

ABSTRACT

This paper shows that no simple, common-sense rule of thumb can be used to identify a most-vital arc, even in a simple maximum-flow problem. The correct answer requires analysis equivalent in difficulty to completely solving the maximum-flow problem, perhaps repeatedly. This insight generalizes to finding a most-vital component, or set of components, in a system whose operation is described by a more general model. Our paper shows how to evaluate the criticality of sets of components, how to assess the worst-case set of components that might be lost to a given number of simultaneous hostile attacks (or engineering failures, or losses to Mother Nature), and how to allocate limited defensive resources to minimize the maximum damage from a subsequent attack. Collateral insights include the fact that *there is no way to prioritize individual components by criticality*, and that the analysis that determines critical component sets also yields objective assessments of operational system resilience and can provide constructive advice on how to increase it.

INTRODUCTION

When determining how best to protect infrastructure systems from attack, a natural question is, "What components are most critical?" or, equivalently, "which set of components will be most disruptive to the system if lost?" A *critical component* (or set of components) is one whose loss would significantly reduce system function relative to the reduction from losing other components. For example, consider a maximum-flow model, in which a network of capacitated arcs is used to model the possible flows of a single commodity (such as highway traffic, or water, natural gas, rail traffic, telecommunications traffic, jobs in a job shop, etc.) from an origin node to a destination node through a set of intermediate nodes. The *system operator* seeks to move as much material through the network from origin to destination as the arc capacities will allow. A classic result in the theory of network flows states that the maximum flow volume is equal to the minimum capacity of any cut, where a cut is a set of arcs such that every path from the

origin to the destination passes through at least one arc in the set, and the capacity of that cut is the sum of the capacities of the arcs in the cut. The loss of all of the arcs in any cut therefore reduces the maximum-flow volume to zero, and so any cut is a set of critical arcs. But what about smaller sets of arcs?

In this paper, we revisit the definition of a *most-vital arc* (or, more generally, *component*) as one whose removal decreases the resulting maximum flow by the greatest amount, illustrating it using a historical example: the Soviet railroad system in the 1950s. Despite the conceptual simplicity of this definition, *no simple rule exists for actually identifying such an arc*. We use this example to motivate an *attacker-defender system interdiction model*, which identifies the worst-case disruption that an intelligent and observant adversary can mount given limited attack capability. We then show that tracing out the worst-case disruption as a function of attack capability provides a natural means to assess the resilience of the system as a whole. This analysis yields a corollary result that common-sense rules of thumb for ranking the importance or criticality of individual system components are invalid. Finally, we introduce a definition of "operational resilience" that follows naturally from this early work on most-vital arcs and maximum-flow problems.

BACKGROUND

The study of vital arcs is intimately tied to the study of network flow problems, and both have their roots in military operations research. As documented by Schrijver (2002), early work on the maximum-flow minimum-cut (*max-flow min-cut*) theorem for network flows was conducted at the RAND Corporation (e.g., Ford and Fulkerson 1954, Fulkerson and Dantzig 1954, Dantzig and Fulkerson 1955) alongside a study that specifically investigated the carrying capacity of the Soviet railway system to convey military materiel from the Soviet Union to confront North Atlantic Treaty Organization forces (Harris and Ross 1955). Later at RAND, Wollmer (1963) studied rail systems "to find the link, which if removed, would reduce the capacity of the network the most." Such a link became known as the *most-vital arc*.

Sometimes There Is No "Most-Vital" Arc: Assessing and Improving the Operational Resilience of Systems

Dr. David L. Alderson,
Dr. Gerald G. Brown, and
Dr. W. Matthew Carlyle

Naval Postgraduate School
dlalders@nps.edu,
ggbrown@nps.edu,
mcarlyle@nps.edu

Dr. Louis Anthony (Tony)
Cox, Jr.

Cox Associates
TCoxDenver@aol.com

APPLICATION AREAS:
Resilience, Interdiction,
Attacker-Defender,
Defender-Attacker-
Defender, Vulnerability,
Operator model,
Optimization

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

Wollmer (1964) considers a more general version that we might call the *k most-vital arcs* problem: “given a maximum flow network from which n links are to be removed, which n arcs, if removed, would reduce the maximum flow from source to sink the most and what would be the maximum flow?” (Wollmer used n instead of k , but we prefer the latter to avoid confusion with the standard definition of n as the number of nodes in a network flow model.) Wollmer solves this problem by taking the topological dual of the original maximum-flow problem, so that finding the minimum cut is equivalent to finding the shortest path through the dual. A drawback of Wollmer’s technique is that it requires the original graph to be planar, meaning that it can be drawn so that no two arcs intersect each other except at nodes. This transformation converts the original maximum-flow problem to a shortest-path formulation where one seeks the k arcs in the dual that when assigned zero length reduce the shortest path the most.

This early work spawned a flurry of activity in model extensions for maximum-flow interdiction problems and improvements to the algorithms for solving them. Wollmer (1968) studies a stochastic variation of this problem in which the reduction in capacity on each interdicted arc is a random variable with known mean and variance, and the overall goal is to identify within specified confidence intervals the k arcs that maximally reduce the expected capacity of the network. Lubore et al. (1971) provide a more efficient algorithm for solving Wollmer’s original (1963) problem. McMasters and Mastin (1970) introduce a “budgetized” version of the problem: given a cost for removing each arc and an overall interdiction budget, find the set of arcs whose removal decreases the maximum flow the most. Ratliff et al. (1975) provide a technique for finding k most-vital arcs that works for both planar and nonplanar networks. Corley and Chang (1974) consider the problem of finding the k most-vital *nodes* that, if removed, would reduce the maximum flow the most. They show that this can be solved by augmenting the original flow network such that each node is replaced by a pair of nodes connected by a single arc, and then solving for the k most-vital of these augmented arcs.

Not surprisingly, the notion of “most vital” has also been studied from the perspective of shortest path problems. Fulkerson and Harding (1977) show how to use a limited budget for lengthening arcs in order to maximize the shortest path. Golden (1977) solves for the least-cost means of lengthening arcs so as to increase the shortest path in a network above a specified length. Corley and Sha (1982) consider the problem of finding the most-vital arc (and node) within a shortest path problem, where all arc costs are the same. Malik et al. (1989) provide an improved algorithm for solving this problem. Ball et al. (1989) establish the NP-hardness of most-vital-arc (and most-important-arc) problems in a shortest path context.

The study of vital arcs has recently continued in the context of *network interdiction problems*, starting with Wood (1993). An important part of this work has been the connection to two-person zero-sum games (Washburn and Wood 1995), and their application to stochastic network interdiction (Cormican et al. 1998), shortest path problems (Israeli and Wood 2002), and multicommodity network models (Lim and Smith 2007). Most recently, these ideas have been applied to the study of critical infrastructure systems (e.g., Brown et al. 2005, 2006), with specific attention toward electric power systems (Salmerón et al. 2004, 2009), facility location problems (e.g., Church and Scaparra 2006, Scaparra and Church 2008), supply chain networks (Snyder et al. 2006), telecommunication systems (Murray et al. 2007), and transportation problems (Alderson et al. 2011). Lunday and Sherali (2012) pose and solve some min-max models depicting interdiction planning to maximize the probability of intercepting a lone evader attempting to traverse a network from some source to some destination. Both overt and covert search efforts are considered, and types of resources, when combined, can return super-additive improvements in search effectiveness.

THE 1950S SOVIET RAIL SYSTEM

Harris and Ross (1955) model the movement of military materiel from the Soviet Union

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

into Europe as a network flow problem, using vertices to represent geographically distributed “railway divisions,” and arcs to abstract the aggregate capacity of the rail connections between each pair of adjacent divisions. In this case the minimum cut represents not only the capacity of the network as a whole, but also identifies the arcs whose removal would yield a complete interdiction of network flows. Figure 1 depicts the rail system studied in Harris and Ross.

Whereas a maximum-flow problem of this scale was considered large at the time, modern modeling languages and computing power make it the kind of problem that students might solve as a homework assignment. A much more difficult problem to solve, and one that is more aligned with modern interests, is: which arc or subset of arcs is most vital to the movement of materiel through this network? The operational importance of this question is immediate. An adversary

looking to use limited attack resources wants to plan effects-based targeting (e.g., DoD 2002, p. I-5). Or a defender of this system wants to know where to invest limited defensive resources in order to obtain mission assurance, i.e., the ability to maintain throughput capacity even in the presence of limited disruptions. We proceed in support of these objectives.

Minimizing Maximum Flow

Consider a transportation system operator who is moving some commodity (materiel, fuel, etc.) through a capacitated flow network consisting of a directed graph $G = (N, E)$, where N is a set of nodes, E is a set of undirected edges connecting node pairs (where we assume $i < j$ for all edges $(i, j) \in E$), and each edge has two associated directed arcs $(i, j) \in A$ and $(j, i) \in A$, one in each direction, and the combined flows on these two arcs has an upper bound u_{ij} . The

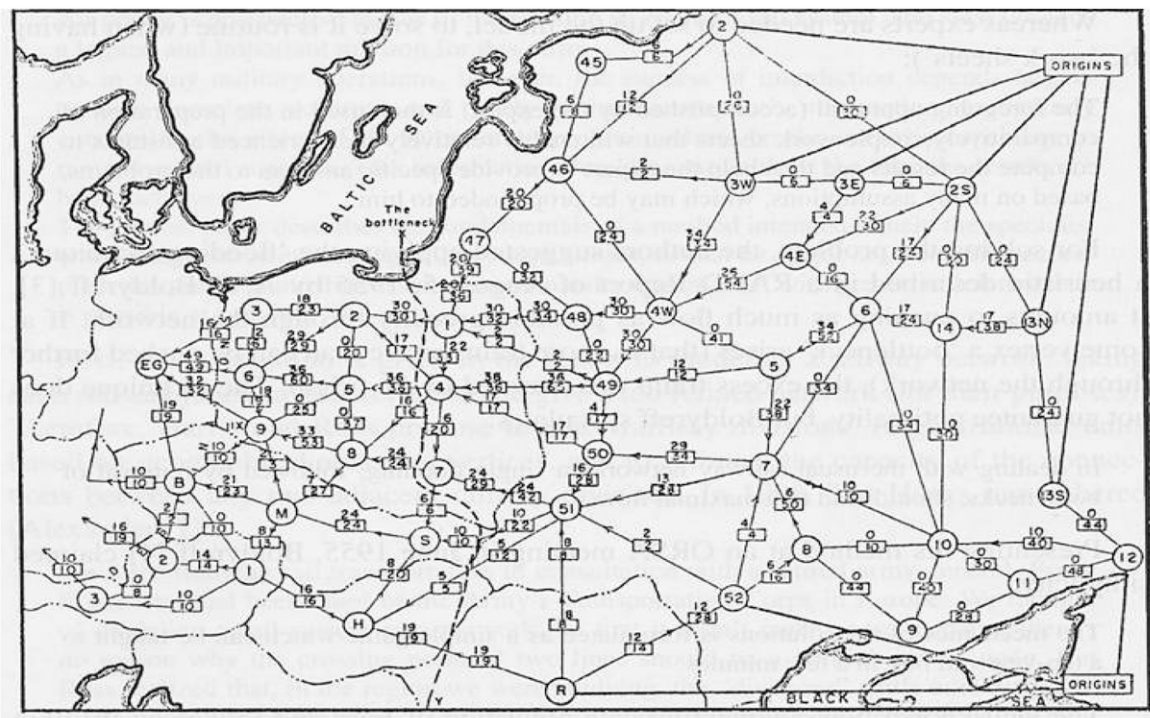


Figure 1. The Soviet rail system, circa 1955, as presented by Harris and Ross (1955). Nodes represent organizational units called “divisions,” and arcs represent the aggregate capacity to move cargo (measured in thousands of tons) between divisions.

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

operator’s objective is to maximize the flow through this network from some distinguished source node s to some other distinguished terminal node t .

Suppose that an attacker has the capability to damage a limited number of edges, rendering each arc associated with such a damaged edge useless, and must decide which edges in the network to destroy so that the operator’s maximum flow is minimized—perhaps to zero. We formulate problem MINMAXFLOW as follows.

Index Use

$i \in N$	node (alias j); where $n = N $
s, t	origin (source) node, and destination (terminal) node
$(i, j) \in E$	undirected edge between nodes i and j ; where $m = E $ $i < j \forall (i, j) \in E$
$(i, j) \in A$	arc directed from node i to node j $(i, j) \in E \Leftrightarrow i < j \wedge ((i, j) \in A \wedge (j, i) \in A)$

Data [units]

$u_{i,j}$	upper bound on total (undirected) flow on edge (i, j) [flow]
$v_{i,j}$	per-unit penalty cost on damaged arc $(i, j) \in A$ [cost/flow]
$num_attacks$	maximum number of edges the attacker can destroy [cardinality]

Decision Variables [units]

$Y_{i,j}$	defender flow on directed arc $(i, j) \in A$ [flow]
$X_{i,j}$	1 if attacker destroys undirected edge $(i, j) \in E$, 0 otherwise [binary]

Minimax optimization of flow [dual variables]

$$\min_{X \in \Xi} \left\{ \begin{array}{l} \max_Y Y_{t,s} - \sum_{(i,j) \in E} (v_{i,j} Y_{i,j} + v_{j,i} Y_{j,i}) X_{i,j} \\ s.t. \quad \sum_{(i,j) \in A} Y_{i,j} - \sum_{(j,i) \in A} Y_{j,i} = 0 \quad \forall i \in N \quad [\alpha_i] \\ 0 \leq Y_{i,j} + Y_{j,i} \leq u_{i,j} \quad \forall (i,j) \in E \quad [\beta_{i,j}] \end{array} \right\}$$

where

$$X \in \Xi = \left\{ \begin{array}{l} \sum_{(i,j) \in E} X_{i,j} \leq num_attacks \\ X_{i,j} \in \{0, 1\} \quad \forall (i,j) \in E \end{array} \right\}$$

We have added dual variables, α and β , to the balance-of-flow and capacity constraints in the maximum-flow inner problem. These will help us reformulate (and solve) the min-max problem. The (finite) penalty cost $v_{i,j}$ can be chosen to be any number greater than 1; any unit of flow across an attacked edge will contribute one unit of flow to the objective (indirectly via the balance of flow constraints and $Y_{t,s}$), but will cost at least that much in terms of penalties paid directly on that arc. If $v_{i,j} = 1$, then the operator is completely indifferent to sending flow over arcs associated with the interdicted edge, and the resulting problem may therefore have many equivalent optimal solutions. For any value $v_{i,j} > 1$, he will be penalized for that flow, and therefore will not send any flow across the interdicted arc. Because we typically require that our data be integer, $v_{i,j} = 2$ is an obvious choice. The case in which $0 < v_{i,j} < 1$ might be interpreted as fractional losses across a particular arc (perhaps from a leak in a pipe), but, unfortunately, this doesn’t work out; the balance of flow constraints still deliver all of the flow to $Y_{t,s}$, and all that a fractional penalty does is change the objective function without changing the arc flows in the maximum-flow solution.

The limitations on the attacker’s actions are simple cardinality constraints; however, we can easily adapt these to situations in which some edges are more costly to destroy than others in terms of some resource limiting the attacker, and there could even be multiple constraints on various attacker resources such as manpower, ordnance, delivery capacity, etc. Neither of these poses any conceptual or algorithmic difficulty for solving these problems.

If we wish to make an arc (or set of arcs) invulnerable, we just set the penalty cost for each invulnerable arc to $v_{i,j} = 0$. Then interdiction of the edge associated with that arc has no effect on the operator’s flow across the arc, and would be wasted effort for the attacker.

By taking the dual of the inner (maximization) problem, we obtain an equivalent mixed integer linear program minimizing flow, denoted DUAL-ILP.

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

$$\begin{aligned}
 & \min_{\alpha, \beta, X} \sum_{(i,j) \in E} u_{i,j} \beta_{i,j} \\
 \text{s.t.} \quad & \alpha_i - \alpha_j + \beta_{i,j} + v_{i,j} X_{i,j} \geq 0 & \forall (i,j) \in E \\
 & \alpha_j - \alpha_i + \beta_{j,i} + v_{j,i} X_{i,j} \geq 0 & \forall (i,j) \in E \\
 & \alpha_t - \alpha_s + \beta_{t,s} \geq 1 \\
 & \sum_{(i,j) \in E} X_{i,j} \leq \text{num_attacks} \\
 & \alpha_s = 0 \\
 & \beta_{i,j} \geq 0 & \forall (i,j) \in E \\
 & X_{i,j} \in \{0,1\} & \forall (i,j) \in E
 \end{aligned}$$

Here we fix $\alpha_s = 0$ as is customary in min-cut formulations (see Ahuja et al. 1993): the dual variables α only appear as pairwise differences, and therefore have an extra degree of freedom we can eliminate by fixing any one of them to a constant value. Using a feasible binary attack plan X^* from this mixed integer linear program, one can recover the operator’s residual flows Y^* by solving the operator’s maximizing

linear program for this fixed X^* . (The values of the dual variables do not directly support calculation of the optimal flows; they can, in fact, be noninteger, even though we would expect to be able to interpret them as node and arc labels as they would be in a typical max-flow, min-cut formulation.) The mixed-integer linear program can be embellished by any ILP restrictions on the X variables.

Figure 2 presents a cosmetically revised version of the Soviet rail network in Harris and Ross (1955), suitable for use as input to the mixed integer linear program interdiction problem. (We have renumbered some of the nodes, and removed redundant capacity information on some arcs).

We explore the effect of an increasing number of worst-case attacks on the ability of the operator to move materiel through this system. That is, by solving MINMAXFLOW for $\text{num_attacks} = 1, 2, \dots, n$, we discover how the system will perform under an increasing number of attacks. Figure 3 presents the results.




Figure 2. Network used as input to the network interdiction problem. The set of origin nodes is $\{1, 2, 3, 5, 10, 15, 16, 17, 26\}$, and the set of destination nodes is $\{42, 43, 44, 48, 49, 50\}$. Edge labels represent capacities (in 1,000s of tons). (This figure displays all data necessary to reproduce the computational experiments reported in this paper.)

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS




Figure 3. Maximum flow through the Soviet rail system after an increasing number of worst-case attacks. The solution of 163,000 tons for zero attacks is exactly the solution to the original maximum-flow problem. It takes seven simultaneous attacks to achieve a complete interdiction of network flow (i.e., a cut).

The results in Figure 3 show that an attacker obtains approximately linear returns with each additional attack. This is not good news for the system operator, but it could be worse. Many transport systems are built with minimal redundancy, which, in the extreme case of a spanning tree (e.g., many pipeline systems), means that any single attack can yield a complete interdiction.

Table 1 shows the edges associated with each worst-case attack. We observe that the sets of edges for one through five attacks are *monotone* (or *nested*, or *prioritizable*), in the sense that the set of edges for $k+1$ attacks includes all of the edges in the set for k attacks, plus one additional edge. This type of result suggests the use of priority lists as a natural means for organizing a list of potential targets. However, the set of edges associated with $num_attacks = 6$ does not contain the set for $num_attacks = 5$.

This problem of minimizing the maximum flow is a specific instance of a much broader class of problems involving network (or system) interdiction. Such models have become popular tools for studying the interaction of a strategic attacker and defender. But in the case of our Soviet rail example, this seems like a lot of work for what seems intuitive to anyone who has studied network flows.

Identifying the Most-Vital Arc

The connection between the maximum flow and the minimum capacity cut is so fundamental that it seems obvious how to identify the bottlenecks, and therefore the “most-vital” arcs, in a maximum-flow problem. Indeed, the intuitive nature of the problem suggests several appealing rules for identifying them (Ahuja et al. 1993, p. 244):

Table 1. Edges associated with worst-case attacks. For one to five attacks, the set of edges is monotone. With seven attacks, we have a complete interdiction of network flow.

Attacks	Maxflow	Attacked Edges							
0	163								
1	127	(35,40)							
2	97	(35,40)	(34,39)						
3	73	(35,40)	(34,39)	(37,45)					
4	49	(35,40)	(34,39)	(37,45)	(36,41)				
5	32	(35,40)	(34,39)	(37,45)	(36,41)	(35,39)			
6	15	(34,35)	(34,39)	(26,38)	(22,35)	(24,36)	(23,35)		
7	0	(35,39)	(34,39)	(35,40)	(35,41)	(36,41)	(37,45)	(38,46)	

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

- An arc having the largest capacity is most vital;
- An arc carrying the largest flow in an optimal solution is most vital;
- An arc having the largest capacity in a minimum-capacity cut is most vital; or
- Any most-vital arc is in some minimum-capacity cut.

This section shows that *none* of these intuitive criteria correctly identifies the most-vital arc.

Figure 4 illustrates a maximum-flow problem with seven nodes and nine arcs, where the system operator seeks a set of feasible arc flows, y_{ij} , to maximize the total quantity sent from node 1 to node 7 per unit time, while abiding by individual arc capacity limits, u_{ij} , and material balance (inflow = outflow) at all intermediate nodes. For this network, the maximum flow is 20 units, with arc flows as indicated on the diagram, and there are two minimum-capacity cuts (illustrated by dashed curved lines), each with capacity 20, proving that the flow is optimal. In fact, the most-vital arc is the one from node 1 to node 3, denoted as arc 1-3.

If an arc carries zero flow in any optimal solution, then removing that arc will not reduce the maximum volume of flow. But if an arc carries at least some minimum flow in every optimal solution, then its removal will reduce the optimal solution by the amount of that minimum flow. If we let $m(a)$ represent the minimum flow across arc a over *all* maximum-flow solutions, then an arc a is a *most-vital arc* if it

maximizes $m(a)$ over A . In Figure 1, arc 1-3 has a minimum possible flow of 15 units (and a maximum of 20) over all optimal solutions. Arc 5-7 could also carry as much as 20 units of flow in an optimal solution, but, as illustrated, it could also carry *no* flow in an optimal solution, and hence cannot be a most-vital arc.

This definition of a most-vital arc (maximizing over all optimal solutions the minimum flow over all arcs) is conceptually simple, but computationally complex: We know of no simple rule of thumb that avoids lengthy computations and that can successfully identify which arc(s) satisfy this definition. Determining a most-vital arc (and, more generally, a *set* of most-vital arcs of any given size) requires solving a maximum-flow interdiction problem like MINMAXFLOW that is of size comparable to the original maximum-flow problem. Although straightforward to formulate, larger instances of this integer linear programming problem can be very difficult to solve.

In some cases, it may seem easier or more convenient to use *random sampling* as an alternate means for selecting “vital” components of the system. The idea is to specify a probability distribution for the possible combinations of arc failures and then select a large sample of (presumably) representative scenarios, keeping track of the worst ones. There are two potential problems with doing this. First, it is unclear how to choose a “good” probability distribution. Second, because there are $\binom{m}{k} = O(m^k)$




Figure 4. Network with a maximum flow of 20 units from node 1 to node 7 and with the two minimum-capacity cuts indicated (dashed curves), illustrating a contradiction for each of several proposed characterizations of a “most-vital” arc. Notation: The two numbers on each arc are respectively (flow, capacity). Arc 1-3 is the most-vital arc; removing it reduces the maximum flow to five, and removing any other single arc allows at least ten units of flow from 1 to 7 in the resulting network. Arc 5-7 has the largest capacity, arcs 5-6 and 6-7 have the maximum flow, and arcs 3-4, 3-5, and 4-5 are the only arcs in any minimum capacity cut.

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

ways to have k failed arcs, for $k = 1, 2, \dots, m$, the number of combinations could be so large that such sampling is ineffective even on fast computers.

To explore this idea of sampling a bit further, we plot in Figure 5 the residual maximum flow of the network under a large number of possible disruptions according to the number of damaged edges k . In general, a system with k damaged components will achieve a range of performance values, depending on which components are damaged. For each value of k , we randomly sample (with replacement) 10,000 possible configurations of disrupted edges, where each edge failure is equally likely and

independent of the others, and then solve each configuration for the remaining maximum flow. This allows us to compare the distribution of residual maximum-flow values from random sampling with that of the optimal (worst-case) attack. For small k , this random sampling obtains optimal or near-optimal solutions, largely because the random sample is equal in size or larger than the total number of configurations and therefore is nearly performing an exhaustive enumeration of the solution space. However, for larger k , random sampling fails to find anything close to the worst-case attack because of the enormous number of possible attack configurations (e.g., there are more than 10^{10} ways

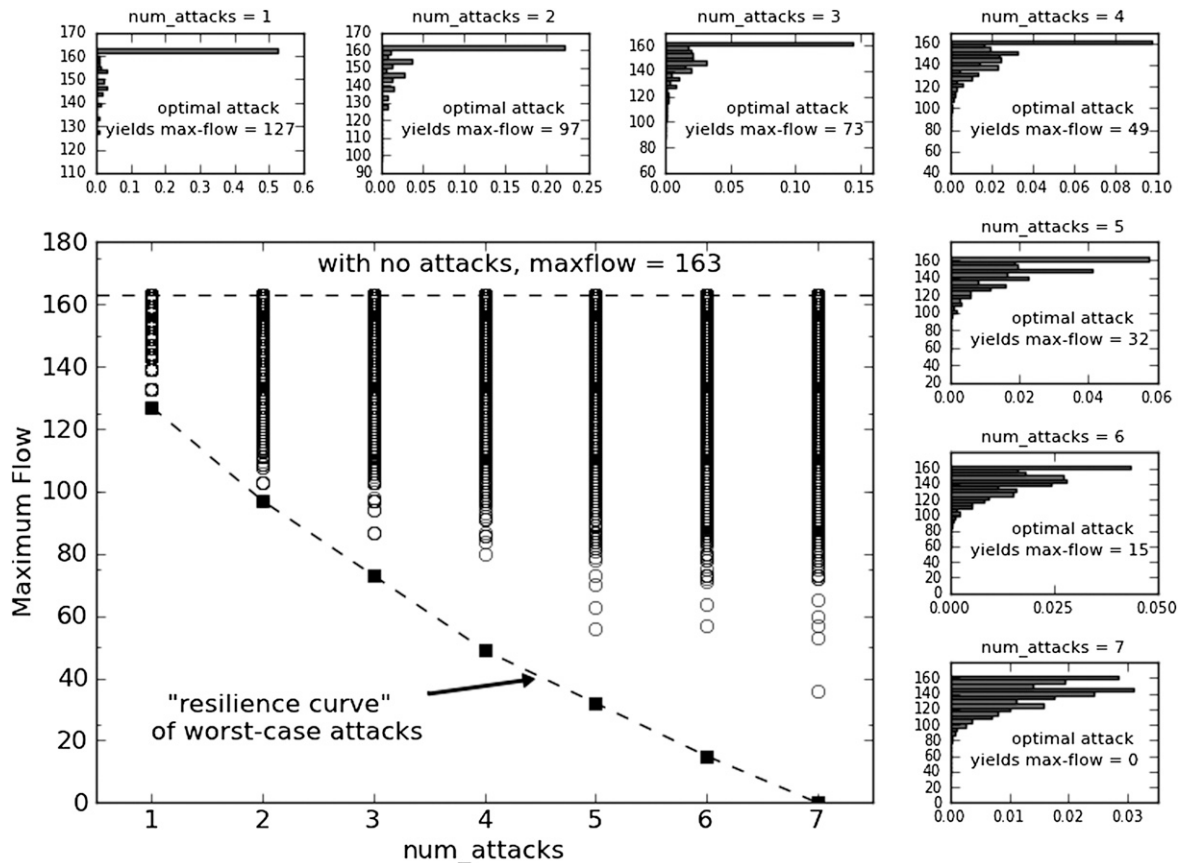


Figure 5. Random attacks on the Soviet Railway, compared with optimal ones. Maximum flow (system performance) degrades with an increasing number of damaged edges. For $num_attacks = 1, 2, \dots, 7$, we present the worst-case disruption, along with 10,000 randomly generated attacks. Each subfigure shows a histogram of the frequency with which random attacks impact the maximum flow. As the number of possible attack combinations increases, it becomes harder and harder to find the worst-case attack by random sampling. In the main figure, the dashed line connects the worst-case disruptions for this system. (Here the phrase “resilience curve” really refers to a discrete frontier of points.)

of choosing seven attacks) and because, for this network, there are only a small number of optimal or near-optimal attacks.

The implications of Figure 5 are striking: we should not expect random sampling of failure scenarios (e.g., via simulation) to reliably answer the question, “How bad could it be?” Rather, we need to solve explicitly for worst-case disruptions.

All this effort to find the absolute worst-case disruption might seem like unnecessary work, were it not for the key role that worst-case disruptions play in assessing the resilience of the system.

A Measure of Operational Resilience

The notion of *resilience* has become an important concept in discussions about critical infrastructure. In its 2006 report, the Critical Infrastructure Task Force of the Homeland Security Advisory Council defined resilience as “the capability of a system to maintain its functions and structure in the face of internal and external change and to degrade gracefully when it must.” The 2007 National Strategy for Homeland Security recognizes that although we cannot prevent all disruptions, deliberate or nondeliberate, we can work to ensure “the structural and operational resilience of critical infrastructures and key resources” (HSC 2007, p. 27), adding that “We must now focus on the resilience of the system as a whole—an approach that centers on investments that make the system better able to absorb the impact of an event without losing the capacity to function” (HSC 2007, p. 28).

There is now a host of competing definitions for resilience that span applications in human and organizational behavior (e.g., Bennis and Heifetz 2003), system safety (e.g., Hollnagel, Woods, and Leveson 2006, and references therein), systems engineering (e.g., Haimes 2009), and control theory (e.g., Vugrin et al. 2010).

In the context of maximum-flow problems, (and, more generally, for infrastructure systems), we propose a definition for operational resilience that follows directly from our discussion of most-vital arcs, and we introduce the notion of a resilience curve that defines the response of a system to a range of disruptions.

Consider a network in which each edge is operating, but which might be damaged such that some edges are transformed from operating to nonfunctional, and where the performance of the network (here, the maximum-flow volume) is a function of the collective state of its edges. Let c be a vector of length m , each element of which represents the binary state (operating or not) of an edge, such that there is a total of 2^m distinct configurations of the network. For any scenario defined by the vector c we define the *magnitude of disruption* as

$$\Delta c = \sum_{i=1}^m (1 - c_i),$$

which, for a vector of binary values, is simply the number of failed edges. In general, the performance of the network will degrade with an increasing number of failed edges, but not all edges are equally critical to the maximum flow. What does this mean for operational resilience?

We introduce the *resilience curve* for the system as that which plots worst-case performance as a function of disruption magnitude (see Figure 5). Although these plots might be more properly referred to as “resilience frontiers” due to their discrete nature, we retain the term *resilience curves* to maintain a connection to prior work on *risk curves* (Kaplan and Garrick 1981), and because we will use stylized continuous approximations when the discreteness of the attacker’s (or defender’s) level of effort is not critical to the discussion. The resilience curve communicates considerable information about the response of the system to worst-case disruptions of increasing magnitude, where magnitude is simply the number of failed edges. One immediately discerns how the maximum flow degrades as additional edges are damaged. Intervals where there is little or no change in maximum flow are called “more resilient,” and intervals where there is greater change are called “less resilient.”

When used for relative comparison between networks, these resilience curves allow us to draw conclusions about dominating alternatives just as in Kaplan and Garrick (1981). For example, in the simple case where the resilience curve of one system (say, System A) dominates the resilience curve of another (say,

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

System B), one can make assertions such as, “System A is more resilient than System B” (Figure 6a). However, when neither curve dominates (Figure 6b), this type of simple assertion is not possible. In this way, we paraphrase the comment about risk by Kaplan and Garrick, namely that “[resilience] is a concept bigger than a single number.”

Resilience curves as defined here show how the performance of a network changes in response to the loss of edges. When the size of disruption is a linear function of the resources (e.g., the number of people, materials, and money) or level of effort (e.g., number of simultaneous attacks) required to cause it, our resilience curves yield a novel interpretation: how will the system respond in the face of increasing disruption? In the case of an intelligent adversary, the resilience curve of the system reflects the attacker’s *return-on-investment* (ROI)—showing how system performance degrades with incremental expenditure of attack resources. Such information can be vital for defensive planning purposes, in that investment to “raise the resilience curve” can effectively *deter* an adversary looking to use limited resources to disrupt system performance.

Our interest in assessing the operational resilience of the system under *any disruption* suggests that we focus on the worst-case performance of the system with k damaged components. But assessment is only part of the problem: we need to know how to invest limited resources to maximally improve the resilience of the system. We again show that simple intuitive solutions often fall short of the best that we can do.

The Problem with Prioritized Lists

The Department of Defense and Department of Homeland Security typically prioritize critical infrastructure assets (system components) into “tiers” to help inform protection decisions (DoD 2002). A prioritized list of system components can be a helpful planning tool for an attacker or defender when budgets are not completely known beforehand, because this defines a simple rule for allocating resources: If we can afford to attack (or defend) k targets, choose the top k components from the list. This “greedy” (myopic) decision rule is easy to implement and sometimes serves as a reasonable first guess at a solution, but only in special circumstances does it optimize use of resources (e.g., Magazine et al. 1975).

Figure 7 illustrates a weakness inherent in creating any “prioritized list” of targets (or assets or system components) to protect. In this example, the most-vital single arc is *not* among the best choices if two (or more) arcs are to be removed. More generally, we observe that none of these most-vital sets is prioritizable, in the sense that none of the k most-vital arcs is included in the set of $k+1$ most-vital arcs. Therefore, the concept of identifying a priority list with a “most-vital” arc followed by a “second-most-vital” arc (or set of arcs), etc., is fundamentally flawed, because how “vital” an arc is depends (nonmonotonically) on *how many arcs an attacker can afford to target simultaneously*—or, more generally, on the attacker’s resources available for inflicting damage. If, instead, we seek a set of k arcs whose simultaneous removal most reduces the resulting maximum flow, then

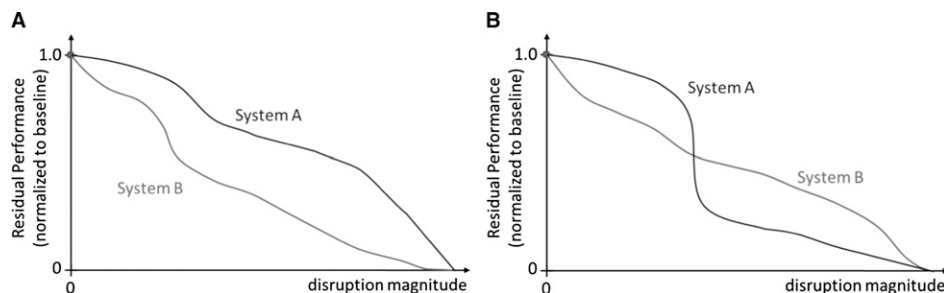


Figure 6. (a) System A has greater resilience than System B. (b) Neither System A nor System B can be said to be more resilient. System A is more resilient to a smaller number of attacks, but System B is more resilient to a larger number of attacks.

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS




Figure 7. Why prioritized lists don’t always work. This network has four different minimum cut sets, each yielding a maximum flow of 25 units. The single most-vital arc is the one with capacity 9; losing it reduces the maximum flow to 16. The two most-vital arcs are the pair having capacity 8 units; losing them reduces the maximum flow to 9. The three most-vital arcs are those with capacity 7, and the four most-vital arcs are those with capacity 6. None of these most-vital sets is nested within another, meaning that there is no single priority list that accurately characterizes the importance of each arc.

in general we need to abandon tiers or priority rankings, and, instead, actually solve a maximum-flow interdiction problem for each value of k . There is no simple rule of thumb for identifying critical sets of arcs (or, in general, components) without analyzing the system’s operation as a whole. Although a well-designed set of priority tiers can sometimes be useful as an approximate guide to how to allocate resources, especially when system performance is not strongly affected by multiway interactions among components or subsets of components, simple rules cannot find useful approximations (Magazine et al. 1975).

In the Appendix, we show how to construct, for any integer n , a maximum-flow problem for which the sets of k -most-vital arcs for $k = 1, \dots, n$ are pairwise disjoint. We then discuss how far from this optimal sequence a prioritized list (which induces a monotone collection of sets of arcs as targets) can be, even if the best prioritized list is found. Finally, we show that determining the best prioritized list is a difficult problem in its own right, even when the optimal sequence is known.

Improving Resilience with a Limited Budget

Assume that we have the ability to defend a limited number of edges in our network, whereby defending an edge makes it invulnerable to attack. Which edges should we defend, and what does this do to improve the resilience of the system? We can formulate this decision problem as follows.

Additional Index [cardinality]

$d \in D$ defense options [few, 2-10]

Additional Data [units]

v_{ij}^d increased cost per unit flow on directed arc $(i, j) \in A$ if attacked under defense option $d \in D$ [cost/kton]

u_{ij}^d capacity of edge $(i, j) \in E$ under defense option $d \in D$ [tons]

$num_defenses$ maximum number of edges the defender can protect [cardinality]

Decision Variables [units]

Y_{ij}^d Flow of traffic on directed arc $(i, j) \in A$ under defense option $d \in D$ [tons]

$W_{ij}^d = 1$ if defense option d chosen for edge $(i, j) \in E$, 0 otherwise [binary]

Formulation MAXRESILIENCE

$$\max_W \min_X \max_Y \sum_{d \in D} Y_{t,s}^d - \sum_{\substack{(i,j) \in E \\ d \in D}} (v_{i,j}^d Y_{i,j}^d - v_{j,i}^d Y_{j,i}^d) X_{i,j} \quad (D0)$$

$$s.t. \sum_{\substack{(i,j) \in A \\ d \in D}} Y_{i,j}^d - \sum_{\substack{(j,i) \in A \\ d \in D}} Y_{j,i}^d = 0 \quad \forall i \in N \quad (D1)$$

$$0 \leq Y_{i,j}^d + Y_{j,i}^d \leq u_{i,j}^d W_{i,j}^d \quad \forall (i, j) \in E, \forall d \in D \quad (D2)$$

$$\sum_{(i,j) \in E} X_{i,j} \leq num_attacks \quad (D3)$$

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

$$X_{i,j} \in \{0, 1\} \quad \forall (i, j) \in E \quad (D4)$$

$$\sum_{d \in D} W_{i,j}^d = 1 \quad \forall (i, j) \in E \quad (D5)$$

$$\sum_{\substack{(i,j) \in E \\ d \in D}} W_{i,j}^d \leq \text{num_defenses} \quad (D6)$$

$$W_{i,j}^d \in \{0, 1\} \quad \forall (i, j) \in E, \forall d \in D \quad (D7)$$

Discussion

The objective function (D0) represents the value of the maximum flow, for a defense option chosen for each edge, an attack plan, and a set of flows in the resulting network. We assume there is a “do nothing” defense plan, $d_0 \in D$, that grants each edge its original capacity, $u_{ij}^{d_0} = u_{ij}$, and unhardened attack penalties, $v_{ij}^{d_0} = v_{ij}$, from the prior models. There could be several ways to defend a particular edge, each with a different penalty and capacity, but for simplicity of exposition we assume there are exactly two: one, d_0 , which does nothing to reduce the damage of an attack ($v_{ij}^{d_0} = 2$), and one, d_1 , which completely nullifies any attack ($v_{ij}^{d_1} = 0$). Constraints (D1) enforce balance of flow at each node. Constraints (D2) limit the total flow on the two directed arcs associated with edge $(i, j) \in E$ to not exceed the total capacity granted by the defense option chosen for that edge. Constraint (D3) limits the number of edges that can be attacked, and (D4) enforces binary decisions about which edges are attacked. Constraints (D5) force the defender to choose exactly one defensive plan per edge. Constraint (D6) limits the number of edges that can be defended. Of course, as was the case with the attacker cardinality constraint, these could be generalized to include several different types of defender budgets. Stipulations (D7) specify that selecting a defense option for each edge is a binary decision.

Formulation MAXRESILIENCE is an example of a *defender-attacker-defender* (DAD) model (Brown et al. 2006, Alderson et al. 2011). For a given level of defensive effort, represented here as *num_defenses*, an optimal solution identifies which edges should be defended (syn. hardened) and by how much

this helps mitigate the impact of deliberate attacks.

Figure 8 shows the resulting resilience curve when defending *num_defenses* = 1, 2, ..., 7 possible edges. We observe that a single defended edge provides little benefit in terms of the residual throughput following a worst-case attack, but it does increase the number of attacks needed for complete interdiction from 7 to 8. A second defense increases this number to 10, and a third defense increases this number to more than 10. However, in many cases these defenses provide only a relatively small increase to the actual maximum flow for a given number of attacks. This chart reveals the ROIs (in each attack scenario) that a defender faces when planning defensive investments. If the benefits of increased flow volume can be stated in the same units as the costs of the defenses (e.g., dollars), then the tradeoffs he faces in each attack scenario can be evaluated in terms of absolute benefits.

RELATED WORK

There are a number of other techniques in use for assessing criticality of infrastructure components and prioritizing their protection.

Risk Scoring Techniques

Simple formulas such as $Risk = Threat \times Vulnerability \times Consequence$, or $Risk = Frequency \times Impact$, with judgment-based rating scales or scores used to assess the inputs on their right-hand sides, are widely applied to assess and compare the importance of individual infrastructure components (e.g., see ASME 2008). Most such rating systems do not account for uncertainties, correlations, or dependencies among the inputs; exploit portfolio effects; help to diversify protective investments against common uncertainties in inputs; or optimize risk reductions for resources spent (Cox 2009). Hence, they are seldom satisfactory for supporting objectively good risk management decisions, although their popularity suggests that they may satisfy other needs, such as an urge to impose simple structures on complex problems.

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

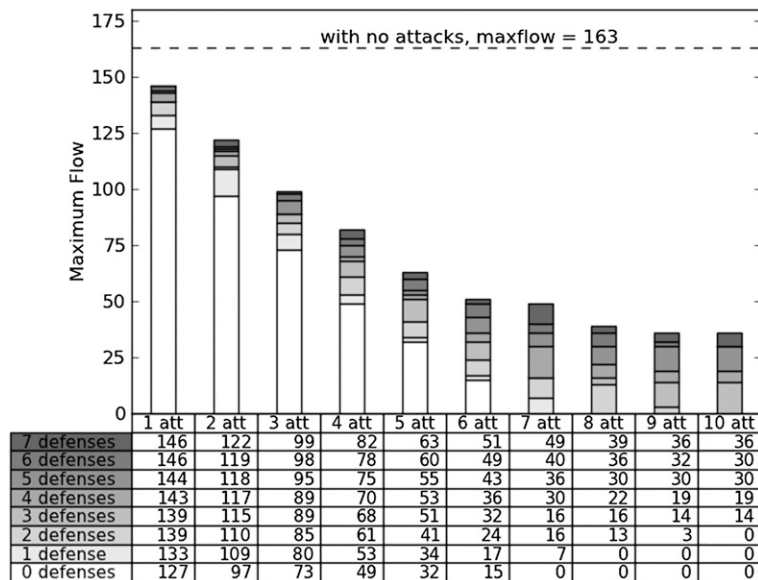


Figure 8. Increases to the resilience curve for the Soviet Railway. Defending edges provides marginal increases to the resilience curve when there are only a small number of attacks, but it serves to increase the number of attacks required for complete system interdiction.

Graph Connectivity and Network Science

Owing in part to the massive size of modern infrastructure networks and the vast amounts of data now being collected about them, recent efforts in the study of “network science” (NRC 2006) have focused on the structure and behavior of very large networks in physical, biological, and social systems. Network science measures function in these systems primarily in terms of graph connectivity statistics (Newman 2003). In this context, vital arcs are those that contribute most to these graph theoretic measures, such as the average path length between every pair of nodes or the size of the largest connected component (e.g., Albert et al. 2000; Holme et al. 2002). A drawback of this is that when applied to real systems (e.g., Albert et al. 2004; Schneider et al. 2011), these simple measures of connectivity often fail to capture the most salient features of network function (e.g., Doyle et al. 2005, Hines et al. 2010), making them of limited value to operators of real network-centric infrastructure systems (Alderson 2008, Alderson and Doyle 2010).

CONCLUSION

No simple rule(s) of thumb can identify the most-critical system components when the components work together to deliver system capacity, such as throughput. Rather, the *criticality of components depends on which sets are damaged or destroyed by an attack* (or by reliability failures, natural disasters, and other nonadversarial hazards).

The maximum-flow problem is an example of a simple, even primitive, operator’s model that shows how the system responds to losses of sets of arcs (components). Manipulation of such a model can reveal a system’s functional capability and remaining vulnerabilities, and can guide measures to identify protective investments that will maximize “the resilience of the system as a whole” for resources spent.

REFERENCES

Ahuja, R. K., Magnanti, T. L., and Orlin, J. B. 1993. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall.

Albert, R., Albert, I., and Nakarado, G. L. 2004. Structural Vulnerability of the North

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

- American Power Grid, *Physical Review E*, Vol 69, 025103.
- Albert, R., Jeong, H., and Barabási, A.-L. 2000. Attack and Error Tolerance of Complex Networks, *Nature*, Vol 406, 378–382.
- Alderson, D.L. 2008. Catching the “Network Science” Bug: Insight and Opportunity for the Operations Researcher, *Operations Research*, Vol 56, No 5, 1047–1065.
- Alderson, D.L., Brown, G.G., Carlyle, W.M., and Wood, R.K. 2011. Solving Defender-Attacker-Defender Models for Infrastructure Defense, *Operations Research, Computing, and Homeland Defense.*, R.K. Wood and R.F. Dell, eds. Institute for Operations Research and Management Science (INFORMS), 28–49.
- Alderson, D.L., Doyle, J.C. 2010. Contrasting Views of Complexity and Their Implications for Network-Centric Infrastructures, *IEEE Transactions on Systems, Man, and Cybernetics-Part A*, Vol 40, No 4, 839–852.
- ASME Innovative Technologies Institute. RAMCAP, Risk Analysis and Management for Critical Asset Protection. 2008. <http://www.asme-iti.org/RAMCAP>.
- Ball, M.O., Golden, B.L., and Vohra, R.V. 1989. Finding the Most Vital Arcs in a Network, *Operations Research Letters*, Vol 8, No 2, 73–76.
- Bennis, W.G., and Heifetz, R.A. 2003. *Harvard Business Review on Building Personal and Organizational Resilience*. Harvard Business Press.
- Brown, G., Carlyle, W.M., Salmerón, J., and Wood, K. 2005. Analyzing the Vulnerability of Critical Infrastructure to Attack, and Planning Defenses. *Tutorials in Operations Research: Emerging Theory, Methods, and Applications*, Greenberg, H., and Smith, J., eds. Institute for Operations Research and Management Science (INFORMS), 102–123.
- Brown, G., Carlyle, M., Salmerón, J., and Wood, K. 2006. Defending Critical Infrastructure, *Interfaces*, Vol 36, No 6, 530–544.
- Church R. L., and Scaparra, M. P. 2006. "Protecting Critical Assets: The r-interdiction Median Problem with Fortification. *Geographical Analysis*, Vol 39, No 2, 129–146.
- Corley, H. W., and Chang, H., 1974. Finding the n Most Vital Nodes in a Flow Network, *Management Science*, Vol 21, No 3, 362–364.
- Corley, H. W., and Sha, D. Y. 1982. Most Vital Links and Nodes in Weighted Networks, *Operations Research Letters*, Vol 1, No 4, 157–160.
- Cormican, K. J., Morton, D. P. and Wood, R. K. 1998. Stochastic Network Interdiction, *Operations Research*, Vol 46, No 2, 184–197.
- Cox, L. A. Jr. 2009. What’s Wrong with Hazard-Ranking Systems? An Expository Note, *Risk Analysis*, Vol 29, No 7, 940–948.
- Dantzig, G. B., and Fulkerson, D. R. 1955. On the Max Flow-Min Cut Theorem of Networks. The RAND Corporation, Research Memorandum RM-1418-1. Rev. April 15, 1955.
- Department of Defense. Joint Doctrine for Targeting. 2002. Joint Publication 3–60. Washington, DC.
- Doyle, J. C., Alderson, D. Li, L., Low, S., Roughan, M., Shalunov, S., Tanaka, R., and Willinger, W. 2005. The “Robust Yet Fragile” Nature of the Internet, *Proceedings of the National Academy of Sciences*, Vol 102, No 41, 14497–14502.
- Ford, L. R., and Fulkerson, D. R. 1954. Maximal Flow through a Network. The RAND Corporation, Research Memorandum RM-1400. November 19.
- Fulkerson, D. R., and Dantzig, G. B. 1954. Computation of Maximal Flows in Networks. The RAND Corporation Research Memorandum RM-1400. November 19.
- Fulkerson, D. R., and Harding, G. C. 1977. Maximizing the Minimum Source-Sink Path Subject to a Budget Constraint, *Mathematical Programming*, Vol 13, No 1, 116–118.
- Golden, B. 1977. A Problem in Network Interdiction, *Naval Research Logistics Quarterly*, Vol 25, No 4, 711–713.
- Haimes, Y. Y. 2009. On the Definition of Resilience in Systems, *Risk Analysis*, Vol 29, No 4, 498–501.
- Harris, T. E., and Ross, F. S. 1955. Fundamentals of a Method for Evaluating Rail Net Capacities. Research Memorandum RM-1573, The RAND Corporation.

- Hines, P., Cotilla-Sanchez, E., and Blumsack, S. 2010. Do Topological Models Provide Good Information About Electricity Infrastructure Vulnerability? *Chaos*, Vol 20, No 3, 033122.
- Hollnagel, E., Woods, D. D., and Leveson, N., eds. 2006. *Resilience Engineering: Concepts and Precepts*, Ashgate Press.
- Holme P., Kim, B. J., Yoon, C. N., and Han, S. K. 2002. Attack Vulnerability of Complex Networks, *Physics Review E*, Vol 65, 056109–056123.
- Homeland Security Advisory Council, 2006, *Report of the Critical Infrastructure Task Force*, U.S. Department of Homeland Security, Washington D.C.
- Homeland Security Council (HSC). 2007. *National Strategy for Homeland Security*, The White House, Washington, DC.
- Israeli, E. and Wood, R. K. 2002. Shortest-Path Network Interdiction, *Networks*, Vol 40, No 2, 97–111.
- Kaplan, S., and Garrick, B. J. 1981. On the Quantitative Definition of Risk, *Risk Analysis*, Vol 1, No 1, 11–27.
- Lim, C., and Smith, J. C. 2007. Algorithms for Discrete and Continuous Multicommodity Flow Network Interdiction Problems, *IIE Transactions*, Vol 39, No 1, 15–26.
- Lubore, S. H., Ratliff, H. D., and Sicilia, G. T. 1971. Determining the Most Vital Link in a Flow Network, *Naval Research Logistics Quarterly*, Vol 18, No 4, 497–502.
- Lunday, B. J., and Sherali, H. D. 2012. Network Interdiction to Minimize the Maximum Probability of Evasion with Synergy between Applied Resources, *Annals of Operations Research*, Vol 196, No 1, 411–442.
- Magazine, M. J., Nemhauser, G. L., and Trotter Jr., L. E. 1975. When the Greedy Solution Solves a Class of Knapsack Problems, *Operations Research*, Vol 23, No 2, 207–217.
- Malik, K., Mittal, A. K., and Gupta, S. K. 1989. The k Most Vital Arcs in the Shortest Path Problem, *Operations Research Letters*, Vol 8, No 4, 223–227.
- McMasters, A. W., and Mastin, T. M. 1970. Optimal Interdiction of a Supply Network, *Naval Research Logistics Quarterly*, Vol 17, No 3, 261–268.
- Murray, A. T., Matisziw, T. C., and Grubescic, T. H. 2007. Critical Network Infrastructure Analysis: Interdiction and System Flow, *Journal of Geographical Systems*, Vol 9, No 2, 103–117.
- National Research Council (NRC), Committee on Network Science for Future Army Applications. 2006. *Network Science*. The National Academies Press.
- Newman, M. E. J. 2003. The Structure and Function of Complex Networks, *SIAM Rev.*, Vol 45, 167–256.
- Ratliff, D. H, Sicilia, G. T., and Lubore, S. H. 1975. Finding the n Most Vital Links in Flow Networks, *Management Science*, Vol 21, No 5, 531–539.
- Salmerón, J., Wood, K., and Baldick, R. 2004. Analysis of Electric Grid Security under Terrorist Threat, *IEEE Transactions on Power Systems*, Vol 19, No 2, 905–912.
- Salmerón, J., Wood, K., and Baldick, R. 2009. Worst-Case Interdiction Analysis of Large-Scale Electric Power Grids, *IEEE Transactions on Power Systems*, Vol 24, No 1, 96–104.
- Scaparra, M. P., and Church, R. L. 2008. A Bilevel Mixed-Integer Program for Critical Infrastructure Protection Planning, *Computers and Operations Research*, Vol 35, No 6, 1905–1923.
- Schneider, C. M, Moreira, A. A., Andrade, J. S., Havlin, S., and Herrmann, H. J. 2011. Mitigation of Malicious Attacks on Networks, *Proceedings of the National Academy of Sciences*, Vol 108, No 10, 3838–3841.
- Schrijver, A. 2002. On the History of the Transportation and Maximum Flow Problems, *Mathematical Programming Series B*, Vol 91, No 3, 437–445.
- Snyder, L. V., Scaparra, M. P., Daskin, M. S., and Church, R. L. 2006. Planning for Disruptions in Supply Chain Networks, *Tutorials in Operations Research: Models, Methods, and Applications for Innovative Decision Making*, Johnson, M.P., Norman, B., and Secomandi, N., eds. Institute for Operations Research and Management Science, (INFORMS).
- Vugrin, E. D., Warren, D. E., Ehlen, M. A., and Camphouse, R. C. 2010. A Framework for

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS

- Assessing the Resilience of Infrastructure and Economic Systems, *Sustainable and Resilient Critical Infrastructure Systems: Simulation, Modeling, and Intelligent Engineering*, Gopalakrishnan, K., and Peeta, S., eds. Springer-Verlag, 77–116.
- Washburn, A., and Wood, K. 1995. Two-Person Zero Sum Games for Network Interdiction, *Operations Research*, Vol 43, No 2, 243–251.
- Wollmer, R. D. 1963. Some Methods for Determining the Most Vital Link in a Railway Network. RAND Memorandum, RM-3321-ISA. April.
- Wollmer, R. D. 1964. Removing Arcs from a Network, *Operations Research*, Vol 12, No 6, 934–940.
- Wollmer, R. D. 1968. Stochastic Sensitivity Analysis of Maximum Flow and Shortest Route Networks, *Management Science*, Vol 14, No 9, 551–564.
- Wood, R. K. 1993. Deterministic Network Interdiction, *Mathematical and Computer Modelling*, Vol 17, No 2, 1–18.

APPENDIX

In this appendix we show how to construct, for any number of attacks, a , an explicit counterexample to the notion that a prioritized list of targets can be sufficient for attack planning, for any number of attacks from 1 to a . For this example, the target system consists of a maximum-flow problem on a directed graph, and the optimal attack consisting of k arcs has *no arc in common* with the optimal attack containing k' arcs, for any $k \neq k'$ from 1 to a . We will use parallel, directed arcs to illustrate this property, but it is straightforward to replicate these results for undirected arcs, and to add extra nodes in each of these arcs to create equivalent examples that do not contain parallel arcs. As a consequence of this result, any “prioritized list” of targets for this counterexample can only yield the optimal attack plan for one value of k , and the attack plan it suggests for any other number of attacks (except zero or a) will be provably suboptimal.

Given an integer, n , we construct our instance as follows. Define a set N of n nodes, numbered 1 to n . Between nodes k and $k + 1$ (for $0 < k < n$) define n parallel arcs. The first k of these arcs have capacity k , and the remaining $n - k$ arcs have capacity $k + n$. Each of the $n - 1$ “layers” of parallel arcs between two consecutive nodes defined in this way has total capacity $k^2 + (n - k)(n + k) = n^2$, and therefore the maximum flow on this graph is n^2 , with

flow on every arc at its capacity (i.e., all arcs are saturated).

The k -most-vital arcs in this maximum-flow problem are the k largest arcs from layer $n - k$, (each of which has capacity k); if they are removed the remaining $n - k$ arcs in that layer each have capacity $n - k$, yielding a resulting optimal flow of size $(n - k)^2$. This result follows because any set of k arcs taken from any other layer will have less total capacity; layers that have larger-capacity arcs will have fewer of those arcs available, and layers with only smaller-capacity arcs don't need to be considered. Any set of k most-vital arcs must come from the same layer (because removing arcs from more than one layer at a time yields a resulting maximum flow that is equivalent to removing fewer total arcs from one of those two layers alone), and choosing arcs from any layer other than $n - k$ will yield a higher resulting maximum flow. Figure 9 illustrates the case $n = 6$, where each layer has capacity $25 = n^2$, and the optimal one-, two-, three-, and four-arc interdictions reduce the maximum flow by 9, 16, 21, and 124 units, respectively, yielding optimal resulting flows 16, 9, 4, and 1.

Any prioritized list that has a chance of being optimal must only involve arcs from a single layer, say, k , (because if not, then for at least one k' the attack for k' arcs will be no more damaging than the attack with $k' - 1$ arcs). That prioritized list will feature the largest arcs first, followed by the smallest, and will have a linear decrease in capacity of size $(n + k)$ units of flow per arc until $n - k$ arcs are chosen, at which point

SOMETIMES THERE IS NO “MOST-VITAL” ARC: ASSESSING AND IMPROVING THE OPERATIONAL RESILIENCE OF SYSTEMS




Figure 9. Values for optimal (diamonds) and prioritized (squares) interdictions from zero to 6 attacks for an example on $n = 6$ nodes, where the prioritized list is taken from layer 3. (The three largest arcs in layer 3 have capacity $3 + 6 = 9$, and the three smaller arcs have capacity 3.)

it will be the optimal $(n - k)$ -arc interdiction, followed by a linear decrease (by k units of flow per arc) until the capacity is zero. This yields a piecewise linear approximation to the optimal sequence of interdiction, agreeing with the optimal result only at zero, $n - k$, and n interdiction. Figure 9 illustrates the resulting maximum-flow capacity for an example with $n = 6$, for

the optimal interdiction of each size and the interdiction resulting from the prioritized list that would be built from layer 3.

Because any prioritized list can only agree with the optimal attacks for exactly three numbers of attacks, a prioritized list cannot be optimal for the corresponding counterexample for any n greater than three.