enter a state with larger shielding cost. At last, T tends to 0 and the algorithm terminates at a state which has smaller shielding cost compared to all its neighbors. As suggested by [28], $T(t) = d/\log(t+1)$, where t is the number of iterations and d is a positive constant. Large enough d, which is at least the amount of cost increase at any state along the way from any local minimum to the global minimum, guarantees convergence to the global minimum at the cost of more iterations for annealing. We set $d = 10$ in our algorithm to enable the escape from a local minimum with a reasonable probability and avoid large computational cost.

If the operations to find a neighbor state are done sequentially, the algorithm may end up in cycles. For example, it is possible that both operations 1 and 2 cannot find a feasible shielding state, and the possible neighbor state is to shield an extra link. Starting from the neighbor state, the only link that can be removed without causing infeasible shielding is the link that was just added. Therefore, after some iterations when many redundantly shielded links are removed, the operations are done randomly to avoid such cycles.

V. Guaranteeing partial connectivity

In most networks, a significant number of links need to be shielded to guarantee the full connectivity of the network. In fact, even in the tiny disk failure case, links need to be shielded according to minimum edge covering for cycle structures. The number of shielded link is at least half the number of nodes. In larger failure cases or if the network has tree structures, even more links need to be shielded.

If the connectivity constraint is relaxed and some nodes are allowed to be separated from the others, shielding cost may be significantly reduced. The reduction in shielding depends on the failure model and network topology. For example, if one node is allowed to be disconnected from the rest, in cycle structures no link need to be shielded in the tiny disk failure case, since only the node within the failure region is disconnected from the others. Similarly, in tree structure, links incident to degree 1 nodes do not need to be shielded, since the failure of a link incident to degree 1 node only separates a degree 1 node from the others.

We determine the optimal shielding to guarantee partial network connectivity under the general failure model, using average two terminal reliability (ATTR) as a measure of the connectivity level. ATTR is calculated by dividing the number of connected SD pairs after a failure by the total number of SD pairs in the original network, and represents the fraction of SD pairs connected after a failure. Compared to MILP (4), the unit flow constraints are not imposed to every SD pair. Instead, only a fraction of SD pair are guaranteed to carry unit flow. In constraints (11), I^{sd} can either take the value 0 or 1, where $I^{sd} = 1$ guarantees the connectivity of the SD pair. The total number of connected SD pair should be at least a fraction α of all the $N(N-1)/2$ SD pairs, guaranteed by constraints (12), where N is the total number of nodes.

$$\min \quad \sum_{(i,j)\in E} c_{ij} h_{ij}/2 \tag{10}$$

$$\text{s.t.} \sum_{\{j|(i,j)\in E\}} x_{ij}^{(z)sd} - \sum_{\{j|(j,i)\in E\}} x_{ji}^{(z)sd} = \begin{cases} I^{(z)sd}, & \text{if } i = s \\ -I^{(z)sd}, & \text{if } i = d \\ 0, & \text{otherwise} \end{cases} \quad \forall z,s,d \tag{11}$$

$$\begin{aligned} & x_{ij}^{(z)sd} - h_{ij} \le 1 - t_{ij}^{(z)} && \forall (i,j) \in E, z, s, d \\ & \sum_{sd} I^{(z)sd} \ge \alpha N(N-1)/2 && \forall z \\ & h_{ij} - h_{ji} = 0 && \forall (i,j) \in E \\ & x_{ij}^{(z)sd} \ge 0 && \forall (i,j) \in E, z, s, d \\ & h_{ij} = \{0,1\} && \forall (i,j) \in E \\ & I^{sd} = \{0,1\} && \forall (i,j) \in E \end{aligned} \tag{12}$$

The above MILP has more variables and constraints than MILP (4). First, there are the additional variables I^{sd}. Moreover, in MILP (4), we only need to check the connectivity between node 1 and nodes $2,3,\ldots,N$, which is enough to guarantee the connectivity of the entire network. However, in MILP (10), we need to check $N(N-1)/2$ SD pairs. We again use simulated annealing to find near-optimal solutions. The only difference compared with the algorithm which guarantees full connectivity is in determining whether the shielding is feasible. As long as the ATTR is above α, the shielding is feasible and is a candidate for the next state.

VI. Numerical results

A. Full connectivity

We first compare the running time of solving the MILP using Benders decomposition, solving the MILP directly, and solving its LP relaxation in Fig. 6. The results are averaged over 10 instances of Erdos-Renyi random graphs, which suffice to show the growth of the running time of these approaches. The number of nodes of the graph is varied from 10 to 30, with average degree 5. We consider failures that affect the links adjacent to two nodes. The number of possible attacks for each graph is $\binom{N}{2}$, where N is the number of nodes in a graph. Note that solving the MILP directly and solving the MILP using

Benders decomposition both give the optimal solutions, while solving the LP relaxation only gives lower bounds of optimal shielding costs. It can be observed that solving the MILP using Benders decomposition works even faster than solving the LP relaxation of the MILP directly for larger networks.

Moreover, the modified Benders decomposition (adding multiple violated constraints in each iteration) reduces the running time further by more than 50% in most cases as shown in Table I.

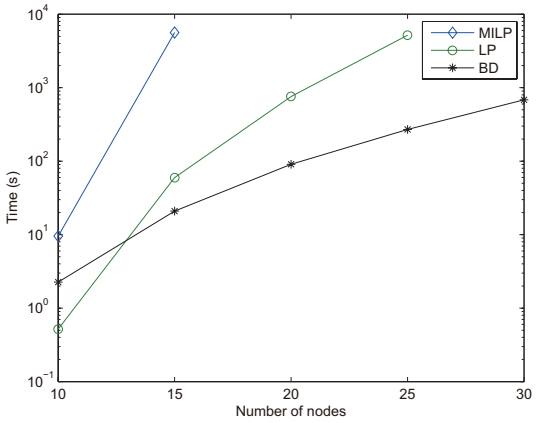

Fig. 6. Running time comparisons of Benders decomposition, directly solving MILP and its LP relaxation.

Next we compare the performance of simulated annealing with the modified Benders decomposition. We observe in Tables I and II that the running time for simulated annealing is about 1/10 of that of modified Benders decomposition in larger network cases, while the relative error is only $3\% \sim 6\%$.

TABLE I
RUNNING TIME COMPARISONS OF SA, BD AND THE MODIFIED BD FOR RANDOM GRAPHS

Number of nodes	Degree	SA Time (s)	BD Time (s)	Modified BD Time (s)
10	5	0.79	2.26	1.54
15	5	1.70	20.89	10.85
20	5	3.96	90.77	34.96
25	5	10.86	270.69	103.32
30	5	19.20	684.83	195.20

TABLE II
SHIELDING COST COMPARISONS OF SA RESULTS AND EXACT SOLUTIONS FOR RANDOM GRAPHS

Number of nodes	Degree	SA Cost	Exact Cost	Relative Error
10	5	67.2	64.8	0.037
15	5	141.4	136.2	0.038
20	5	248.0	240.0	0.033
25	5	394.4	374.2	0.054
30	5	551.4	532.8	0.035

Finally we apply our algorithm to obtain the optimal shielding for the XO communication network. Fig. 7 illustrates the optimal shielding to guarantee the connectivity of the entire XO backbone network after any disk failure with radius 1°.

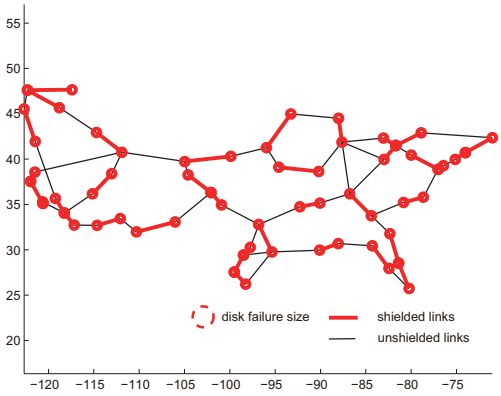

Fig. 7. Optimal shielding in the XO network to guarantee full connectivity after any disk failure with radius 1°.

Simulated annealing also has good performance in solving the shielding problem for the XO network. The results are shown in Table III. This advocates the use of simulated annealing in larger size problems where the exact solution is too expensive to obtain.

TABLE III
COMPARISON OF SA AND THE MODIFIED BD ALGORITHMS FOR THE XO NETWORK

Attack radius	SA Time (s)	Modified BD Time (s)	SA Cost	Modified BD Cost	Relative error
1°	51.98	109.19	106.6	99.5	0.071
2°	64.73	875.12	129.3	121.3	0.065

B. *Partial connectivity*

The MILP for partial connectivity involves a large number of variables and constraints, and can only be solved for small problem instances (were able to solve for networks which have fewer than 15 nodes). On the other hand, simulated annealing algorithm which guarantees partial connectivity has comparable running time with the algorithm which guarantees full connectivity. The only difference is in determining whether the shielding is feasible. The shielding costs obtained by the simulated annealing algorithm are nearly identical with the exact solutions by solving the MILP for different levels of ATTR requirement, and are omitted.

Shielding cost is significantly reduced by relaxing the connectivity constraint to allow $\alpha = 95\%$. This corresponds to the case that one node is allowed to be disconnected from the others in XO network ($\alpha = 29/30 > 95\%$). Figure 8 depicts the shielded links in the case where the disk failure has radius 1°. Table IV suggests that the cost reduction is larger for smaller failure, and depicts the shielding cost for $\alpha = 92\%$, which corresponds to the case that two nodes are allowed to be disconnected.

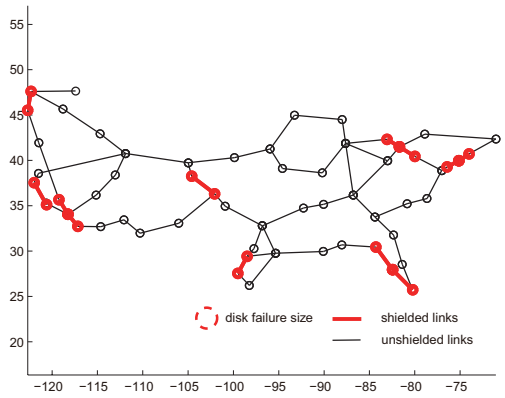

Fig. 8. Optimal shielding in the XO network to guarantee that $\alpha = 95\%$ after any disk failure with radius 1°.

TABLE IV
COST REDUCTION OF PARTIAL CONNECTIVITY FOR XO NETWORK

Attack radius	SA Cost (full connectivity)	SA Cost ($\alpha = 95\%$)	SA Cost ($\alpha = 92\%$)
1°	106.5	26	5
2°	129	86	49.7

VII. CONCLUSION

In this paper we considered the network shielding problem under geographical and general failure models. We developed MILP formulations to obtain the minimum cost shielding to guarantee the connectivity of a single SD pair and the entire network under the general failure model. To guarantee the connectivity of a single SD pair under the geographical failure model, we develop an algorithm to decompose the problem to multiple subproblems, each of which determines the optimal shielding for links in a geographical region. The MILP that guarantees the connectivity of the entire network has partial separable structure, for which Benders decomposition can be applied to significantly reduce the running time. A slightly modified Benders decomposition reduces the running time further by more than 50%. In addition, simulated annealing is used to obtain near-optimal solutions with much shorter running time.

Significantly less shielding cost is required to guarantee partial connectivity, even in the case where only one node is allowed to be disconnected from the others. Moreover, we observe larger reduction in shielding cost if the size of a failure region is small. The algorithms can be easily modified to solve the problem which guarantees the connectivity of a selected set of SD pairs in a network. For example, in the MILP, the flow constraints can be imposed only for the selected set of SD pairs. The methodologies in this paper can be used to construct new and upgrade existing networks to guarantee the connectivity after any single geographical or general failure.

REFERENCES

[1] J. S. Foster *et al.*, "Report of the commission to assess the threat to the united states from electromagnetic pulse (EMP) attack, critical national infrastructures," Apr. 2008.
[2] 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, Dec. 2011.
[3] P. Agarwal, A. Efrat, S. 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, pp. 1525–1538, Oct 2013.
[4] S. Borne, E. Gourdin, B. Liau, and A. Mahjoub, "Design of survivable IP-over-optical networks," *Annals of Operations Research*, vol. 146, no. 1, pp. 41–73, 2006.
[5] H. Kerivin and A. R. Mahjoub, "Design of survivable networks: A survey," *Networks*, vol. 46, no. 1, pp. 1–21, Aug. 2005.
[6] G. Maier, A. Pattavina, S. D. Patre, and M. Martinelli, "Optical network survivability: Protection techniques in the WDM layer," in *Photonic Networks Communications*, 2002, pp. 251–269.
[7] M. Medard, D. Marquis, R. Barry, and S. Finn, "Security issues in all-optical networks," *IEEE Network*, vol. 11, no. 3, pp. 42–48, May 1997.
[8] Booz Allen Hamilton Inc, "Electromagnetic pulse survivability of t-elecommunication assets," Feb 1987.
[9] B. Overton *et al.*, "Microbend-resistant optical fiber," Mar. 27 2012, US Patent 8,145,027.
[10] O. Tatat, "Optical fiber telecommunications cable," Jan. 12 2010, US Patent 7,646,954.
[11] S. Gibbs, "Google reinforces undersea cables after shark bites," The Guardian, Aug. 14 2014. [Online]. Available: http://gu.com/p/4vm78/stw
[12] M. Dziubinski and S. Goyal, "Network design and defence," *Games and Economic Behavior*, vol. 79, no. 0, pp. 30 – 43, 2013.
[13] C. Liu, Y. Fan, and F. Ordez, "A two-stage stochastic programming model for transportation network protection," *Computers and Operations Research*, vol. 36, no. 5, pp. 1582 – 1590, 2009.
[14] L. V. Snyder, M. P. Scaparra, M. S. Daskin, and R. L. Church, "Planning for disruptions in supply chain networks," in *Tutorials in Operations Research. INFORMS*, 2006.
[15] W. H. Cunningham, "Optimal attack and reinforcement of a network," *J. ACM*, vol. 32, no. 3, pp. 549–561, Jul. 1985.
[16] R. L. Church, M. P. Scaparra, and R. S. Middleton, "Identifying critical infrastructure: The median and covering facility interdiction problems," *Annals of the Association of American Geographers*, vol. 94, no. 3, pp. pp. 491–502.
[17] G. Brown, M. Carlyle, J. Salmerón, and K. Wood, "Defending critical infrastructure," *Interfaces*, vol. 36, no. 6, pp. 530–544, Nov. 2006.
[18] P. Gray, "Exact solution of the fixed-charge transportation problem," *Operations Research*, vol. 19, no. 6, pp. pp. 1529–1538, 1971.
[19] M. Pióro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*. Morgan Kaufmann Publishers Inc., 2004.
[20] D. Bertsimas and J. Tsitsiklis, *Introduction to Linear Optimization*, ser. Athena Scientific series in optimization and neural computation. Athena Scientific, 1997.
[21] C. Cao, M. Zukerman, W. Wu, J. H. Manton, and B. Moran, "Survivable topology design of submarine networks," *IEEE/OSA J. Lightwave Technol.*, vol. 31, no. 5, pp. 715–730, March 2013.
[22] D. Papadimitriou *et al.*, "Inference of shared risk link groups," Working Draft, IETF Secretariat, Internet-Draft draft-many-inference-srlg-02.txt, Nov. 2001.
[23] XO Communications, "Network map." [Online]. Available: http://www.xo.com/about/network/Pages/maps.aspx
[24] S. Neumayer, A. Efrat, and E. Modiano, "Geographic max-flow and min-cut under a circular disk failure model," in *INFOCOM*, 2012, pp. 2736–2740.
[25] J. Plesnik, "Equivalence between the minimum covering problem and the maximum matching problem," *Discrete Mathematics*, vol. 49, no. 3, pp. 315 – 317, 1984.
[26] J. Zhang, "Enhancing network robustness via shielding," Master's thesis, MIT, 2014.
[27] D. Bertsimas and J. Tsitsiklis, "Simulated annealing," *Statistical science*, vol. 8, no. 1, pp. 10–15, 1993.
[28] B. Hajek, "Cooling schedules for optimal annealing," *Mathematics of Operations Research*, vol. 13, pp. 311–329, 1988.