12 Bounding Network Interdiction Vulnerability Through Cutset Identification

Timothy C. Matisziw[1], Alan T. Murray[1,2], and Tony H. Grubesic[3]

[1] Center for Urban and Regional Analysis, The Ohio State University, USA; Email: matisziw.1@osu.edu
[2] Department of Geography, The Ohio State University, USA; Email: murray.308@osu.edu
[3] Department of Geography, Indiana University, USA; Email: tgrubesi@indiana.edu

12.1 Introduction

Assessing the vulnerability of network infrastructure to disruptive events is recognized as an important component of network planning and analysis. Motivations for this type of research range from searching for the most effective/ efficient means of disrupting a network (e.g., preventing drug trafficking – see Wood 1993) to assessing possible threats to critical network infrastructures so that adequate protective measures can be devised to limit potential disruption (see Wu 1992). In such analysis, the disruptive activity being examined, whether due to natural disaster, accident, or sabotage, can be generically referred to as network interdiction.

Traditionally, approaches for modeling network interdiction have focused on identifying nodes or linkages most critical to some interpretation of system performance. For instance, increasing the cost associated with routing flow between an origin-destination (O-D) pair is a common goal. Given the objective of increasing transportation costs, the impact of total or partial interdiction of linkages/ nodes can be considered as either: 1) decreasing network capacity, preventing flow or forcing it over more costly alternate paths; or, 2) increasing the cost associated with minimal cost paths. Both aspects of interdiction rely on negatively affecting network connectivity in some way. A classic network analysis approach to impacting connectivity between an O-D pair is through the identification of a cutset, or a set of linkages whose removal prevents O-D flow. Provided that interdiction efforts are limited by available resources, it is reasonable to focus on components of the smallest cutset possible (Wood 1993). It has been well established that solution of the maximum-flow model corresponds to a minimum capacity cut; hence, it is no surprise that this relationship has been exploited in the formulation of many interdiction models (Wollmer 1964; McMasters and Mustin 1970; Ghare et al. 1971; Corley and Chang 1974; Ratliff et al. 1975; Cunningham 1985; Phillips 1993; Wood 1993; Burch et al. 2003).

Models based upon a maximum-flow model generally seek to apply limited interdiction resources to minimize the network's capacity to move flow between origins and destinations. To achieve this goal, minimal cutsets can be identified for an O-D pair(s). No other cutset can be contained within a minimal cutset. A minimum capacity cutset then is a cutset of the smallest total weight (however defined). The usefulness of the maximum flow-minimum cut theorem is that the total capacity of a minimum cutset corresponds to the maximum amount of flow capable of moving between an O-D on the network (Ford and Fulkerson 1962; Colbourn 1987; Evans and Minieka 1992). Once minimum capacity cuts are found, linkages in these cuts are likely candidates for attack. The task then becomes determining which component linkages would be interdicted under a budgetary scenario. In this type of model, lower flow capacity remaining in the network indicates a more effective interdiction plan. An algorithmic approach to this problem is presented by Phillips (1993), while Wood (1993) implements this basic idea as an integer program. Though a minimum cutset may indeed be effective for interdiction in certain circumstances, it has been suggested that solution to some problems may require assessment of other minimal cutsets. For instance, if multiple interdiction objectives exist, a minimum capacity cut for each O-D may not necessarily be the most effective option (Boyle 1998; Balcioglu and Wood 2003).

Interdiction of network capacity is indeed an important consideration in assessing a network's vulnerability to interdiction; however, other criteria are also of interest. For instance, how actual origin-destination flow activity may be impacted by interdiction efforts is of obvious concern when addressing network survivability. Discussion on this topic can be found in Wu (1992) and Doyle et al. (2005). More recent analysis of this issue is found in Myung et al. (2004), Matisziw et al. (2006), Grubesic et al. (2006), and Murray et al. (2007). Another fundamental measure of attack vulnerability, therefore, is how network connectivity is impacted by an intentional disruption (see for instance, Holme et al. 2002; Grubesic et al. 2003). The argument is that given an attack on network facilities, higher potential connectivity loss equates to a more vulnerable network infrastructure. Furthermore, assessment of connectivity underlies the notions of network capacity and flow; hence, interdiction of connectivity is a valid concern when safeguarding network operation. As is clear from the previous discussion, network connectivity is directly related to the concept of minimal cutsets. Obviously, if all elements of a minimal O-D cutset are removed, then connectivity cannot be preserved. From an interdiction standpoint, just as it makes sense to target a minimum capacity cutset, minimum cardinality cutsets are also of interest because they reflect a scenario where limited resources are expended to cause the greatest damage possible. Colbourn (1987) describes one way of deriving minimum cardinality cutsets.

Regardless of the vulnerability measure(s) of concern (e.g., connectivity, capacity, flow), it is vital to understand the outcomes of potential interdiction to better support planning and management of network risk. One way of reducing risk is through the identification of the most disruptive interdiction schemes (those causing maximal damage) to establish an upper bound on vulnerability. If these worst-case scenarios can be identified, then administrators and managers can better plan for protection against threats and system improvement to minimize risks. In fact, many models developed for identifying optimal interdiction plans have their roots in network vulnerability assessment. For instance, the modeling efforts of Wollmer (1964), Corley and Chang (1974), Ratliff et al. (1975); Corley and Sha (1982), Ball et al. (1989), Malik et al. (1989), Church et al. (2004), and Murray-Tuite and Mahmassani (2004) all deal with finding infrastructure components of greatest importance to network operation, or rather the most vital links/ nodes.

However, complete focus on mitigating worst-case damage may not be entirely warranted as many near-optimal interdiction plans may also exist (Grubesic et al. 2006; Matisziw et al. 2006). Evaluation of the range of possible interdiction outcomes is undoubtedly beneficial in this regard, especially if multiple objectives are involved (see Boyle 1998). Hence, aside from an upper (worst-case) performance bound on interdiction severity, establishment of a lower bound (best-case) is also important to guide planning efforts. A higher lower bound may be more indicative of greater interdiction tolerance, as an example. Valid upper and lower bounds can also benefit simulations geared at generating a representative range of potential interdiction outcomes (see Matisziw et al. 2006).

To address the generation of bounds on interdiction of network flows, the flow interdiction model (FIM) has been recently proposed by Murray et al. (2007). The FIM permits assessment of maximally destructive node-based interdiction efforts on network operation. In other words, the FIM can produce an upper bound on the amount of network activity that may be lost due to a node-based disruption. Interdiction impacts for multiple origins and destinations are easily considered in this modeling framework. As suggested by Murray et al. (2007), the FIM enables either maximization or minimization of flow disruption to be evaluated. This is possible because O-D paths are explicitly tracked.

12.2 Modeling Linkage-based Interdiction

The focus of this chapter is the development of a model capable of producing upper and lower bounds on the loss of connectivity that may result from interdiction efforts aimed at network linkages. In other words, the goal is to identify a cutset of cardinality p that either minimizes or maximizes connectivity of origin and destination pairs. Given an uncapacitated network and the following notation, the p-cutset problem (PCUP) can be formulated:

j = index of linkages, entire set denoted J
k = index of paths, entire set denoted K
o = index of origins, entire set denoted O
d = index of destinations, entire set denoted D
N_{od} = set of paths providing o-d flow
p = number of linkages interdicted
Φ_k = set of linkages along path k

$$X_j = \begin{cases} 1 \text{ if linkage } j \text{ is interdicted} \\ 0 \text{ otherwise} \end{cases}$$

$$Y_k = \begin{cases} 1 \text{ if path } k \text{ remains uneffected by interdiction} \\ 0 \text{ otherwise} \end{cases}$$

$$Z_{od} = \begin{cases} 1 \text{ if no flow possible between } o \text{ - d} \\ 0 \text{ otherwise} \end{cases}$$

p-Cutset Problem (*P*CUP)

$$\textit{Minimize/Maximize} \quad \sum_o \sum_d Z_{od} \tag{1}$$

Subject to:

$$\sum_{k \in N_{od}} Y_k + Z_{od} \geq 1 \qquad \forall o,d \tag{2}$$

$$Z_{od} \leq (1 - Y_k) \qquad \forall\, k \in N_{od} \tag{3}$$

$$Y_k \geq 1 - \sum_{j \in \Phi_k} X_j \qquad \forall k \tag{4}$$

$$Y_k \leq (1 - X_j) \qquad \forall k, j \in \Phi_k \tag{5}$$

$$\sum_j X_j = p \tag{6}$$

$$X_j = \{0,1\} \qquad \forall j \tag{7}$$

$$Y_k = \{0,1\} \qquad \forall k$$

$$Z_{od} = \{0,1\} \qquad \forall o,d$$

Objective (1) is to either minimize or maximize O-D connectivity loss in a network. Constraints (2)-(3) track O-D path availability. Constraints (4)-(5) account for whether a given path is available given the loss of links. The number of linkages to be interdicted is stipulated in Constraint (6). Integer restrictions are specified in Constraints (7).

The PCUP formulation is similar to the FIM detailed in Murray et al. (2007). There are two fundamental differences with the PCUP, however. First, interdiction is considered only for arcs. Second, connectivity is addressed in the PCUP rather than flow.

The PCUP is beneficial in that it permits both minimization and maximization of (1). This is a convenient property since reformulation is not necessary given either goal. One key assumption of the model is that all paths permitting movement or flow between an origin and destination are accounted for. This is necessary for ensuring that a minimal or maximal cutset is identified. That is, here it is assumed that if *any* path connecting an O-D pair exists, then interaction between the two nodes is possible. Otherwise, if *no* path is available, then interaction between the pair cannot occur. Though use of a subset of O-D paths (e.g., *k*-shortest, arc/node disjoint) can reduce problem size, there is no guarantee that an identified cutset is optimal if all O-D paths are not accounted for. The PCUP deals explicitly with total interdiction of linkages (e.g., linkage is either available or is completely disabled) and partial disruption of a linkage is not possible. Worth noting as well is that a special case of the PCUP is the approach proposed in Myung et al. (2004) capable of addressing the maximization version of the problem. However, in their paper a minimization version of the model is not provided and cannot be obtained via a straightforward extension of their formulation. Furthermore, Myung et al. (2004) propose a heuristically derived bound and only consider a subset of possible O-D paths. The PCUP is an integer program and as such can be solved directly using a commercial optimization package. Here ILOG's CPLEX 6.6 mixed integer optimizer was utilized for solving problem instances. An issue that may arise though is that due to the number of constraints and integer decision variables, achieving optimality may be a computationally demanding task. Murray et al. (2007) discuss some ways in that these issues may be resolved. For example, integer requirements on Y_k and Z_{od} can be relaxed and some constraints can be consolidated (e.g., (5)). Additionally, some constraints could be eliminated from the general model depending on the objective orientation.

12.3 Application of the PCUP

Analysis of *p*-cutsets is conducted on the Abilene Internet2 backbone. The Abilene backbone is a high capacity fiber-optic Internet network connecting member universities within the U.S. (Abilene 2005). The backbone itself consists of 11 routers (nodes) connected by 14 linkages as shown in Figure 12.1.

Here, the PCUP is used to identify those cutsets capable of causing minimal and maximal damage to the network. All nodes in the Abilene network are both origins and destinations of flow and interact with each other. Given this, the network contains 121 interacting O-D pairs. In this network, intra-nodal interaction is present, meaning that flow can move into and out of the same node. Since nodes are not targeted for removal, only 110 O-D pairs (inter-nodal interactions) can potentially be disrupted given link-based interdiction. The O-D paths were obtained by enumerating all simple (loopless) paths for each O-D pair. 896 O-D exist, requiring approximately 2 seconds of computational time. Both the maximization and minimization cases are examined here for a range of interdiction scenarios.

Table 12.1 and Figure 12.2 illustrate results maximizing O-D connectivity loss. Since every node in the Abilene backbone is directly connected to at least two other nodes (a 2-degree node), the interdiction of a single linkage can not disconnect any O-D pair. However, when two linkages are rendered inoperative, more than half (60) of the O-D pairs lose connectivity. For example, the PCUP identifies the Kansas City-Indianapolis and

Houston-Atlanta linkages in the 2-cutset (p=2) causing the greatest impact. This cutset essentially partitions the network such that the number of nodes in each half is as balanced as possible, thereby maximizing disruption. For p=3, the PCUP identifies the Sunnyvale-Los Angeles, Denver-Kansas City, and Los Angeles-Houston linkages as a 3-cutset causing maximum connectivity loss. Los Angeles is consequently disconnected from the net-

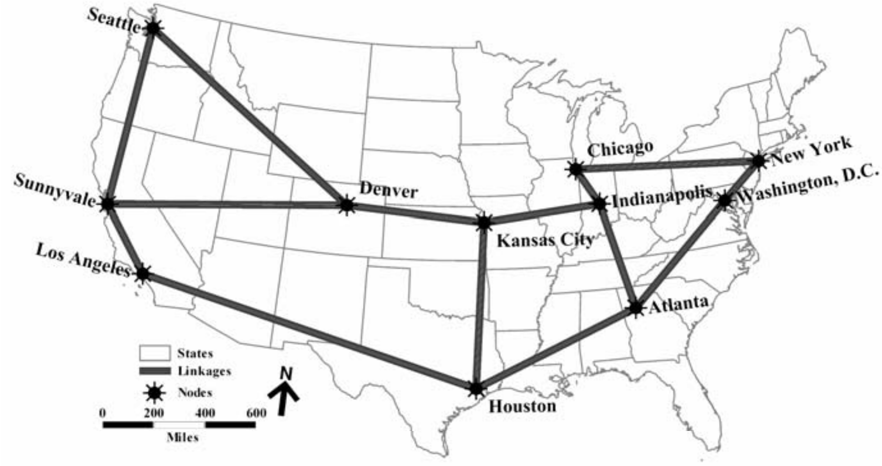

Fig. 12.1 Abilene Internet2 network backbone

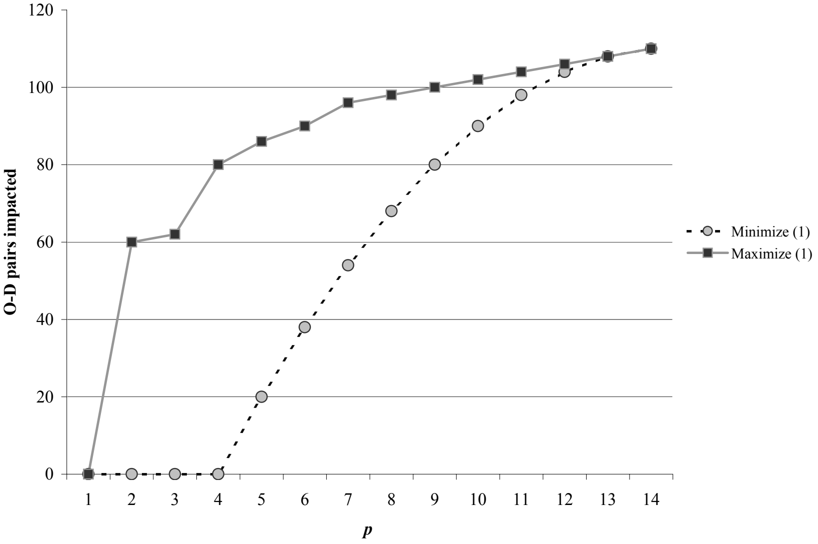

Fig. 12.2 Connectivity impact for minimal and maximal p-cutsets

work in this instance (see Table 12.1). Figure 12.3 shows the maximum impact of a 7 linkage failure/attack on the backbone. This particular interdiction plan fragments the network into 5 components, disconnecting all but 14 O-D pairs. Given a linkage-based interdiction plan, the maximum number of O-D pairs that can be interdicted is 110 since intra-nodal interaction cannot be impacted by a linkage-based attack. Thus, some 87% of O-D flow interactions are impacted in this case.

Fig. 12.3 Maximum 7-cutset (*p*=7)

Minimization of connectivity loss produces very different outcomes. Analysis for the minimization case of the PCUP is presented in Table 12.2 and Figure 12.2. The results presented in Table 12.2 illustrate the best-case situations for system performance given the occurrence of each interdiction scenario. In these cases, the model tries to preserve connectivity between the O-D pairs to the greatest extent possible in a rather intuitive manner. Given a set of nodes V in a network, it is well-known that minimum network connection occurs when $|J| = (|V|-1)$. In the case of the Abilene backbone there are 11 nodes, so a minimum of 10 linkages are needed to maintain connectivity. Since the backbone is composed of 14 linkages, up to four can be removed without causing connectivity loss. This is exactly the result found using the PCUP. After enough linkages are eliminated to reduce the network to a spanning tree, then reduction of each additional linkage disconnects exactly one node from the network, retaining a minimally connected network between the remaining nodes. That is, for values of $p > |J|-(|V|-1)$ in an undirected network where all nodes interact with each other, the maximum connectivity remaining in the network

Table 12.1 Maximizing connectivity loss

p	**OB**	**IT**	**BR**	**TM***	**Linkages Interdicted**
1	0	1301	10	9.890	CH-NY
2	60	348	0	1.391	IN-KC;AT-HO
3	62	1649	28	11.156	DE-KC;SV-LA;LA-HO
4	80	285	0	1.046	DE-KC;CH-IN;LA-HO;WA-AT
5	86	181	2	1.938	DE-KC;CH-IN;SV-LA;LA-HO;WA-AT
6	90	240	4	2.406	DE-KC;CH-IN;SV-LA;LA-HO;WA-AT;CH-NY
7	96	163	8	2.688	DE-KC;CH-IN;SV-LA;KC-HO;AT-HO;IN-AT;NY-WA
8	98	240	15	2.438	DE-KC;CH-IN;SV-LA;KC-HO;AT-HO;IN-AT;NY-WA;CH-NY
9	100	48	0	0.453	DE-KC;CH-IN;SE-SV;DE-SV;LA-HO;KC-HO;IN-AT;WA-AT;NY-WA
10	102	49	0	0.469	DE-KC;CH-IN;SE-SV;DE-SV;LA-HO;KC-HO;AT-HO;IN-AT;WA-AT;NY-WA
11	104	42	0	0.469	DE-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO;IN-AT;WA-AT;NY-WA
12	106	42	0	0.562	DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO;IN-AT;WA-AT;NY-WA
13	108	41	0	0.453	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO;IN-AT;WA-AT;NY-WA
14	110	0	0	0.234	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO;IN-AT;WA-AT;NY-WA;CH-NY

OB Objective, *IT* Iterations, *BR* Branches, *TM* Time, *AT* Atlanta, *CH* Chicago, *DE* Denver, *HO* Houston, *IN* Indianapolis, *KC* Kansas City, *LA* Los Angeles, *NY* New York, *SE* Seattle, *SV* Sunnyvale, *WA* Washington, D.C.
*Solution times in seconds for Pentium III parallel processor with 1.0 GB RAM

given the removal of p linkages can be determined as follows: 1) compute the number of nodes that become separated from the network $V_S = p\text{-}(|J|\text{-}(|V|\text{-}1))$, 2) compute the number of nodes retained in the network $V_R\text{=}(|V|\text{-}V_S)$, 3) since the network design preserving connectivity is a tree, it is known that $V_R(V_R\text{-}1)$ node pairs will remain connected, and 4) the difference between original network connectivity and that remaining after p linkages are removed gives the connectivity lost (PCUP's objective). Figure 12.4 illustrates the minimum connectivity loss resulting from a 7-cutset interdiction, in contrast to the maximum scenario shown in Figure 12.3. Note that through a minimum network connection, interaction between 8 nodes (56 O-D pairs) can be preserved. While identifying a spanning tree may be an alternative and attractive way of solving this minimization problem, it is unclear whether such a technique will always result in an optimal solution if: 1) all nodes do not interact with each other, and/or 2) actual O-D flow activity is incorporated within the model.

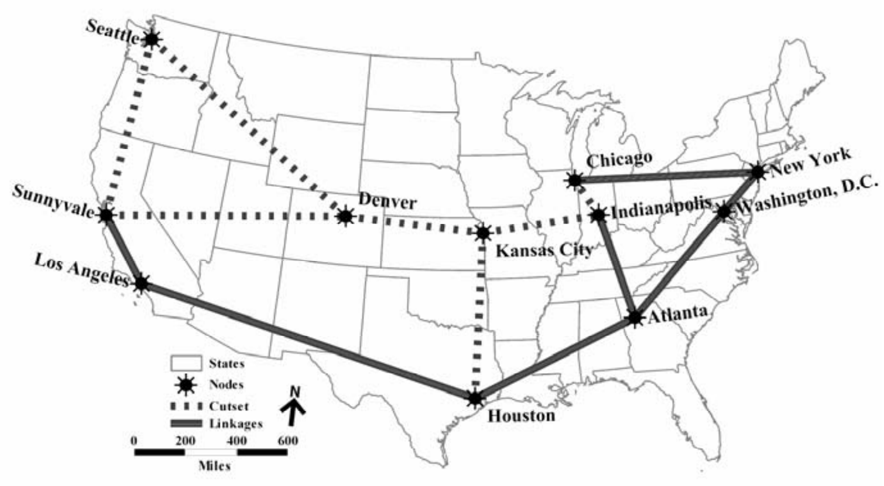

Fig. 12.4 Minimum 7-cutset (p=7)

In the existing literature, it is commonly assumed that individual characteristics of network nodes or linkages (e.g., degree) can be used as proxies to infer the importance of a facility to network operation (Holme et al. 2002, for example). However, the result of an interdiction is in fact strongly dependant upon the spatial structure of the network, as can be seen in the model results. As an example, Table 12.2 shows that a 2-cutset resulting in minimum O-D disruption involves the Denver-Kansas City and Atlanta-Houston linkages that are both rooted by different nodes of degree 3. In this case, no O-D connectivity is lost. On the other hand, the

Table 12.2 Minimizing connectivity loss

p	OB	IT	BR	TM*	Linkages Interdicted
1	0	915	0	4.375	DE-KC
2	0	667	1	23.344	DE-KC;AT-HO
3	0	433	1	22.141	LA-HO;AT-HO;IN-AT
4	0	356	1	12.547	SE-DE;DE-SV;IN-AT;WA-AT
5	20	96118	1246	151.953	DE-KC;IN-KC;CH-IN;DE-SV;CH-NY
6	38	110907	1891	195.063	SE-DE;DE-KC;IN-KC;CH-IN;DE-SV;IN-AT
7	54	88637	1649	179.047	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;KC-HO
8	68	60221	1346	140.125	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO
9	80	30563	602	85.609	IN-KC;CH-IN;DE-SV;KC-HO;AT-HO;IN-AT;WA-AT;NY-WA;CH-NY
10	90	17946	325	66.250	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO
11	98	10537	157	53.125	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO;IN-AT
12	104	5858	91	58.235	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO;IN-AT;WA-AT
13	108	2694	22	41.219	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO;IN-AT;NY-WA;CH-NY
14	110	0	0	0.235	SE-DE;DE-KC;IN-KC;CH-IN;SE-SV;DE-SV;SV-LA;LA-HO;KC-HO;AT-HO;IN-AT;WA-AT;NY-WA;CH-NY

OB Objective, *IT* Iterations, *BR* Branches, *TM* Time, *AT* Atlanta, *CH* Chicago, *DE* Denver, *HO* Houston, *IN* Indianapolis, *KC* Kansas City, *LA* Los Angeles, *NY* New York, *SE* Seattle, *SV* Sunnyvale, *WA* Washington, D.C.
*Solution times in seconds for Pentium III parallel processor with 1.0 GB RAM

2-cutset identified in Table 12.3 (discussed above) also involves different nodes of degree 3. However, their selection forms a spatial partition, maximizing O-D connectivity loss. While both of these 2-cutsets have similar physical characteristics, their impact on network functionality (if interdicted) is not at all similar. In fact, these two cutsets form upper and lower bounds on possible connectivity loss due to 2-link interdiction and serve to illustrate the range of interdiction outcomes possible. Measures based on proxies for connectivity are not likely to be good approximations for these bounds and could obscure the true extent of network vulnerability. Nonetheless, many other feasible interdiction outcomes undoubtedly occur for each interdiction scenario (e.g., p=2) and may include near-optimal solutions or alternate-optima. Hence, from the perspective of managing network vulnerabilities, there is still a clear benefit in characterizing the range of possible outcomes between the upper and lower bounds through simulation as discussed in Matisziw et al. (2006).

12.4 Discussion and Conclusion

This chapter has focused on identifying minimal or maximal p-cutsets for a system of origins and destinations. The goal is to obtain the set of p cardinality cuts, or linkages, capable of maximizing or minimizing network connectivity loss. This distinction is important in that other models have focused primarily on the interdiction of capacity, not connectivity or other measures of network vulnerability. Furthermore, models that have approached connectivity have typically done so using proxies for connectivity (e.g., nodal degree, betweeness, etc.) and have not modeled it exactly as is done here. The motivation for this problem follows directly from the need to assess a network's vulnerability to interdiction. Effective planning and management of network risks must consider the range of interdiction scenarios possible if appropriate mitigation measures are to be devised.

Recent events emphasize the importance of such analysis. For instance, single-link failures are a common occurrence in the operation of many networks and hence a common consideration in network design (Wu et al. 1988). Although many networks are resilient to single-link attack/ failure, additional, simultaneous disruptions can be very problematic, leading to wide-spread service outages. A recent example of this type of service outage is that caused by a 2-cut in the Sprint Nextel fiber optic network in the southwestern U.S. (C|net 2006; CNN 2006).

In order to identify the upper and lower bounds on post-interdiction network connectivity, the p-cutset problem (PCUP) was proposed. The

PCUP is an extension of the flow interdiction model of Murray et al. (2007) aimed at explicitly accounting for linkage-based interdiction. Unlike other interdiction models, the PCUP's structure permits both maximization and minimization of network connectivity loss. This is accomplished by enumeration of O-D paths, permitting identification of cutsets of a stipulated cardinality that disconnect or preserve the greatest number of O-D relationships. Through application to a real world network, the PCUP is shown to be effective for identifying bounds on potential interdiction scenarios.

Acknowledgements

Project funding for Matisziw and Murray is provided through the Center for Urban and Regional Analysis at The Ohio State University.

References

Abilene. 2005. http://abilene.internet2.edu.

Balcioglu, A. and R.K. Wood. 2003. Enumerating near-min s-t cuts. In *Network Interdiction and Stochastic Integer Programming*. Edited by D.L. Woodruff. Boston: Kluwer Academic Publishers, 51-69.

Ball, M.O., B.L. Golden, and R.V. Vohra. 1989. Finding the most vital arcs in a network. *Operations Research Letters*. 8(2), 73-76.

Burch, C., R. Carr, S. Krumke, M. Marathe, C. Phillips, and E. Sundberg. 2003. A decomposition-based approximation for network inhibition. In *Network Interdiction and Stochastic Integer Programming*. Edited by D.L. Woodruff. Boston: Kluwer Academic Publishers, 51-69.

Boyle, M.R. 1998. *Partial-Enumeration for Planar Network Interdiction Problems*. M.S. Thesis: Naval Postgraduate School.

C|net News.com. 2006. Sprint Nextel suffers service outage. http://www.news.com. Monday, Jan. 9.

Church, R.L., M.P. Scaparra, and R.S. Middleton. 2004. Identifying critical infrastructure: the median and covering facility interdiction problems. *Annals of the Association of American Geographers*. 94(3), 491-502.

Colbourn, C.J. 1987. *The Combinatorics of Network Reliability*. New York: Oxford.

Corley, H.W. and H. Chang. 1974. Finding the *n* most vital nodes in a flow network. *Management Science*. 21, 362-364.

Corley, H.W. and D.Y. Sha. 1982. Most vital links and nodes in weighted networks. *Operations Research Letters*. 1(4), 157-160.

CNN.com. 2006. Cut cable quiets Sprint service in West. http://www.cnn.com. Jan. 9.

Cunningham, W.H. 1985. Optimal attack and reinforcement of a network. *Journal of the Association for Computing Machinery*. 32(3), 549-561.

Doyle, J.C., D.L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger. 2005. The "Robust Yet Fragile" Nature of the Internet. *Proceedings of the National Academy of Sciences of the United States of America*. 102(4), 14497-14502.

Evans, J. and E. Minieka. 1992. *Optimization Algorithms for Networks and Graphs*. 2nd Edition. New York: Marcel Dekker, Inc.

Ford, L.R. and D.R. Fulkerson. 1962. *Flows in Networks*. Princeton Press.

Ghare, P.M., D.C. Montgomery, and W.C. Turner. 1971. Optimal interdiction policy for a flow network. *Naval Research Logistics Quarterly*. 18, 37-45.

Grubesic, T.H., M.E. O'Kelly, and A.T. Murray. 2003. A geographic perspective on commercial internet survivability. *Telematics and Informatics*. 20, 51-69.

Grubesic, T.H., T.C. Matisziw, and A.T. Murray. 2006. Targeted attacks and survivability in critical network infrastructure. *Submitted for publication*.

Holme, P., B.J. Kim, C.N. Yoon, and S.K. Han. 2002. Attack vulnerability of complex networks. *Physical Review E*. 65, 056109.

Malik, K., A.K. Mittal, and S.K. Gupta. 1989. The most vital arcs in the shortest path problem. *Operations Research Letters*. 8, 223-227.

Matisziw, T.C., A.T. Murray, and T.H. Grubesic. 2006. Exploring the vulnerability of network infrastructure to interdiction. *Submitted for review*.

McMasters, A.W. and T.M. Mustin. 1970. Optimal interdiction of a supply network. *Naval Research Logistics Quarterly*. 17, 261-268.

Murray, A.T., T.C. Matisziw, and T.H. Grubesic. 2007. Critical network infrastructure analysis: interdiction and system flow. *Journal of Geographical Systems*, 39.

Murray-Tuite, P.M. and H.S. Mahmassani. 2004. Methodology for determining vulnerable links in a transportation network. *Transportation Research Record*. 1882, 88-96.

Myung, Y-S. and H. Kim. 2004. A cutting plane algorithm for computing *k*-edge survivability of a network. *European Journal of Operational Research*. 156, 579-589.

Phillips, C.A. 1993. The network inhibition problem. *Proceedings of the Annual Association for Computing Machinery STOC*. California, May.

Ratliff, H.D., G.T. Sicilia, and S.H. Lubore. 1975. Finding the *n* most vital links in flow networks. *Management Science*. 21(5), 531-539.

Wollmer, R. 1964. Removing arcs from a network. *Operations Research*. 12, 934-40.

Wood, R. K. 1993. Deterministic network interdiction. *Mathematical Computer Modelling*. 17(2), 1-18.

Wu, T-H. 1992. *Fiber Network Service Survivability*. Boston: Artech House.

Wu, T-H., D.J. Kolar, and R.H. Cardwell. 1988. Survivable network architectures for broad-band fiber optic networks: model and performance comparison. *Journal of Lightwave Technology*. 6(11), 1698-1709.