A Resource Allocation Model for Improving the Resilience of Critical Transportation Systems

July 1, 2014

Jussi Kangaspunta[1] and Ahti Salo
Systems Analysis Laboratory
Department of Mathematics and Systems Analysis
Aalto University School of Science
P.O.Box 11100, 00076 Aalto, Finland
email: firstname.lastname@aalto.fi

Abstract

The functioning of society and economy depends on critical infrastructures which are vulnerable to disruptions: energy distribution networks and water supply systems, for example, can be exposed to power outages and supply chain disturbances. Analyses of these vulnerabilities are therefore needed to assess and secure the performance of critical infrastructures. Towards this end, we model critical transportation systems as networks which consist of nodes and edges and whose performance is measured by the extent to which the transportation objectives are achieved even if part of the network may be disrupted. Specifically, we develop decision analytic methods to determine (i) how different combinations of disruptions would impact the performance of the network and (ii) which combinations of risk management actions are cost-efficient in maintaining the performance of the network at as high a level as possible with regard to multiple evaluation criteria.

Keywords: Cost-efficiency analysis, critical infrastructure, portfolio analysis, risk analysis, transportation systems.

[1]Corresponding author

1 Introduction

Critical infrastructures consist of all assets, comprising both systems and processes, that are needed to produce and distribute vital goods and services to people. These assets are crucial for the health, wealth, and security of society: transportation systems, for example are needed to deliver goods for people. The performance of these systems may be compromised by adversities such as natural disasters and terrorist attacks, which makes it necessary to understand how vulnerable they are (see e.g. Burns and Slovic 2012, Kleinmuntz and Willis 2009, Brown et al. 2006).

In this paper, we develop methods for assessing the vulnerability of critical transportation systems and for improving the resilience of these networks through cost-efficient risk management actions. Technically, we model transportation systems as networks consisting of nodes that are connected by edges. For example, in the modeling of railroads, railway stations would correspond to nodes while tracks between stations could be treated as edges. In our methods, the performance of the network is measured by the extent to which relevant transportation objectives are achieved. For instance, one objective could be that of minimizing the time needed to transport medical supplies to patients, while another one could be that of ensuring the existence of at least one functioning route between two cities.

Transportation networks, like other critical infrastructures, are vulnerable to disruptions caused by natural events such as extreme weather conditions or intentional attacks such as sabotage. Due to such disruptions, the performance of the network may deteriorate so that some transportation objectives are no longer met. For instance, traveling times may increase if trains must be rerouted due to damaged railway stations.

From the viewpoint of risk analysis, not all parts of the network are equally important, because disruptions of some network nodes will cause a greater decrease in network performance than disruptions of other nodes. Furthermore, because the impacts depend on what the status of other nodes in the network

is, it may be necessary to evaluate the impacts of *combinations* (or portfolios) of disruptions. For instance, if a single node is disrupted, then the performance of the network may not change much; but the performance may decrease dramatically if two nodes are disrupted at the same time. It is also pertinent to account for disruption probabilities, given that some disruptions are more probable than others.

Specifically, our methods help identify which (i) individual network nodes and (ii) combinations of network nodes are most important in the sense that disruptions of these nodes would impact the performance of the network most; moreover, they help identify which (iii) risk management actions are cost-efficient in ensuring that the performance of the network stays at as high a level as possible. Examples of such risk management actions include the fortification of selected network nodes against disruptions or the introduction of new edges between nodes. Technically, we aggregate disruption probabilities so that it becomes possible to associate with any given level of network performance the probability with which this performance level will be reached. Thus, for example, it is possible to ascertain whether or not the probability of meeting the targeted level of network performance is 95% or higher.

Because it is usually possible to implement many risk management actions jointly, we analyze combinations or *portfolios* of such actions. The results of such analyses support the allocation of resources to those risk management actions which are cost-efficient in ensuring that the network performs as well as possible relative to the costs that arise from the implementation of risk management actions (c.f. Kangaspunta et al. 2012). For instance, if the budget for maintaining the road network were to be reduced, then one would like to spend the remaining budget so that the reduction in the performance level would remain as small as possible.

The rest of this paper is organized as follows. Section 2 reviews earlier approaches to evaluating the vulnerability of the critical infrastructures. Section 3 presents our methods for evaluating the performance of transportation networks and for guiding the allocation of resources to risk management actions. Section

4 gives an illustrative example. Section 5 discusses our results and identifies directions for future work. Section 6 concludes.

2 Earlier Methods to Evaluate Networks

Much of the research on the resilience of network infrastructures has been carried out after the turn of the millennium. Albert et al. (2000) discuss the tolerance of real-world networks such as the World-Wide Web and social networks against errors due to random failures and intentional attacks. Even if the nodes and edges of these networks regularly fail, they tend to exhibit a high degree of robustness in that they maintain a fair degree of functionality despite high failure rates. Still, Albert et al. (2000) note that some networks are vulnerable to attacks directed at few vital network components.

Brown et al. (2006) present methods which help identify vulnerabilities in different sectors of critical infrastructures and assist in planning defensive measures against an intelligent attacker. They stress that infrastructures which resist failures at single points may not survive intentional attacks. They note that there is plenty of publicly available information for planning of disastrous attacks on infrastructures. Focusing on the vulnerability of electric power networks, Holmgren (2006) finds that complex systems can be represented as networks, and that interdependencies between different infrastructures such as power and communication systems are of particular interest.

Latora and Marchiori (2005) present a method to find the critical components of an infrastructure network represented by nodes and links. They also analyze how improvements such as the introduction of additional links between nodes increases the performance of the network. As in Latora and Marchiori (2001) the performance is defined as how efficiently information can be exchanged over the network.

Israeli and Wood (2002) define a shortest-path network-interdiction problem in which an attacker seeks to maximize the shortest path length between two given nodes of the network. In this problem, the attacker uses limited re-

sources to interdict paths between these two nodes for instance by destroying paths or by increasing their lengths by reducing the capacities of these paths. Cappanera and Scaparra (2011) consider the problem of allocating resources to protect shortest-path networks with the aim of maximizing the resilience of these networks against disruptions. They present methods to identify which components need to be protected to minimize the length of the shortest path after a worst-case disruption.

The methods we develop in this paper have novel features for evaluating the importance of network nodes and for allocating resources to maintain the performance of transportation networks:

1. The performance of the network is measured holistically by the extent to which the transportation objectives of the network are achieved.

2. It is possible to account for multiple objectives which, for example, can represent the level of performance for transporting different commodities or for servicing different stakeholders.

3. Weight information about the relative importance of the different transportation objectives need not be complete. Thus, for example, results can be produced even on the basis of an ordinal ranking information about the relative importance of objectives. This is useful, because the elicitation of complete weight information can be difficult and time-consuming.

4. Cost-efficiency analyses help identify which portfolios of risk management actions outperform others at different levels of total cost. These kinds of analyses can be performed in order to define the appropriate budget level as well. For instance, if a small increase in the budget for risk management actions helps ensure the performance of the network, then increasing the budget may be justified. Conversely, budget reductions could be warranted if the same performance level can be achieved at a smaller budget.

3 Evaluating Framework for Transportation Networks

Let $G(V; E)$ denote a network consisting of a set of nodes $V = \{1, \ldots, m\}$ and a set of undirected edges $E \subseteq \{(i, j) \mid i, j \in V\}$ between the nodes. If a node is disrupted, then all the edges connecting it to other nodes are removed from the network. Thus, if there is a disruption at node $k \in V$, then the disrupted network is $G(V^d; E^d)$, where $V^d = V \setminus \{k\}$ and $E^d = \{(i, j) \in E \mid i, j \in V^d\} \subseteq E$.

The state of the network is a vector $x = [x_1, \ldots, x_m] \in \mathcal{X} = \{0, 1\}^m$ such that $x_k = 1$ if node $k \in V$ is disrupted and $x_k = 0$ if it functions. The total number of possible states of network nodes is 2^m. Figure 1 shows an example of disruptions in a network consisting of nine nodes and thirteen edges.

Insert Figure 1 around here.

Edges, too, can be vulnerable to disruptions. Edge disruptions can be modeled by replacing each edge $(i, j) \in E$ with an additional node $k \notin V$ and by introducing two edges (i, k) and (k, j). The augmented network then becomes $G(V^a, E^a)$, where $V^a = \{1, \ldots, m + |E|\}, E^a \subseteq \{(i, k) \mid i \in V, k \in V^a \setminus V\} \cup \{(k, j) \mid j \in V, k \in V^a \setminus V\}$ and $|E|$ is the number of edges in the original network. Thus, without losing generality, we consider node disruptions only.

3.1 Assessing the Impacts of Network Disruptions

We consider transportation networks whose performance is measured by the extent to which the transportation objectives are attained. For example, the objective can be that of maximizing the number of connections between suppliers and customers; minimizing the travel time or cost between factories and retailers; or maximizing the number of delivered shipments. The objective can also be based on topological measures such as the average of the shortest dis-

tances between network nodes (see e.g. Latora and Marchiori 2001).

Disruptions decrease the performance of the network. For example, routing adjustments required by a node disruption may cause delays so that the objectives concerning delivery times are not attained. Disruptions at different nodes do not impact the performance of the network similarly, and the impacts caused by a node disruption often depend on what other nodes are disrupted. The total impact of disruptions is therefore a function of the *combination* of nodes that are disrupted, and consequently the performance of the transportation network must be considered in view of combinations of possible disruptions.

Technically, we use a value function v to map the network state $x \in \mathcal{X}$ onto a performance scale so that

$$v = v(x) : \mathcal{X} \to \mathbb{R}. \tag{1}$$

For instance, the performance of a network can be measured using the measure by Latora and Marchiori (2001) so that

$$v(x) = \frac{1}{m(m-1)} \sum_{i \neq j} \frac{1}{d_{i,j}(x)} \in [0, 1], \tag{2}$$

where m is the number of nodes and $d_{i,j}(x)$ is the length of the shortest path between nodes i and j as a function of $x \in \mathcal{X}$. If there is no path between nodes i and j, then $1/d_{i,j}(x) = 0$. The performance (2) is zero if there is no path between any pair of nodes. Conversely, it is one if the average length of shortest paths between all pairs nodes is one.

Often, it is necessary to account for multiple objectives, such as that of maximizing the number of connections *and* that of minimizing transportation costs. In our modeling approach, we assume that there are n objectives which are measured using value functions such that $v_j(x)$ is the value with regard to objective $j = 1, \ldots, n$ when the network is in state $x \in \mathcal{X}$. Without losing generality, these criterion specific values can be normalized onto the range $[0, 1]$ so that $v_j(x) = 0$ if all nodes are disrupted (i.e., $x_k = 1$ for all $k = 1, \ldots, m$) and $v_j(x) = 1$ if there are no disruptions (i.e., $x_k = 0$ for all $k = 1, \ldots, m$).

The different objectives are not necessarily equally important. To gauge the relative importance of these objectives, we employ criterion weights which reflect the increase in the overall value when the criterion specific performance changes from its minimum level to the maximum level. The relative weight of criterion $j = 1, \ldots, n$ is denoted $w_j \in [0, 1]$. Following the usual convention, we assume that these weights are normalized so that they add up to one.

If the objectives are mutually preferentially independent and some technical assumptions hold (see e.g. Dyer and Sarin 1979), the value of network performance associated with the network state $x \in \mathcal{X}$ can be represented by an additive value function

$$v(x, w) = \sum_{j=1}^{n} w_j v_j(x) \in [0, 1].$$

In this additive representation, the additional value brought by an incremental improvement in the performance level on one criterion does not depend on what performance the network provides on the other criteria.

Recognizing that it can be difficult to elicit complete information about criterion weights, we assume that weight information is characterized by a feasible set $\mathcal{S}_w$ rather than exact point estimates (Salo and Hämäläinen 1992). This feasible weight set is a subset of all possible weights

$$\mathcal{S}_w = \{w \in \mathbb{R}^n \mid A_w w \leq B_w\} \subseteq \left\{ w \in \mathbb{R}^n \mid w_j \geq 0 \ \forall j, \sum_{j=1}^{n} w_j = 1 \right\} = \mathcal{S}_w^0, \quad (3)$$

where the constraint matrices $A_w \in \mathbb{R}^{a \times n}$ and $B_w \in \mathbb{R}^a$ contain the coefficients that are implied by preference statements concerning the relative importance of objectives. For instance, if criterion 1 is judged to be at least as important but no more than twice as important than criterion 2, then the constraints $w_1 \geq w_2$ and $w_1 \leq 2w_2$ would apply.

For purposes of illustration, consider the network in Figure 2 which has three suppliers (S_1, S_2, and S_3), three customers (C_1, C_2, and C_3), and six intermediate nodes. A supplier is connected to a customer if there is a route from the supplier to the customer through edges and nodes that are operational.

The objective is to maximize the number of connections between suppliers and customers. The relative importance of a connection between supplier S_i and customer C_j is denoted by $w_{i,j}$. If all connections between suppliers and customers are equally important, then $w_{i,j} = 1/9$ for all $i, j = 1, 2, 3$. The resulting value function is therefore the following function of the network state $x \in \mathcal{X}$

$$v(x, w) = \sum_{i=1}^{3} \sum_{j=1}^{3} w_{i,j} g_{i,j}(x) \in [0, 1], \tag{4}$$

where $g_{i,j}(x) = 1$ if there is a connection between supplier $i = 1, 2, 3$ and customer $j = 1, 2, 3$ and $g_{i,j}(x) = 0$ if there is no such connection.

Insert Figure 2 around here.

We next evaluate which combinations of disruptions are most critical to the performance of this network. For purposes of illustration, we assume that only the intermediate nodes $1, \ldots, 6$ are vulnerable to disruptions. The number of connections between suppliers and customers is shown in Figure 3 as a function of the number of disrupted network nodes.

Insert Figure 3 around here.

The following minimization problem gives those combinations of node disruptions that decrease the performance of the network most for a given number $n_d \in \mathbb{Z}$ of disruptions and $w \in \mathcal{S}_w$

$$\begin{array}{lrcl} \min\limits_{x} & v(x, w) & & \\ \text{s.t.} & \sum_{i=1}^{m} x_i & = & n_d \\ & x & \in & \mathcal{X}. \end{array} \tag{5}$$

Solutions for the problem (5) are marked with squares in Figure 3.

For example, if node 5 in Figure 2 is disrupted, then seven connections remain. Moreover, if one of the two combinations {1,5} or {4,5} are disrupted, then only three connections remain. This can be contrasted with the disruption

of four nodes {1,3,4,6} which leaves all connections intact. This highlights that the nodes are not equally important for the performance of the network.

3.2 Risk Profiles for Network Performance

To evaluate the risk profile of the network, assume that the weight vector w is selected from the set of feasible weights $\mathcal{S}_w$. Here, we provide results on the risk profiles for this selected weight vector, in the recognition that results for the entire set of feasible weights $\mathcal{S}_w$ can be obtained by examining all the extreme points of the feasible weight set.

If node disruptions occur independently, uncertainties about node disruptions can be represented by the vector $p = [p_1, \ldots, p_m] \in \mathcal{P} = [0,1]^m$ where p_k is the probability of node disruption at node k. Thus, the probability of network state $x \in \mathcal{X}$ is

$$p(x) = \prod_{k=1}^{m} [x_k p_k + (1 - x_k)(1 - p_k)] \in [0,1]. \tag{6}$$

By construction, summing the probabilities of all network states equals one.

The probability (density) function of network performance for $w \in \mathcal{S}_w$ and performance level $v \in [0,1]$ is

$$f(v,w) = \sum_{\{x \in \mathcal{X} | v(x,w) = v\}} p(x) \in [0,1]. \tag{7}$$

The corresponding cumulative probability function gives the probability that network performance is less than or equal to a given performance level v. For the given weight vector $w \in \mathcal{S}_w$ and $v \in [0,1]$, this function is

$$F(v,w) = \sum_{\{x \in \mathcal{X} | v(x,w) \leq v\}} p(x) \in [0,1]. \tag{8}$$

Various risk constraints can be introduced by employing these probability distributions. For example, if network performance is measured using the value function (4), then requiring that the probability of having three or fewer connections out of nine must not exceed 5% corresponds to the inequality $F(3/9, w) \leq 0.05$.

For a given weight vector $w \in \mathcal{S}_w$, the expected network performance is

$$\mathbb{E}_p\left[v(x,w)\right] = \sum_{x \in \mathcal{X}} p(x)v(x,w) \in [0,1]. \tag{9}$$

The Decision Maker (DM) may be interested in minimizing risks by modifying the structure of the network or by carrying out maintenance actions to fortify network nodes. For such purposes, risk measures such as Value-at-Risk (VaR) can be used to quantify risks and to evaluate risk management actions. VaR is defined for a given confidence level $\alpha \in (0,1]$ and $w \in S_w$ as the greatest value in the worst α-quantile of the network performance

$$\mathrm{VaR}_p^{\alpha}(v,w) = \sup\{v \in [0,1] \mid F(v,w) \leq \alpha\} \in [0,1]. \tag{10}$$

Although VaR is a widely used, it is not coherent unlike Conditional Value-at-Risk (CVaR) (Rockafellar and Uryasev 2000, Sarykalin et al. 2008). For a given $\alpha \in (0,1]$ and $w \in S_w$, CVaR is defined as the expected value in the worst α-quantile of the network performance

$$\mathrm{CVaR}_p^{\alpha}(v,w) = \mathbb{E}\left[v(x,w) \mid v \leq \mathrm{VaR}_p^{\alpha}(v,w)\right] \in [0,1]. \tag{11}$$

To illustrate how risk measures can be employed, we reconsider the example in Figure 2 and assume that the probability of disruption is 20% at each node (i.e. $p_k = 0.20$ for $k = 1,\ldots,6$). The probability density function (7) and the cumulative probability function (8) are shown in Figure 4(a) and Figure 4(b), respectively.

The expected number of connections, for example, is 7.97 and the probability of having exactly seven connections between suppliers and customers is about 10%. There is a 16% probability that there are fewer than seven connections. Now, if it were to be required that the probability of having at least seven connections must be higher than 90%, this network would not fulfil such a requirement and it would be necessary to identify nodes at which risk management actions could be implemented.

Insert Figure 4 around here.

3.3 Importance Measures for Network Nodes

We evaluate the relative importance of network nodes in view of questions such as: How much does the expected performance of the network improve if a node is fortified by eliminating the possibility of disruption? How much does the expected performance of the network decline if a node disruption does occur? Will the performance of the network decline below some given performance level if a selected set of nodes has been fortified or disrupted?

We start by considering the *status quo* network in which node k disrupts with probability p_k. Now, let $p^{k=0}$ denote the probability vector of node disruptions which is identical to p expect in that node k has been fortified, meaning that the possibility of a disruption at node k has been eliminated so that $p_k^{k=0} = 0$ and $p_l^{k=0} = p_l,\ l \neq k$. Similarly, let $p^{k=1}$ be the probability vector which, again, is the same as p except in that there is a sure disruption at node k, meaning that $p_k^{k=1} = 1$ and $p_l^{k=1} = p_l,\ l \neq k$.

From the viewpoint of directing actions that help *maintain* the expected performance of the network, it is of interest to identify at which nodes disruptions would cause the greatest decline in expected network performance. According to this measure, called *disruption impact*, node k would be viewed as more important than node l if $\mathbb{E}_{p^{k=1}}[v(x,w)] < \mathbb{E}_{p^{l=1}}[v(x,w)]$, given that the disruption at node k would cause a greater decline in the expected performance of the network than that at node l (here, the probability distribution with regard to which the expectation is taken is shown in the subscript).

In order to *improve* expected network performance, one is interested in determining those nodes the fortification of which would result in the largest improvement in the expected network performance. Specifically, this *fortification impact* would be higher for node k than for node l if $\mathbb{E}_{p^{k=0}}[v(x,w)] > \mathbb{E}_{p^{l=0}}[v(x,w)]$.

Insights into the relative importance of different network nodes can be generated by comparing (i) the expected performance of networks in which selected nodes are disrupted or fortified with (ii) the expected performance of the status quo network. That is, the disruption impact $I^1(k) = \mathbb{E}_p[v(x,w)] - \mathbb{E}_{p^{k=1}}[v(x,w)]$

of node k shows how much the the expected network performance declines if there is a disruption at node k. Similarly, the fortification impact $I^0(k) = \mathbb{E}_{p^{k=0}}[v(x,w)] - \mathbb{E}_p[v(x,w)]$ indicates how much the expected network performance can be improved by fortifying node k.

For example, consider the transportation network in Figure 5 in which there are twelve edges and nine nodes with disruption probabilities

$$p_1 = p_2 = p_3 = 0.05, p_4 = p_5 = p_6 = 0.10, \text{ and } p_7 = p_8 = p_9 = 0.20.$$

We assume that the performance of this network is measured by employing the value function in (2). The disruption and fortification impacts for this network are shown in Figure 6. For example, the expected performance decreases most if there is a disruption at node 4. Conversely, the expected performance increases most if node 8 is fortified to remain operational.

Insert Figure 6 around here.

Disruption impact and fortification impact are both results of sensitivity analyses that can be extended to examine networks in which several nodes are disrupted (i.e. $p_k = 1$ for some $k \in V_1$) or fortified (i.e. $p_k = 0$ for some $k \in V_0$). That is, if the nodes in the set $V_0 \subseteq V$ remain operational and those in the set $V_1 \subseteq V$ are disrupted (we assume that the sets V_0 are V_1 disjoint), then the probabilities of node disruptions in this *modified* network is $\tilde{p} = [\tilde{p}_1, \ldots, \tilde{p}_m] \in [0,1]^m$, where

$$\tilde{p}_k = \begin{cases} 0, & \text{if } k \in V_0, \\ 1, & \text{if } k \in V_1, \\ p_k, & \text{otherwise.} \end{cases}$$

If $V = V_0 \cup V_1$, all uncertainties about node disruptions will be eliminated, because the network state will be $x = [x_1, \ldots, x_m]$ where $x_k = 0$ for all $k \in V_0$ and $x_k = 1$ for all $k \in V_1$.

The disruption impact and fortification impact have here been defined in terms of changes in the expected performance of the network. Yet, one could

equally well examine what impacts possible disruptions and fortifications would have on different risk measures of network performance, most notably on VaR or CVaR. Also, when using the additive value function (1) with incomplete information about attribute weights, then these impacts could be computed for all the extreme points of the feasible weight set S_w in order to communicate to the DM the ranges within which the resulting impacts would reside for different choices of feasible weights (see e.g. Toppila and Salo 2013).

3.4 Securing the Performance of Networks

The performance of the network can be improved through risk management actions. These actions can, for example, decrease the probability of node disruptions or they may entail the introduction of additional nodes and edges. Figure 7 illustrates the decision to implement an action to protect a network node. Without this action, the initial probability of disruption at node k is p_k. If the action is implemented, this probability becomes smaller, denoted by $p'_k < p_k$.

Insert Figure 7 around here.

The introduction of additional nodes and edges may establish new routes in the network so that possible disruptions have less impact. For example, to illustrate how the addition of a node can improve the resilience of the network, assume that the objective in the network of Figure 8 is to secure the connection between nodes 1 and 3. Initially, there is only one connecting node 2 between nodes 1 and 3. However, if node 4 and two additional edges are added, nodes 1 and 3 remain connected even if the node 2 is disrupted.

The introduction of additional nodes and edges can be viewed as risk management actions that decrease the disruption probabilities of the nodes. For instance, if the disruption probabilities of nodes 2 and 4 are p_2 and p_4, respectively, then there is a connection between nodes 1 and 3 with probability $1 - p_2 p_4$. Therefore, this additional node can be modeled by updating the

disruption probability of node 2 to $p_2' = p_2 p_4 < p_2$.

Insert Figure 8 around here.

Usually, there are many actions that the DM can take to improve the resilience of the network, and the probability with which different levels of network performance can be attained depends on which *portfolio* of actions is implemented. We assume that there are r possible risk management actions and $q_i = 1$ if action i is implemented and $q_i = 0$ if it is not, $i = 1, \ldots, r$. The action portfolio is given by a vector $q = [q_1, \ldots, q_r] \in \mathcal{Q} = \{0, 1\}^r$. The total cost of portfolio q is $c(q) \in \mathbb{R}_+$, and the available budget is b.

A portfolio of risk management actions is *feasible* if its cost is within budget (i.e., $c(q) \leq b$) and it satisfies possible logical constraints. For instance, if actions 1 and 2 are mutually exclusive, the constraint $q_1 + q_2 \leq 1$ must hold. The set of feasible portfolios is denoted by $\mathcal{Q}_F \subseteq \mathcal{Q}$. The objective is to determine which feasible portfolio $q \in \mathcal{Q}_F$ satisfies the relevant risk constraints and is best in terms of the selected objective, such as the maximization of expected performance. Once the action portfolio $q \in \mathcal{Q}_F$ has been implemented, the probability of the network state $x \in \mathcal{X}$ is

$$p(x \mid q) = \prod_{k=1}^{m} [x_k p_k(q) + (1 - x_k)(1 - p_k(q))] \in [0, 1], \tag{12}$$

where $p_k(q)$ is the probability of disruption at node k when portfolio q has been implemented. Likewise, the probability density function and the cumulative probability distribution of network performance are denoted $f(v, w \mid q)$ and $F(v, w \mid q)$, respectively.

This formulation is generic in that the disruption probability of each node can depend on the entire portfolio of actions that are being considered. Thus, if there are two actions which both impact the probability of disruption at a given node and can be taken jointly, it would be necessary to estimate what the probability of node disruption would be if both actions are implemented or if only one of the actions is taken. More generally, if there are K different actions

such that all combinations therefore can impact the probability of disruption at node k, then it would be necessary to elicit 2^K probability parameters for node k.

In more typical case, however, different risk management actions pertain to different nodes. Assuming that there is only one possible action at each node (i.e. $r = m$) and that the action at node k decreases the disruption probability to $p'_k < p_k$, the probability of disruption at node k is

$$p_k(q) = \begin{cases} p_k, & \text{if } q_k = 0, \\ p'_k, & \text{if } q_k = 1. \end{cases} \tag{13}$$

When the DM seeks to maximize the expected network performance, as measured by the value function, the objective is to determine which feasible portfolios outperform others at different cost levels for all feasible weights. The risk management portfolio $q^1 \in \mathcal{Q}_F$ is dominated by portfolio $q^2 \in \mathcal{Q}_F$ if and only if $\mathbb{E}_p\left[v(x,w) \mid q^1\right] \leq \mathbb{E}_p\left[v(x,w) \mid q^2\right]$ for all $w \in \mathcal{S}_w$ and $\mathbb{E}_p\left[v(x,w) \mid q^1\right] < \mathbb{E}_p\left[v(x,w) \mid q^2\right]$ for some $w \in \mathcal{S}_w$.

Definition 1 *Risk management portfolio $q^1 \in \mathcal{Q}_F$ is dominated by portfolio $q^2 \in \mathcal{Q}_F$, denoted by $q^2 \succ q^1$, if and only if $\mathbb{E}_p\left[v(x,w) \mid q^1\right] \leq \mathbb{E}_p\left[v(x,w) \mid q^2\right]$ for all $w \in \mathcal{S}_w$ and (ii) $\mathbb{E}_p\left[v(x,w) \mid q^1\right] < \mathbb{E}_p\left[v(x,w) \mid q^2\right]$ for some $w \in \mathcal{S}_w$.*

Moreover, if $\mathbb{E}_p\left[v(x,w) \mid q^1\right] = \mathbb{E}_p\left[v(x,w) \mid q^2\right]$ for all $w \in \mathcal{S}_w$, then portfolios q^1 and q^2 are equally efficient, which is denoted by $q^1 \sim q^2$. A feasible portfolio is cost-efficient if (i) it is not dominated by any feasible portfolio which costs at most as much and (ii) there is no other equally efficient portfolio that costs less.

Definition 2 *A feasible portfolio $q^1 \in Q_F$ is cost-efficient, denoted by $q^1 \in Q_{CE} \subseteq Q_F$ if and only if $\nexists q^2 \in Q_F$ such that $q^2 \succ_C q^1 \Leftrightarrow$ (i) $q^2 \succ q^1, c(q^2) \leq c(q^1)$ or (ii) $q^2 \sim q^1$ and $c(q^2) < c(q^1)$.*

Although Definitions 1 and 2 are defined in terms of the objective of maximizing the expected network performance, these definitions can be readily extended to account for feasibility constraints that arise from the consideration

of risk measures such as VaR and CVaR, for instance. With incomplete weight information, these constraints could be employed conservatively by requiring that they hold for all extreme points of the feasible weight set.

3.5 Computation of Cost-Efficient Portfolios

The elicitation of preferences about the relative importance of objectives typically leads to linear inequalities which define a polyhedral set of feasible weights $\mathcal{S}_w$. To compare values of network performance $v(x, w)$ with regard to all feasible weights, it is sufficient to examine the extreme points of the feasible weight set $\mathcal{S}_w^{ext}$ (Liesiö et al. 2008). These extreme points of a polyhedral set can be computed using techniques of linear programming (see e.g. Taha 2003).

The following iterative algorithm can be used to determine cost-efficient action portfolios $\mathcal{Q}_{CE}$:

1: Initialize $v(x, w)$ for all $x \in \mathcal{X}$ and $w \in \mathcal{S}_w^{ext}$
2: $\mathcal{Q}^0 \leftarrow \{[0, \ldots, 0]\}$
3: for $k = 1, \ldots, r$
4: $\quad \mathcal{Q}^k \leftarrow \left\{q \in \mathcal{Q}_F \mid q_k = 1 \wedge \exists q' \in \mathcal{Q}^{k-1} : q_l = q'_l, l \neq k\right\}$
5: $\quad \mathcal{Q}^k \leftarrow \mathcal{Q}^k \setminus \left\{q \in \mathcal{Q}^k \mid \exists q' \in \mathcal{Q}^{k-1} : q' \succ q\right\}$
6: $\quad \mathcal{Q}^{k-1} \leftarrow \mathcal{Q}^{k-1} \setminus \left\{q \in \mathcal{Q}^{k-1} \mid \exists q' \in \mathcal{Q}^k : q' \succ q\right\}$
7: $\quad \mathcal{Q}^k \leftarrow \mathcal{Q}^k \cup \mathcal{Q}^{k-1}$
8: end for
9: $\mathcal{Q}_{CE} \leftarrow \mathcal{Q}^r$

In Step 1, the criterion specific values and the overall values for network states are initialized by using the extreme points of the set of feasible weights. This can be time consuming if the values need to be computed through simulation models and therefore, it may be appropriate to consider only a subset of possible states, for instance. In Step 2, the set of portfolios is initialized to contain only an empty portfolio. In Steps 3 to 8, the index k iterates through all risk management actions $1, \ldots, r$. In Step 4, a set $\mathcal{Q}^k$ is constructed using set $\mathcal{Q}^{k-1}$ so that only the values for column k are different. In Step 5 portfolios in

$\mathcal{Q}^k$ are compared to portfolios in $\mathcal{Q}^{k-1}$, and inefficient portfolios are discarded from the set $\mathcal{Q}^k$. These comparisons require the computation of expected network performance or other measure such as CVaR for each $q \in \mathcal{Q}^k$ and $w \in \mathcal{S}_w^{ext}$. In Step 6 inefficient portfolios are discarded from $\mathcal{Q}^{k-1}$ by comparing them to portfolios in $\mathcal{Q}^k$. In Step 7 portfolios in $\mathcal{Q}^{k-1}$ are added to set $\mathcal{Q}^k$. In the final Step 9, the set of cost-efficient portfolios is Q^r.

4 An Illustrative Example

Consider the twelve node transportation network in which nodes S_1, S_2, and S_3 represent suppliers and nodes C_1, C_2, and C_3 represent customers (see Figure 2). The objective is to serve as many customers as possible, subject to the following preference statements about suppliers and customers. First, all connections from supplier S_1 to any one of its customers are at least as important as those between supplier S_2 and its customers; and similarly, connections between supplier S_2 and its customers are at least as important as those between supplier S_3 and its customers. Second, for all suppliers, customer C_1 is at least as important as customer C_2, which in turn is at least as important as customer C_3. These statements lead to the following linear inequalities on feasible weights

$$w_{1,1} \geq w_{1,2} \geq w_{1,3} \geq w_{2,1} \geq w_{2,2} \geq w_{2,3} \geq w_{3,1} \geq w_{3,2} \geq w_{3,3},$$

where $w_{i,j}$ represents the relative importance of the connection between supplier S_i and customer C_j. The probability of disruption is assumed to be $p_k = 0.20$ at all nodes $k = 1, \ldots, 6$. The DM states that she seeks (i) to maximize the expected network performance and (ii) to maximize the probability that there are more than four connections.

For each node, there are two mutually exclusive risk management actions, A and B, of which the cost of action A is one unit and that of B is two units. Thus, the total number of portfolios of actions is 3^6=729, given that at each node it is possible to implement action A, B, or no action at all. Action A decreases the probability of node disruption to 0.10 and action B decreases the probability to

0.05. Cost-efficient portfolios are computed using the algorithm in Section 3.5. The computation took a few seconds on a laptop computer (2GHz/4Gb).

The resulting 42 cost-efficient portfolios are presented in Figure 9 as a function of the total cost of actions. For instance, at the cost level 3, there are three cost-efficient portfolios. These portfolios correspond to action B in node 5 and action A in node 1 or 2; and action A in node 5 and action B in node 1. Furthermore, it can be seen, for example, that for cost levels exceeding 3, it is always action B that is employed to protect node 5.

Insert Figure 9 around here.

The expected network performance of cost-efficient portfolios with regard to cost levels is shown in Figure 10. Moreover, the probabilities of having more than four connections for each of the cost-efficient portfolios are presented in Figure 11.

Insert Figures 10 and 11 around here.

5 Discussion

From the computational perspective, one challenge in the application of this framework is that the number of network states grows quickly as a function of network size. That is, if the network has m nodes which are all vulnerable to disruptions, there are 2^m network states and, in principle, the performance level of each network state would have to be estimated by relying on analytical computations, simulation models or expert judgments (see e.g. Kangaspunta and Salo 2014, Bedford and Cooke 2001).

However, in practice, attention can largely be focused on nodes that are likely to be of greatest concern, in the sense that they have higher disruption probabilities or have a more central position in the network. Also, if the disruption probabilities are small (say, below 10%), then introducing an additional

disruption to an existing portfolio of $m > 1$ disrupted nodes would result in a portfolio of disrupted nodes such that the probability of the augmented portfolio would be by an order of magnitude smaller than that of the initial one. Such observations can be employed to reduce the number of portfolios that would have to be explicitly considered.

In our framework, we have assumed that individual node disruptions do not depend on what other nodes are disrupted. Yet, in some situations there would be interest in modeling dependencies so that the probability of disruption in one node does depend on what other nodes have been disrupted: this would be the case, for example, if the failure of a node increases the load on some other node to the extent that the disruption probability of the latter node increases. Analyses of such interdependencies could be captured through Bayesian analyses (see e.g. Gelman et al. 2009, Langseth and Portinale 2007). One could even admit incomplete information about disruption probabilities: for example, the probabilities of node disruptions could be elicited by asking experts to express statements on verbal scales and by mapping each such statement to an interval of probabilities (see e.g. Toppila and Salo 2013).

The use of binary variables in modeling disruptions means that the nodes are either fully operational or dysfunctional. To model different levels of node performance more comprehensively, one could introduce multi-valued variables in order to capture different gradations of node performance and, for instance, to model the capacity that nodes in a transportation network have during rush-hours or in the presence of minor road accidents (see e.g. Ramirez-Marquez and Rocco S. 2009, Cormican et al. 1998). A potential challenge with the introduction of multi-valued variables is that the overall number of network states would quickly grow. That is, if three states are permitted for each node in a network containing m nodes, the total number of network states would be 3^m as opposed to 2^m in a network with the same number of binary variables.

In attacker-defender situations, there is an interdictor who seeks to determine which combinations of node disruptions would reduce the expected performance of the network most. The objective function of such an interdictor can be mod-

eled by replacing the original value function $v(x, w)$ by $v^{int}(x, w) = 1 - v(x, w)$. The interdictor would seek to determine how he should expend his resources to influence node disruption probabilities so that the network performance deteriorates as much as possible. But then the defender could anticipate that disruptions at these nodes are more probable (because they are more attractive for the interdictor) and therefore prepare by implementing risk management actions at such nodes (see e.g. Powell 2007). Thus, given the strategic nature of such expectations and actions, game theoretic approaches would be needed to capture interactions in these kinds of attacker-defender situations (see e.g. Parnell et al. 2010, Rios Insua et al. 2009, Bell et al. 2008).

One could also extend this framework by building scenarios to characterize different states of the world in which the network is required to operate. For instance, if the node disruptions probabilities depend on the weather, then it could be beneficial to specify scenarios that represent different weather conditions. Then, one could examine risk management actions in order to determine if a given portfolio of risk management actions will offer a sufficiently high level of network performance across all such scenarios; or, assuming that probabilities can be associated with the different scenarios, to determine what the expected level of performance of the network is when taking the expectations over all such scenarios.

6 Conclusion

In this paper, we have developed a framework for assessing the risks of transportation networks consisting of nodes and edges that may be disrupted due to events such as natural hazards, technical failures, or intentional attacks. Starting from estimates about the probabilities of these distributions, we are able to characterize the cumulative probability distribution for the performance level of the network. This resulting distribution synthesizes information about how these disruptions (which may occur at a single node or many nodes simultaneously) affect the performance of the network. It also offers the foundation

for choosing cost-efficient portfolios of risk management actions that best improve the performance of the network, relative to the cost of implementing such actions.

Our framework opens up avenues for further methodological and applied research on the analysis of critical infrastructures and the planning and implementation of risk management actions. For example, although we have focused on a single network, the simultaneous consideration of multiple interlinked networks consisting of, say, energy, transportation, and communication systems calls for further methodological extensions that can be employed to improve the resilience of these interdependent networks. Furthermore, the framework could be extended to situations in which the disruption probabilities are contingent not only on the selected risk management actions but also on the occurrence of disruptions in other parts of the network or the presence of specific environmental conditions depicted by scenarios.

References

R. Albert, H. Jeong, and A.-L. Barabási. Error and attack tolerance of complex networks. *Nature*, 406:378–382, 2000.

T. Bedford and R. Cooke. *Probabilistic Risk Analysis: Foundations and Methods*. Cambridge University Press, New York, 2001.

M.G.H. Bell, U. Kanturska, J.-D. Schmöcker, and A. Fonzone. Attacker-defender models and road network vulnerability. *Philosophical Transactions of the Royal Society*, 366(1872):1893–1906, 2008.

G. Brown, M. Carlyle, J. Salmerón, and K. Wood. Defending critical infrastructure. *Interfaces*, 36(6):530–544, 2006.

W.J. Burns and P. Slovic. Risk perception and behaviors: Antcipating and responding to crises. *Risk Analysis*, 32(4):579–582, 2012.

P. Cappanera and M.P. Scaparra. Optimal allocation of protective resources in shortest-path networks. *Trasportation Science*, 45(1):64–80, 2011.

K.J. Cormican, D.P. Morton, and R.K. Wood. Stochastic network interdiction. *Operations Research*, 46(2):184–197, 1998.

J.S. Dyer and R.K. Sarin. Measurable multiattribute value functions. *Operations Research*, 27(4):810–822, 1979.

A. Gelman, J.B. Carlin, H.S. Stern, and D.B. Rubin. *Bayesian Data Analysis, Second Edition*. Chapman Hall/CRC Texts in Statistical Science, 2009.

Å.J. Holmgren. Using graph models to analyze the vulnerability of electric power networks. *Risk Analysis*, 26(4):955–969, 2006.

E. Israeli and R.K. Wood. Shortest-path network interdiction. *Networks*, 40(2): 97–111, 2002.

J. Kangaspunta and A. Salo. Expert judgments in the cost-effectiveness analysis of resource allocations: A case study in military planning. *OR Spectrum*, 36 (1):161–185, 2014.

J. Kangaspunta, J. Liesiö, and A. Salo. Cost-efficiency analysis of weapon system portfolios. *European Journal of Operational Research*, 223(1):264–275, 2012.

D.N. Kleinmuntz and H.H. Willis. Risk-based allocation of resources to counter terrorism. *CREATE Research Archive, Research Project Summaries, Paper 37*, 2009.

H. Langseth and L. Portinale. Bayesian networks in reliability. *Reliability Engineering and System Safety*, 92(1):92–108, 2007.

V. Latora and M. Marchiori. Efficient behavior of small-world networks. *Physical Review Letters*, 87(19):198701, 2001.

V. Latora and M. Marchiori. Vulnerability and protection of infrasctructure networks. *Physical Review E*, 71:015103, 2005.

J. Liesiö, P. Mild, and A. Salo. Robust portfolio modeling with incomplete cost information and project interdependencies. *European Journal of Operational Research*, 190(3):679–695, 2008.

G.S. Parnell, C.M. Smith, and F.I. Moxley. Intelligent adversary risk analysis: A bioterrorism risk management model. *Risk Analysis*, 30(1):32–48, 2010.

R. Powell. Defending against terrorist attacks with limited resources. *The American Political Science Review*, 101(3):527–541, 2007.

J.E. Ramirez-Marquez and C.M Rocco S. Stochastic network interdiction optimization via capacitated network reliability modeling and probabilistic solution discovery. *Reliability Engineering and System Safety*, 94(5):913–921, 2009.

D. Rios Insua, J. Rios, and D. Banks. Adversarial risk analysis. *Journal of the American Statistical Association*, 104(486):841–854, 2009.

R.T. Rockafellar and S.P. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–42, 2000.

A. Salo and R.P. Hämäläinen. Preference assessment by imprecise ratio statements. *Operations Research*, 40(6):1053–1061, 1992.

S. Sarykalin, G. Serraino, and S. Uryasev. Value-at-risk vs. conditional value-at-risk in risk management and optimization. *Tutorials in Operations Research*, INFORMS 2008:270–294, 2008.

H.A. Taha. *Operations Research: An Introduction*. Pearson Education, Inc, New Jersey, 2003.

A. Toppila and A. Salo. A computational framework for prioritization of events in fault tree analysis under interval-valued probabilities. *IEEE Transactions on Reliability*, 62(3):583–595, 2013.

Tables and Figures

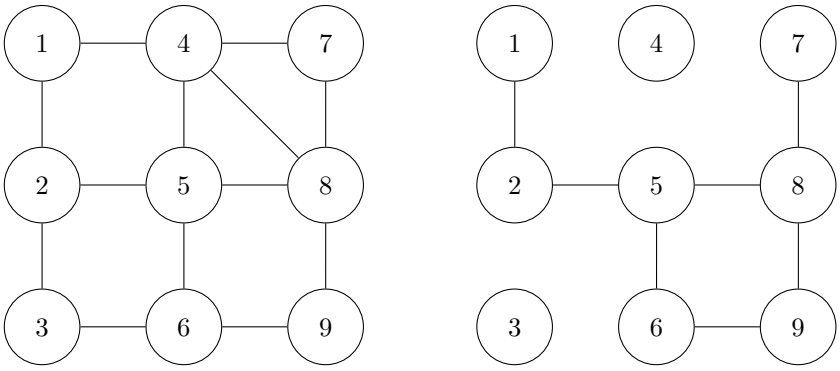

Figure 1: Disrupted network caused by disruptions of nodes 3 and 4.

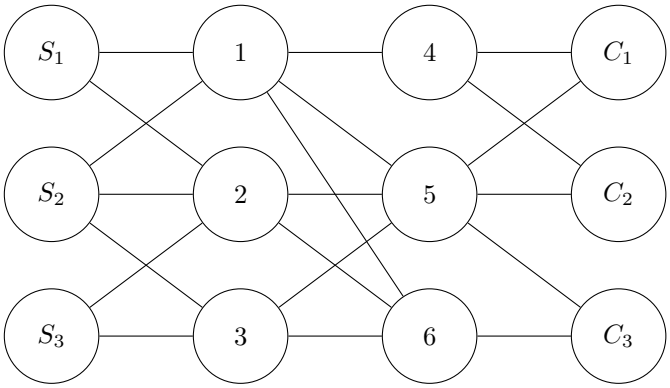

Figure 2: A transportation network of three suppliers ($S_1 - S_3$), three customers ($C_1 - C_3$), and six intermediate nodes (1-6).

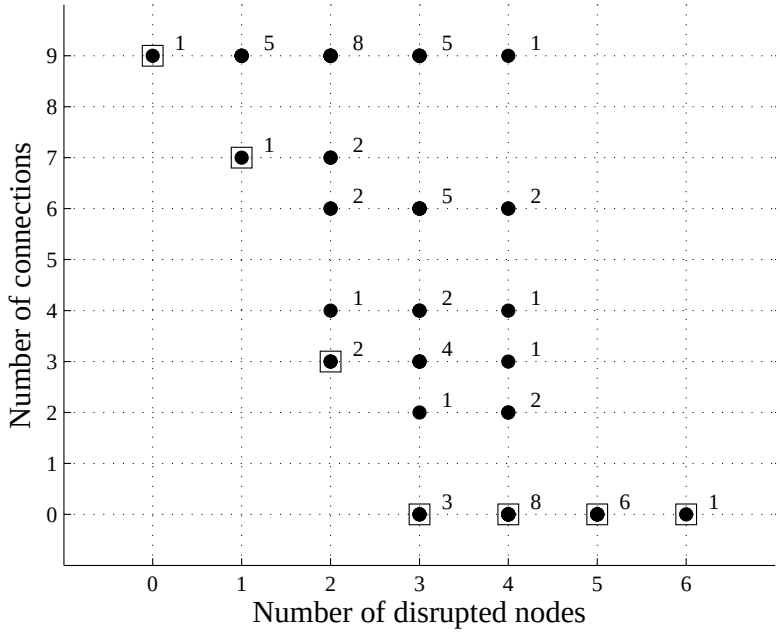

Figure 3: Number of connections between suppliers and customers as a function of the number of disrupted nodes. Numbers on the right side of the dots show how many disruption combinations are at the same point.

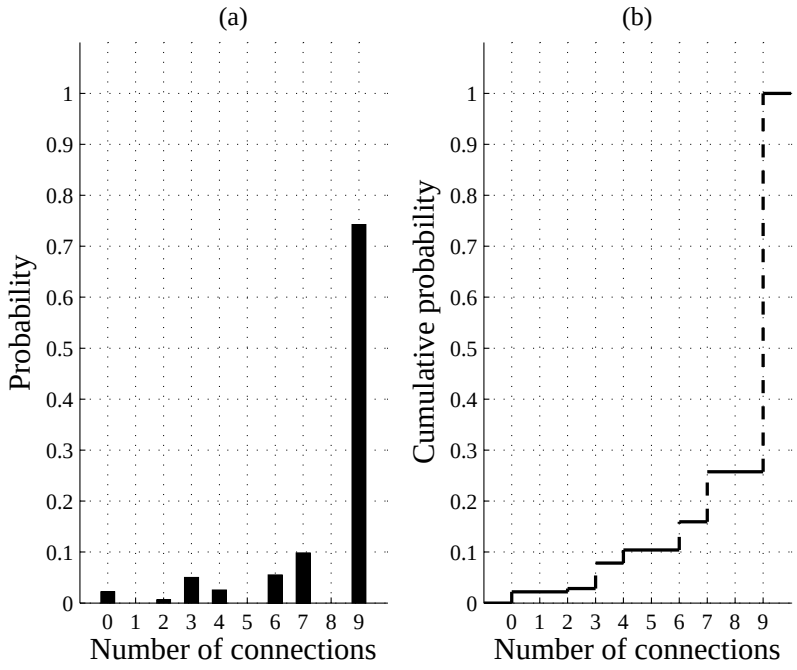

Figure 4: (a) Probability for different number of connections and (b) probability that the number of connections is at most a given level.

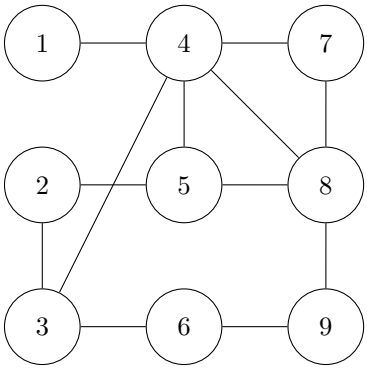

Figure 5: A network of of nine nodes and twelve edges.

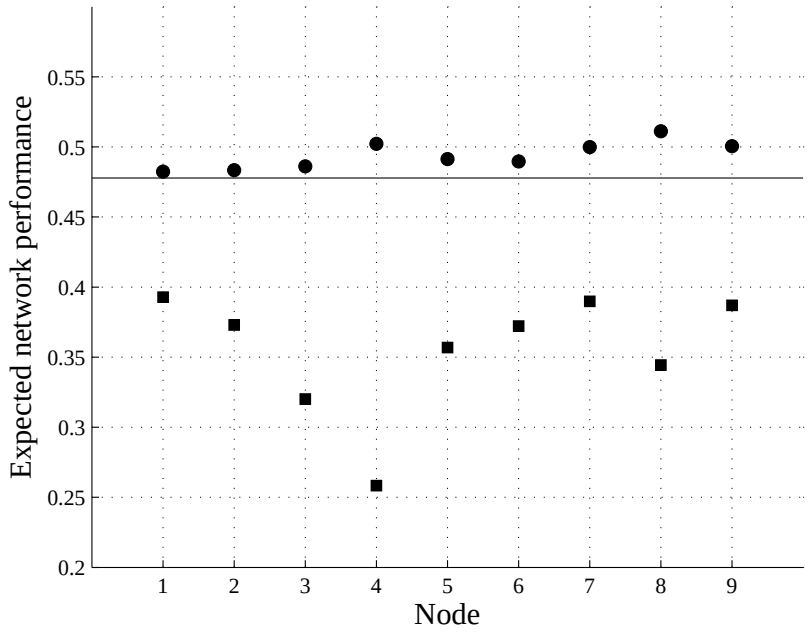

Figure 6: Changes in the expected performance of the network in Figure 5 due to disrupted nodes (squares) and fortified nodes (circles). Horizontal line is the expected performance of the status quo network.

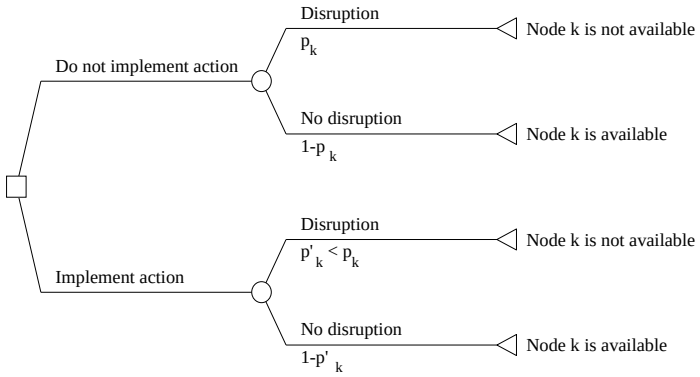

Figure 7: Decision tree for implementing an action to protect network node.

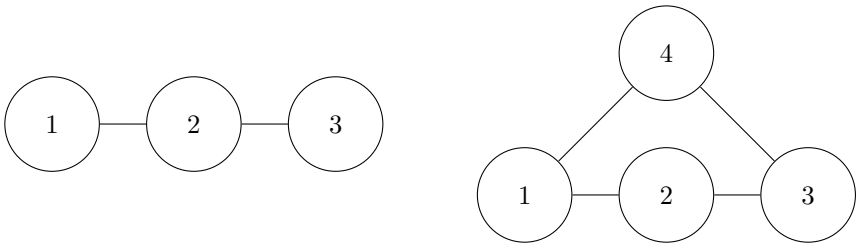

Figure 8: The addition of a node may enhance the resilience of the network.

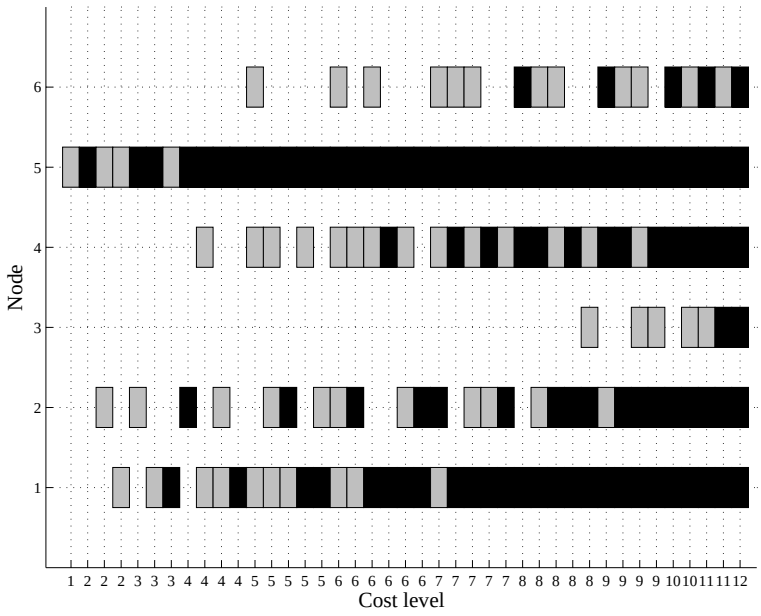

Figure 9: Cost-efficient portfolios of actions for the network in Figure 2. Actions A and B are indicated using grey and black markers, respectively.

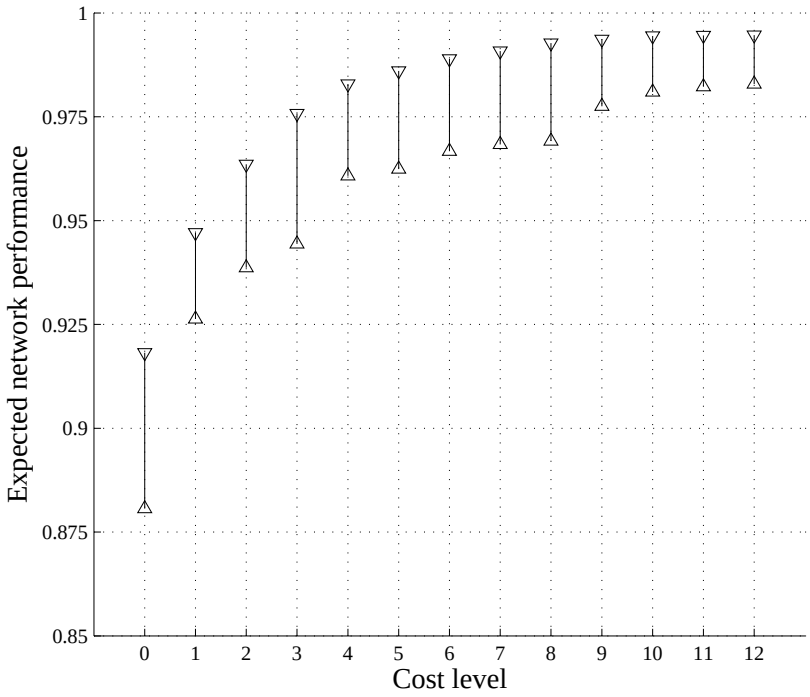

Figure 10: Cost-efficient portfolios with regard to expected network performance.

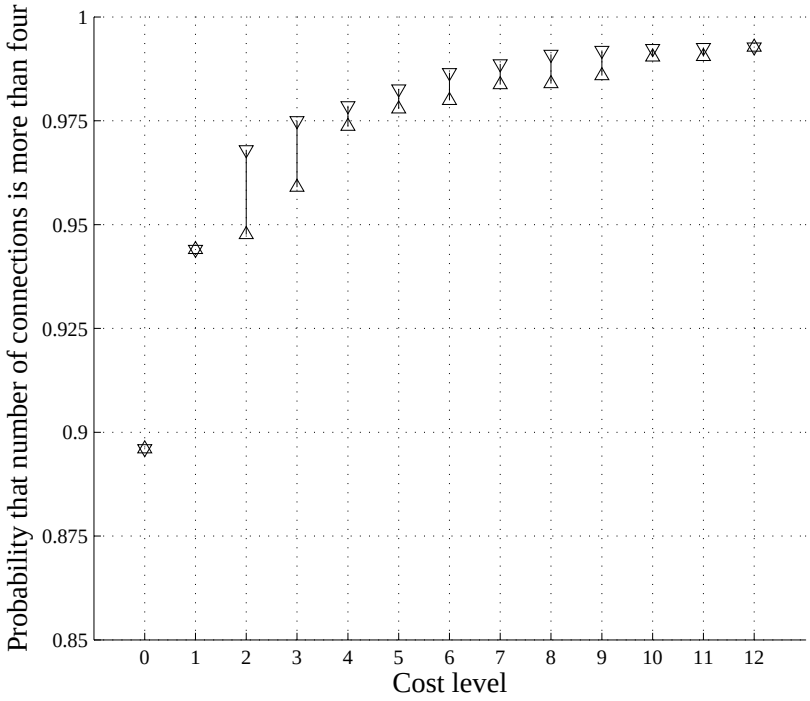

Figure 11: Cost-efficient portfolios with regard to the probability that the number of connections is more than four.