

SIMULATION OF IMPERFECT INFORMATION IN VULNERABILITY MODELING FOR INFRASTRUCTURE FACILITIES

Dean A. Jones

Sandia National Laboratories
PO Box 5800 MS 1138
Albuquerque, NM 87185-1138, U.S.A

Mark A. Turnquist

Linda K. Nozick
School of Civil & Environmental Engineering
309 Hollister Hall
Cornell University
Ithaca, NY 14850, U.S.A.

ABSTRACT

A model of malicious intrusions in infrastructure facilities is developed that uses a network representation of the system structure together with Markov models of intruder progress and strategy. Simulation is used to analyze varying levels of imperfect information on the part of the intruders in planning their attacks. This provides an explicit mechanism to estimate the probability of successful breaches of physical security, and to evaluate potential means to reduce that probability.

1 INTRODUCTION

There is widespread interest in protection of critical infrastructure from malicious attack. The attacks might be either physical intrusions (e.g., to steal vital material, plant a bomb, etc.) or cyber intrusions (e.g., to disrupt information systems, steal data, etc.) and the attackers may be international terrorists, home-grown hackers, or ordinary criminals. In 1997, the report of the U.S. President's Commission on Critical Infrastructure Protection (PCCIP) identified eight critical infrastructures "whose incapacity or destruction would have a debilitating impact on our defense and economic security" (PCCIP, 1997). These eight are: telecommunications, electric power systems, natural gas and oil, banking and finance, transportation, water supply systems, government services and emergency services.

In this analysis, we focus primarily on transportation facilities, but the approach we suggest could also be used in other infrastructure contexts. For example, a similar type of analysis has been applied to information systems by Carlson, et al. (2004). The objective of the analysis presented here is to provide guidance to system owners and operators regarding effective ways to reduce vulnerabilities of specific facilities. To accomplish this, we develop a Markov Decision model of how an intruder might try to penetrate

the various barriers designed to protect the facility. This intruder model provides the basis for consideration of possible strategies to reduce the probability of a successful attack on the facility.

Our primary attention in this paper is on how varying levels of information about the infrastructure system affects the strategies of potential intruders, how the overall probability of intruder success is affected by their level of information, and what implications this has for effective defense of the system against intrusion.

We represent the system of interest as a network of nodes and arcs. Nodes represent barriers that an intruder must penetrate, and arcs represent movements between barriers that an intruder can make within the system. Several previous authors have used graph-based methods to represent attackers or defenders in security analyses. Phillips and Swiler (1998) introduced the concept of an "attack graph" to represent sets of system states and paths for an attacker to pursue an objective in disrupting an information system. Several subsequent papers (e.g., Swiler, et al., 2001; Jha, et al., 2002; Sheyner, et al., 2002) have extended these initial ideas.

The adversaries first must penetrate entry points to the system, and if an attempted penetration at a particular entry node is successful, they can traverse edges from the successfully breached node to other nodes in the network that are connected to the one breached. Traversing an edge entails a risk of detection. The adversary is assumed to make the decision that maximizes what he/she perceives to be the probability of successful attack. If this perception is inaccurate, the strategy pursued may not be optimal and the overall probability of success is reduced.

We can think of this analysis as having three layers. At the bottom layer, the physical characteristics of individual barriers are translated into summary probabilities of detection, success, etc., for use in the middle level model. This mid-

the layer is a Markov Decision Process (MDP) model that represents the optimization of the intruder’s strategy, given the perceived values of detection probabilities, etc. The perceptions may have different levels of accuracy. At this level, we use both simulation and optimization tools. Simulation is used to represent varying perceptions of the system parameters by intruders, and optimization is used to create strategies on the part of intruders, given those perceptions. At the top layer of analysis, the system operator (or defender) examines the probabilities of success on the part of potential intruders and the paths that they are likely to follow through the network, and makes changes to reduce the system vulnerabilities. Those changes may be designed either to reduce the real success rates of intruders in penetrating system barriers, or to decrease the accuracy of the information available to the intruders so that their attempts to optimize strategies are less effective.

At the lowest layer, we use Hidden Markov Models (HMM) to represent an intruder’s actions at a single node (barrier) in a system and the associated “signals” those actions provide that can lead to detection. Then we develop an aggregated representation of that single-node model for inclusion in an MDP model of intruder strategy within a network representation of the entire system at the middle layer. These parts of the analysis are described in detail by Carlson, et al. (2004) and Jones, et al. (2005). In the interests of space, they will not be included here, so that we may focus this paper on the interaction of simulation and optimization analysis at the middle layer.

2 MARKOV DECISION MODEL OF INTRUDER STRATEGY

At the system level, we represent a network of barriers and potential movements as shown in the example in Figure 1, representing a simplified hypothetical attempt by an intruder to place a delayed-action (e.g., altitude detonated) explosive device on an aircraft sitting at a gate in an airport terminal.

The intruder must first gain access to the apron area of the terminal. We postulate that this can occur either by gaining illicit access through the employee gate (e.g., by stealing an employee ID and using it to enter the area), or by entering in a service vehicle at a gate (e.g., in a catering truck). If the intruder is successful in getting access to the area, he/she must then impersonate a legitimate worker in the aircraft gate area – either an airline employee or a service contractor. The “cross-over” arcs between “entry” and “impersonation” in Figure 1 indicate that even if the intruder gained access to the apron area using an employee ID, he/she may switch ID’s and impersonate a service contractor within the area (or vice versa). This impersonation

must be successful for the period of time required to get from the entrance to the aircraft itself.

Approaching the aircraft carries a risk of detection, and the approachable areas on the aircraft if the intruder is impersonating an employee may be different from those that are approachable if he/she is impersonating a service contractor. For example, a person who appears to be an airline maintenance employee might not attract attention approaching the under-wing area around the landing gear, whereas a person who appears to be a catering contractor would. For purposes of this example, we consider three areas of the aircraft where an explosive device might be hidden – inside the wing around the landing gear, in the cargo hold, or in the catering supplies delivered to the galley.

If access to the aircraft is gained, the device must be placed without arousing suspicion. This is represented by the arcs connecting the aircraft area nodes to the exit node. Each of these arcs has a probability of detection.

Finally, if the intruder succeeds in gaining access to the aircraft and placing the device, he/she must exit without detection, and this represents the last barrier. Our modeling premise is that if the intruder is detected after placing the device, it will trigger a thorough search of the aircraft and the device will be discovered, so that the attempted attack will be foiled.

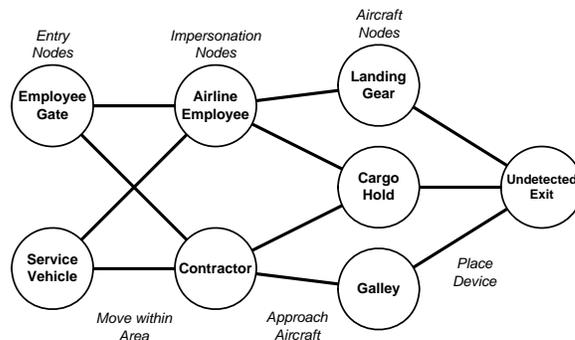


Figure 1: Illustrative network for analyzing an attempted placement of an explosive device on an aircraft

To generalize from this specific example, in a network representing some infrastructure facility or system, if the intruder is successful at breaching a particular barrier, he/she has choices about where to go next (which arc to cross). Crossing arc ij entails a probability of detection δ_{ij} , and this is represented in the transition matrix.

If the intruder is in state i and chooses action a_i , we denote the expected value of the future stream of rewards by $w(i,a_i)$. Each possible action a_i implies a change in the transition probabilities that govern the process. We denote

the elements of the transition matrix resulting from choosing action a_i as $P_{ij}(a_i)$. The MDP we define for this problem is positive bounded, and we can find the optimal policy through either policy iteration or linear programming (Puterman, 1994).

Table 1 summarizes the hypothetical node data used for the example analysis, and Table 2 shows the probabilities of detection used for the arcs in the example network. Note that we assume there is no retreat at the stage of exiting after placing the device – at that stage either the attack is successful or it is detected. Also note that the probability of detection on the arcs leading to the “impersonation” nodes is zero. This is because we are treating impersonation process (and time) as a barrier (node), so the probability of detection is lumped at the nodes, rather than on the arcs.

Table 1 : Example data for network nodes

Node #	Node Description (see Figure 1)	Expected Time For Attempted Breach (min)	Prob. Of Success	Prob. of Detection	Prob. of Retreat
1	Employee Gate	1	0.2	0.65	0.15
2	Service Gate	2	0.25	0.7	0.05
3	Impersonate Employee	10	0.2	0.6	0.2
4	Impersonate Contractor	15	0.4	0.5	0.1
5	Landing Gear	5	0.15	0.8	0.05
6	Cargo Hold	3	0.1	0.75	0.15
7	Galley	15	0.15	0.75	0.1
8	Undetected Exit	10	0.8	0.2	0

Table 2: Probability of detection for possible moves

Arc	Prob. of Detection
Empl. Gate – Impersonate Employee	0
Empl. Gate – Impersonate Contractor	0
Service Gate – Impersonate Empl.	0
Service Gate – Impersonate Contr.	0
Impersonate Empl. – Landing Gear	0.7
Impersonate Empl. – Cargo Hold	0.7
Impersonate Contr. – Cargo Hold	0.6
Impersonate Contr. – Galley	0.6
Landing Gear – Exit	0.4
Cargo Hold – Exit	0.2
Galley – Exit	0.3

If an intruder knew the structure of the network (Figure 1) and the values in Tables 1 and 2, we would consider him/her to be perfectly informed. Under this assumption, an optimal intrusion strategy (i.e., one that maximizes the probability of successful attack) can be constructed by solving the MDP. For the set of input data in Figure 1 and Tables 1 and 2, the solution for the optimal intruder strategy can be summarized as shown in Figure 2. To the left of each node is the probability of successful attack, given that the intruder is “arriving at” that barrier. To the right of each node is the probability of success, given that the intruder has successfully negotiated that barrier. There is only one value shown for the exit node (i.e., the “approaching” probability), because once that node is successfully negotiated, the attack has been a success, by definition.

The light colored arcs indicates the optimal path for an intruder (i.e., the path that maximizes the probability of success). This is the path of greatest vulnerability to the system. In our simple example, we would compute a probability of successful attack of 0.0034 for an intruder whose strategy is to gain entry to the apron area through the service vehicle gate, then impersonate a contractor (probably a catering service worker) to access the aircraft galley and place the device there before exiting.

The existence of this strategy does not mean that all intruders will always proceed in exactly the way indicated. It does mean that if an intruder were perfectly informed, this would be a strategy through which the probability of a successful attack could be maximized. In actuality, the probability of successful attack is likely to be less than this maximum value because intruders will have less-than-complete information and may not optimize their strategy. The solution to the MDP also provides useful information on the conditional probability of success for an attacker that reaches a certain point in the network, regardless of whether or not he/she followed the optimal strategy. For example, if an intruder succeeds in reaching the cargo hold of the aircraft (despite the fact that this is not an optimal strategy), the probability of a successful attack from that point on is 0.064.

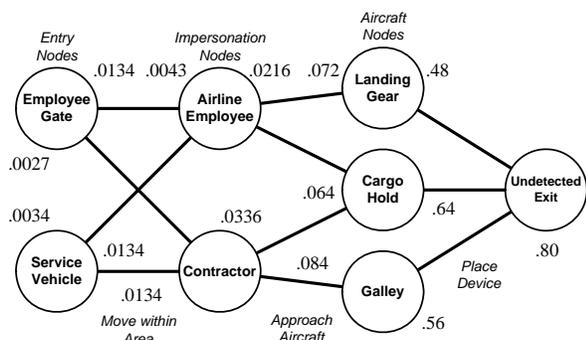


Figure 2: Summary of intruder strategy and probability of success under perfect information

3 REPRESENTING IMPERFECT INFORMATION

One useful representation of imperfect information is to assume that a potential intruder does not know the values of the probabilities in Tables 1 and 2, but has perceptions of those probabilities that contain errors. An intruder with imperfect information will attempt to construct an optimal strategy, but because of errors in perception of detection probabilities, the strategy is likely to actually be suboptimal against the real probabilities. Simulation is an effective tool to explore the effects of imperfect information represented in this way.

Suppose that the perception of a given detection probability is represented as a beta random variable with parameters $a > 0$ and $b > 0$. The mean of such a random variable is

$$\frac{a}{a+b}, \text{ and the variance is } \frac{ab}{(a+b+1)(a+b)^2}.$$

If the intruder's perception of an unknown probability π is unbiased,

$$\frac{a}{a+b} = \pi, \text{ and we can express one of the parameters in terms of the other - e.g., } b = \frac{(1-\pi)a}{\pi}.$$

By varying a , we can change the variance (i.e., the level of uncertainty in the perception of π) and set b in terms of a to maintain the same expected value. A convenient way to create experiments is to set the coefficient of variation for the distribution and then solve for the values of a and b that will maintain the desired mean and achieve the required standard deviation. The coefficient of variation for the beta distribution is

$$\sqrt{\frac{b}{a(a+b+1)}}.$$

Alternatively, we can assume that the intruder's perception of the unknown probability may be biased. If we specify both the coefficient of variation in the distribution and the degree of bias ($\pi - \frac{a}{a+b}$), we can solve for values of a and b to satisfy those requirements.

For any setting of the values for the parameters a and b , we can sample from the perception distribution to simulate an intruder operating with some specified level of imperfect information. Of course, this concept extends to imperfect information with respect to any number of probability estimates. Replicating this simulated sampling leads to varying choices of paths through the network by the imperfectly informed intruder, each of which has a different probability of success. This allows construction of an estimated probability distribution for the likelihood of successful attack by an intruder operating at that level of imperfect information, as well as a probability distribution over possible paths through the network. The distribution of path choices allows us to reach some conclusions regarding the likelihood that an intruder will appear at certain points in the network.

4 ILLUSTRATIVE SIMULATION RESULTS

To illustrate these ideas, we will consider a series of experiments using the basic network from Figure 1, and compare the results to the perfect-information solution in Figure 2. As a first experiment we assume that the intruder's perception of the detection probabilities (at the nodes and along the arcs) is unbiased, but has a coefficient of variation of 0.1 for all non-zero probabilities (i.e., excluding the first four entries in Table 2).

As an example of the beta distribution parameter computations, consider the detection probability for the arc connecting "Impersonate Contractor" to "Cargo Hold" (the seventh row of Table 2). The true value for this probability is 0.6. To determine the a and b parameters of the beta dis-

tribution to represent imperfect information, we establish the two equations:

$$\frac{a}{a+b} = 0.6 \tag{1}$$

$$\sqrt{\frac{b}{a(a+b+1)}} = 0.1 \tag{2}$$

We then solve for a and b, leading to the values a = 39.4 and b = 26.27. This computation is repeated (for different underlying probabilities in equation 1) to produce a and b parameters for all the non-zero detection probabilities.

In each simulation experiment, the success probability for a given node or arc is adjusted to accommodate the sampled value of the detection probability. The retreat probabilities at the nodes are unchanged. This adjustment ensures that the required probabilities sum to 1.0.

Table 3 summarizes the results of 30 replications of the simulation. The path descriptors use the node numbering scheme from Table 1, and are listed in order of decreasing probability of success. The probabilities of use are rounded to two decimal places, and may not add exactly to 1.0. The path found in the perfect-information case (2-4-7-8) is one of the two most likely paths when the intruder has imperfect information, but approximately 63% of the time, the imperfectly informed intruder will choose a suboptimal path, even when the variability in the perceptions of detection probabilities (as measured by the coefficient of variation) is relatively small (0.1). The average probability of success for an intruder with this level of information is .00279, approximately 17% lower than for the perfect information case. This experiment indicates that even a little reduction in information about the system can have a significant effect on reducing the likelihood of a successful attack.

Table 3: Summary of results when probability estimates are unbiased and coefficient of variation is 0.1

Chosen Path	Probability of Use	Probability of Success
2-4-7-8	0.37	.00336
1-4-7-8	0.17	.00269
2-4-6-8	0.37	.00256
1-4-6-8	0.07	.00205
2-3-5-8	0.03	.00108

In addition to information on average probability of success, the path data and probabilities in Table 3 can be used to estimate the likelihood that an intruder will appear at a given point in the network, given the level of imperfect in-

formation hypothesized. This is done simply by summing probabilities for paths that include a given node or arc. For example, we might be particularly interested in the relative likelihoods of attempts to place explosives in the three different areas of the aircraft. In this case, we could use the results in Table 3 to conclude that the probabilities of an intruder attempting to use the landing gear (node 5), the cargo hold (node 6) and the galley (node 7) are .03, .44 and .54, respectively (again rounded to two decimal places).

Further insight into the effects of imperfect information can be obtained by increasing the level of uncertainty. A second experiment increased the coefficient of variation in the detection probability perceptions to 0.3. The perceptions are still considered to be unbiased. Table 4 summarizes the results, again based on 30 replications of the simulation.

Table 4: Summary of results when probability estimates are unbiased and coefficient of variation is 0.3

Chosen Path	Probability of Use	Probability of Success
2-4-7-8	0.33	.00336
1-4-7-8	0.23	.00269
2-4-6-8	0.27	.00256
1-4-6-8	0.07	.00205
2-3-5-8	0.03	.00108
1-3-5-8	0.07	.00086

Comparing Table 4 to Table 3, we see that the increase in uncertainty about the correct detection probabilities causes the optimal path to be chosen less frequently, and a very suboptimal path (1-3-5-8) appears in the list of possibilities. Overall, the average probability of success is .00266. This is a decrease from the case where the coefficient of variation is 0.1, but only about 5%. In this sample problem at least, a small amount of uncertainty in the perceived detection probabilities is important, but making that uncertainty much larger has relatively little effect on the expected probability of successful attack, as long as the perceptions are unbiased.

There is a somewhat more noticeable effect of the increase in uncertainty on the probabilities of the intruder attempting to use different parts of the aircraft. From the results in Table 4, we can compute estimates of the probability that the intruder would attempt to use the landing gear (node 5), the cargo hold (node 6) and the galley (node 7) as 0.1, 0.34, and 0.56, respectively. There is a noticeable shift in likelihood from the cargo hold to the landing gear for less well-informed intruders. This insight can be helpful to security forces.

To test the effects of biased perceptions, we have conducted a third simulation experiment. The coefficients of variation in the detection probability perceptions are set to 0.1, as in the first experiment, but we introduce a bias on two of the perceived probabilities – the detection probabilities associated with a contractor approaching the aircraft, either the cargo hold or the galley. In Table 2, the “true” values are indicated to be 0.6, but we assume that the intruder believes (on average) that the values are 0.9 for both probabilities. Intuitively, we expect that these misperceptions will tend to drive the intruder’s attack path away from paths that use those two arcs, and since one of the two arcs is part of the optimal path under perfect information, the net effect should be a reduction in success probability for the intruder.

Table 5 summarizes the results of the experiment, again based on 30 simulation replications. The overall average probability of success for an attack is reduced to .00144, a reduction of 48% from the value in experiment 1 (.00279), and a reduction of 57% from the original value based on perfect information. The misperception of detection probabilities on the two arcs makes it much less likely that the intruder will attempt to use those arcs (probability of 0.23 versus 0.97 in the first experiment). Attacks are much more likely to be focused on paths (and areas of the aircraft) where the real detection probability is higher, leading to much lower success probability for the intruder. In the results shown in Table 5, the probability of the intruder attempting to use the landing gear area is 0.37, as compared to 0.03 in experiment 1, and the probability of attempts through the galley has decreased from 0.54 to 0.1.

Table 5: Summary of results when probability estimates are biased on arcs 4-6 and 4-7

Chosen Path	Probability of Use	Probability of Success
2-4-7-8	0.10	.00336
2-4-6-8	0.13	.00256
2-3-5-8	0.33	.00108
2-3-6-8	0.33	.00096
1-3-5-8	0.03	.00086
1-3-6-8	0.07	.00077

The level of bias in the perceptions of the detection probabilities on arcs 4-6 and 4-7 used in this experiment is substantial, and smaller assumed biases would create less dramatic results. However, we have only introduced the bias on two arcs in the network. More widespread misperceptions would be likely in a larger system. This experiment does indicate that creating biased perceptions of detection probabilities among potential intruders can be very effective in reducing the likelihood of successful attacks by “steering” those attacks into areas where detection really is very likely.

There are several means through which a system operator might create such misperceptions. Implementing inexpensive, highly visible (though perhaps not really very effective) detection mechanisms might be one means. Supplying disinformation about real operations or procedures may be another, although this has obvious drawbacks as well.

5 EXTENSIONS

Several possible extensions to this analysis are possible. First, other aspects of imperfect intruder information could be included, such as imperfect knowledge about what barriers (nodes) and arcs exist in the system. This type of imperfect information can be incorporated into the general analysis framework described in this paper.

A second useful extension is to consider where improvements in security (i.e., increases in detection probability) would be most effective against several classes of potential intruders (i.e., intruders with differing levels of information about the system). The analysis of possible investments to improve security is a vital part of the overall approach we have outlined here, and this is an active area of current work.

A third useful extension is to create semi-Markov models for the processes of attempted penetration of barriers. This would allow more accurate representation of the uncertain time required to penetrate a given barrier, as well as offer the opportunity for time-dependent detection probabilities (i.e., the longer an intruder is present at a barrier, the more likely it becomes that he/she will be detected). This extension could improve the range of applicability of the model.

6 CONCLUSIONS

We have developed a model of intruder actions in attacking an infrastructure system based on a Markov Decision Process (MDP). Lower level models of intruder detection at barriers (nodes) of the system can be built as Hidden Markov Models, and the results of those lower level models can be aggregated for use in the MDP of intruder strategy for attacking the system. A key aspect of this analysis is representing imperfect information on the part of the intruders, and this paper focuses on that part of the analysis. Simulation is used as a tool to evaluate the effects of varying levels of imperfect information, sampling from distributions of detection probabilities and using those samples to construct distributions of intruder path choices through the network and overall success probability.

A small example problem illustrates that even relatively small amounts of uncertainty in the information the intruders have about the system can significantly affect the prob-

ability that they can mount a successful attack. If the uncertainty is combined with bias in the perceptions of some system parameters, the effect on the intruders is magnified. In the small example studied, biased perceptions of two key detection probabilities combined with small amounts of uncertainty in perceptions of all the detection probabilities reduces the likelihood of a successful attack by a factor of about two. In addition to allowing us to estimate the probability of a successful intrusion, the simulation also allows us to estimate the likelihood of attacks appearing at specific locations in the network. This is very useful information for security forces.

Several extensions are possible within the framework of the model developed here, and efforts to extend and improve the analyses are ongoing. This approach appears to offer a significant new tool for evaluating and improving the security of infrastructure facilities.

REFERENCES

- Carlson, R., Turnquist, M.A. and Nozick, L.K. (2004). Expected Losses, Insurability, and Vulnerability to Attacks, Report SAND2004-0742, Sandia National Laboratories, Albuquerque, NM.
- Jha, S., Sheyner, O., and Wing, J.M. (2002). Minimization and Reliability of Attack Graphs. Research Report CMU-CS-02-109, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.
<http://reports-archive.adm.cs.cmu.edu/anon/2002/CMU-CS-02-109.pdf>
- Jones, D.A., Turnquist, M.A. and Nozick, L.K. (2005). Physical Security and Vulnerability Modeling for Infrastructure Facilities, Working Paper, Sandia National Laboratories, Albuquerque, NM.
- Phillips, C.A., and Swiler, L.P. (1998). "A Graph-Based System for Network-Vulnerability Analysis," In ACM Proceedings for the 1998 New Security Paradigms Workshop, pp. 71-81.
- Puterman, M.L. (1994). Markov Decision Processes. Wiley, New York.
- President's Commission on Critical Infrastructure Protection (1997). Critical Foundations: Protecting America's Infrastructures, available on line at <http://www.ciao.gov>
- Swiler, L.P., Phillips, C.A., Ellis, D., and Chakerian, S. (2001). "Computer-Attack Graph Generation Tool," Proceedings of the 2nd DARPA Information Survivability Conference and Exposition, Vol. 2, pp. 307-321.
- Sheyner, O., Haines, J., Jha, S., Lippmann, R., and Wing, J.M. (2002). "Automated Generation and Analysis of Attack Graphs," Proceedings of the IEEE Computer Society Symposium on Research in Security and Privacy, Berkeley, CA, pp. 273-284.

AUTHOR BIOGRAPHIES

DEAN A. JONES is a Principal Member of Technical Staff in the Operations Research & Computational Analysis Group at Sandia National Laboratories. He received his Master's degree in Applied Mathematics from the University of New Mexico in 1992. He specializes in the application of large-scale network optimization solution techniques to a variety of domains including critical infrastructures and transportation. His e-mail address is [<dajones@sandia.gov>](mailto:dajones@sandia.gov).

MARK A. TURNQUIST is a Professor of Civil & Environmental Engineering at Cornell University. He received his PhD in Transportation Systems Analysis from MIT in 1975, and specializes in large-scale network optimization models for use in transportation, logistics, manufacturing and infrastructure systems, with a particular emphasis on models that include uncertainty. He is also active in network modeling related to project planning under uncertainty. His e-mail address is [<mat14@cornell.edu>](mailto:mat14@cornell.edu).

LINDA K. NOZICK is a Professor of Civil & Environmental Engineering at Cornell University. She received her PhD in Systems Engineering from the University of Pennsylvania in 1992, and specializes in systems engineering and applications to transportation, supply chain logistics, infrastructure and project planning. She was a Presidential Early Career Award winner for her work on uncertainty modeling in transportation systems. Her e-mail address is [<lkn3@cornell.edu>](mailto:lkn3@cornell.edu).