

SANDIA REPORT

SAND2012-6050
Unlimited Release
Printed July 2012

Design for Resilience in Infrastructure Distribution Networks

Eric D. Vugrin and Mark A. Turnquist

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



SAND2012-6050
Unlimited Release
Printed July 2012

Design for Resilience in Infrastructure Distribution Networks

Eric D. Vugrin
Resilience and Regulatory Effects
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS1138

Mark A. Turnquist
School of Civil and Environmental Engineering
Cornell University
Ithaca, New York 14853

Abstract

The recognition that resilience is a critical aspect of infrastructure resilience has caused the national and homeland security communities to ask “How does one ensure infrastructure resilience?” Previous analysis methods have primarily focused on system recovery activities following the occurrence of a disruptive event. In this report, we expand on those methods by including pre-disruption investment options, in addition to post-event recovery activities, as means to infrastructure resilience. The report describes a stochastic optimization model for designing network infrastructure resilience to a variety of uncertain potential disruptions. The model seeks investment-recovery combinations that minimize the overall cost to a regional distribution network. A set of numerical experiments illustrates how changes to disruption scenarios probabilities affect the optimal resilient design investments.

ACKNOWLEDGMENTS

The authors thank Patti Glass for her editorial assistance. This work was supported by Sandia National Laboratories Laboratory Directed Research and Development funding.

CONTENTS

1. Introduction.....	6
2. Resilience Measurement and Optimization	9
3. Distribution Network Description.....	11
4. Computational Experiments.....	22
5. Conclusions.....	35
6. References.....	36
Distribution.....	39

FIGURES

Figure 1. Measurement of Systemic Impact.....	9
Figure 2. Recovery Effort Strategies, Costs and Effects of Pre-event Investment.....	9
Figure 3. Example Distribution System Problem.....	12
Figure 4. Modeled Evolution of DC Capacity Over Time.....	18
Figure 5. Delayed Restoration of DC Capacity.....	19
Figure 6. Probability Distribution of Total Impact Across Scenarios, and Definition of V^{ϕ}	20
Figure 7. . Summary of DC Capacity Added in Experiment 1.....	26
Figure 8. Summary of DC Capacity Added in Experiment 2.....	26
Figure 9. Summary of DC Capacity Added in Experiment 3.....	26
Figure 10. Summary of DC Capacity Added in Experiment 4.....	27
Figure 11. Restorative Capacity Investments in the Four Experiments.....	30
Figure 12. Restoration of Capacity in Scenario 5 of Experiment 1.....	31
Figure 13. Probability Distribution of Total Impacts in Experiment 1.....	32
Figure 14. Capacity Additions in Experiment 5.....	32
Figure 15. Distribution of Total Impact in Experiment 5.....	33

TABLES

Table 1. Customer Locations, Demands and DC Assignments.....	13
Table 2. DC loads and capacities.....	14
Table 3. Scenarios and Probabilities – Experiment 1.....	24
Table 4. Scenarios and Probabilities – Experiment 2.....	24
Table 5. Scenarios and Probabilities – Experiment 3.....	24
Table 6. Scenarios and Probabilities – Experiment 4.....	25
Table 7. Summary of Secondary Connections Established.....	29
Table 8. Unmet Demand in Scenario 5 of Experiment 1.....	31

This page intentionally left blank.

1. INTRODUCTION

Many elements of critical infrastructure take the form of networks. These networks provide service by allowing flows (of materials, information, electric power, fuels, etc.), given a capacity or operability state for nodes and links. Disruptions change the operability state of parts of the network (nodes and/or links), and recovery is a set of actions to restore capacity to damaged parts of the network, allowing system performance to return to nominal levels as quickly as possible. Since the 1980's, U.S. Federal Government policy toward critical infrastructure protection has focused primarily on physical asset protection and hardening (Reagan 1982; Clinton 1998; Bush 2002, 2003), but there is now increasing emphasis on infrastructure resilience – the ability of infrastructure systems to withstand, adapt to, and rapidly recover from the effects of a disruptive event. The U.S. Department of Homeland Security National Infrastructure Protection Plan (NIPP) (DHS, 2009) contains explicit language calling for increasing the resilience of the nation's critical infrastructure.

Increasing network resilience involves three related capabilities – providing absorptive capacity so that the network can withstand disruptions, providing adaptive capacity so that flows through the network can be accommodated via alternate paths, and providing restorative capacity so that recovery from a disruptive event can be accomplished quickly and at minimum cost.

There is considerable literature on system recovery in infrastructure networks following a disruptive event. As a few examples, see the work of Xu, et al. (2007) on electric power restoration, Clausen, et al. (2010) on airline system recovery, Luna, et al. (2011) on water distribution networks, Wang, et al. (2011) for internet protocol (IP) networks, and Chen and Miller-Hooks (2012) related to freight transportation networks. However, analyses of recovery strategies do not directly address the important design question: What capabilities, resources and/or network elements should be present to best provide absorptive capacity, adaptive capacity and restorative capacity in infrastructure networks?

Resilience-enhancing investments made prior to the occurrence of disruptive events are important complements to effective post-event recovery strategies. Models that focus on post-event actions can be used in an ad hoc approach to determine benefits that could result from specific efforts to enhance resilience, but this approach frequently must be an iterative process, can be time-consuming, and does not guarantee that one will identify an optimal or near-optimal set of resilience-enhancing investments.

In order to design infrastructure systems that are maximally resilient to a range of threats, infrastructure planners and managers need a capability that can simultaneously consider the impacts of pre-event resilience-enhancing design investments and post-event recovery actions. Concern with design-for-resilience in infrastructure systems has appeared in a variety of contexts in recent years (e.g., Little 2002; Fiksel 2003; Petersen and Johansson 2008; Mansouri, et al. 2010). However, there is little previous work on using explicit stochastic optimization models to design network infrastructures to be resilient to a variety of uncertain potential disruptions. A primary objective of the work described here is to develop such a capability.

One important class of infrastructure networks is distribution networks. These are networks that move or transform materials to meet demands. Production and distribution networks (including electric power, gas distribution, water supply, food production/processing/distribution, manufacturing supply chains, etc.) are examples of this class. The model in this report focuses on an example of such a network to illustrate several core ideas.

2. RESILIENCE MEASUREMENT AND OPTIMIZATION

Vugrin *et al.* (2010) define system resilience as:

Given the occurrence of a particular disruptive event (or set of events), the resilience of a system to that event (or set of events) is the ability to reduce effectively both the magnitude and duration of the deviation from targeted system performance levels.

Measurement of resilience thus includes both the systemic impact (SI) of deviation from desired performance and the resources needed to reduce those impacts. The resource side of resilience is denoted total recovery effort (TRE). Figure 1 illustrates the concept of SI . The occurrence of an event reduces some performance metric for the system, and through recovery effort this metric returns to its nominal level over time, as shown in Figure 1(a). SI is the area of the degraded performance, as shown in Figure 1(b).




Figure 1. Measurement of Systemic Impact.

Varying strategies for recovery may affect SI , but require different levels of recovery effort (cost), as shown in Figure 2(a). It may also be possible to make investments in the system (design improvements) that will reduce the magnitude of the disruption from a given event occurrence as well as speed system recovery, as shown in Figure 2(b). These expenditures may be defined as resilience-enhancing investments (REI). Thus, as we consider system resilience, and design-for-resilience in particular, it is important to incorporate all three elements: REI , SI and TRE , as they vary across some set of potential disruption scenarios.




Figure 2. Recovery Effort Strategies, Costs and Effects of Pre-event Investment.

In general, a system will have several performance metrics and SI measurement must include all the relevant performance dimensions. In some cases, the performance degradation will be measured as

an increase in costs (for operating and using the system, for example), and the measurement of *SI* will be the area above the nominal cost level, rather than below as shown in the example in Figures 1 and 2. These variations are not conceptual differences, but simply reflect different performance measures that may be relevant in different situations.

The specific infrastructure system of interest in this report is a distribution network where the *SI* measures include the additional costs of moving material through a degraded network and a set of penalty costs for not being able to meet all demand, in scenarios where that occurs. Recovery resources are limited, and in any given scenario, these resources must be allocated in the most effective way. Finally, several types of pre-event investment are available – to increase absorptive capacity, adaptive capacity and restorative capacity – and the optimization model trades off the costs of *REI* against the expected costs of *SI* and *TRE* in a range of scenarios. The following section describes the distribution network and its representation in the optimization model.

3. DISTRIBUTION NETWORK DESCRIPTION

Consider a system that consists of a set of spatially separated distribution centers (DCs) and a set of customer locations that receive some product (or set of products) from the DCs. One obvious application of this structure is where the customers are retail stores that receive products from warehouses, but the structure could also represent electrical substations receiving power from generating stations, or municipal water systems receiving water from reservoirs.

Index the customers by i and the DCs by j . Each customer has a demand q_i , and each DC has a capacity K_j . The distance between customer i and DC j is d_{ij} . Under nominal conditions, each customer is connected to its nearest DC and its demand is met through that connection. Assume that the initial DC capacities are designed so that they can accommodate this operation, so that each DC has an initial capacity that equals the sum of the demands assigned to it.

The focus in this analysis is on disruptions that create inoperability of one or more DC's, and the ability of the system to continue to meet demands at the customer locations. The *system design* options available are:

- 1) An opportunity to expand each DC, adding *absorptive capacity* to the system, allowing it to more easily weather the loss of one or more DCs;
- 2) The ability to connect each customer to a single back-up DC, in the event that its primary DC is inoperable or runs out of capacity. In terms of the resilience assessment framework, this is an example of investment in *adaptive capacity*, as it allows the system to adapt to the loss of DC operation by reconfiguring the channels for movement of material;
- 3) Investment in resources to allow faster recovery from a disruption – i.e., ability to restore lost capacity at the DCs more quickly. This is an example of investment in *restorative capacity*.

The model is an optimization that considers all three types of potential pre-event investments, as well as the effects of post-event recovery decisions, and makes tradeoffs to determine an optimal allocation of overall resources to improve system resilience (i.e., minimize the total impact across a range of possible disruption scenarios). The likelihood that any scenario might occur is considered uncertain, so each scenario is assigned a likelihood of occurrence. For the sake of simplicity, we assume these probabilities are independent, but it should be noted that this assumption is reasonable for natural disaster types of disruptions. A network model described below estimates the impacts to the distribution system when one of the disruption scenarios is assumed to have occurred.

An example of this problem type is shown in Figure 3, representing 39 customers spread throughout nine southeastern states in the United States, served by four DCs located in Memphis, TN; Nashville, TN; Knoxville, TN; and Atlanta, GA. Table 1 summarizes the customer locations, demand quantities and DC assignments. Note that each DC location is also a customer location (customers 36-39 in the list in Table 1).




Figure 3. Example Distribution System Problem.

Table 1. Customer Locations, Demands and DC Assignments

Customer #	City	State	Q	Primary Assignment
1	Opelika	AL	25	Atlanta
2	Huntsville	AL	20	Nashville
3	Muscle Shoals	AL	110	Nashville
4	Ormond Beach	FL	137	Atlanta
5	Hialeah	FL	82	Atlanta
6	Orlando	FL	41	Atlanta
7	Hartwell	GA	61	Atlanta
8	Sarepta	LA	173	Memphis
9	Brookhaven	MS	115	Memphis
10	Clarksville	TN	219	Nashville
11	Athens	TN	45	Knoxville
12	Livingston	TN	28	Nashville
13	Springfield	TN	184	Nashville
14	Smyrna	TN	53	Nashville
15	Galatin	TN	32	Nashville
16	Clinton	TN	198	Knoxville
17	Ripley	TN	135	Memphis
18	Lebanon	TN	110	Nashville
19	Jackson	TN	172	Memphis
20	Manchester	TN	86	Nashville
21	Lewisburg	TN	146	Nashville
22	Pulaski	TN	166	Nashville
23	Winchester	TN	153	Nashville
24	Pine Bluff	AR	122	Memphis
25	Fayetteville	AR	84	Memphis
26	Paragould	AR	32	Memphis
27	Charleston	SC	151	Atlanta
28	Fort Mill	SC	144	Knoxville
29	Spartanburg	SC	66	Knoxville
30	Stanfield	NC	70	Knoxville
31	Fletcher	NC	29	Knoxville
32	Jefferson	NC	136	Knoxville
33	Burlington	NC	99	Knoxville
34	Hamptonville	NC	145	Knoxville
35	Newton	NC	162	Knoxville
36	Nashville	TN	166	Nashville
37	Memphis	TN	107	Memphis
38	Knoxville	TN	124	Knoxville
39	Atlanta	GA	138	Atlanta

The DC capacities are set initially to just meet the total demand assigned (computed from the q_i values in Table 1). The values are shown in Table 2.

Table 2. DC loads and capacities.

DC	Nominal Demand and Capacity
Nashville	1473
Memphis	940
Knoxville	1218
Atlanta	635
Total	4266

We consider a set of disruption scenarios that describe different combinations of how the DCs are rendered nonfunctional. We formulate the model in a way to find both the design decisions and the recovery decisions that minimize an overall cost function. The costs include:

- Pre-event costs (*REI*) for investing in absorptive, adaptive and restorative capacity to reduce the system impact of a disruption;
- Post-event costs (*SI* and *TRE*) for:
 - Increased transportation associated with serving some customers from secondary DCs
 - Penalty costs for not meeting demand if insufficient capacity remains after the disruption
 - Costs of recovering the damaged capacity
 - Additional penalty costs for extreme outcomes (i.e., scenarios in which total system impact and recovery costs exceed some threshold).

The post-event (scenario-specific) costs are weighted by the probabilities of the scenarios.

The model formulation is built around eight sets of variables:

w_j = additional initial capacity provided at DC j

$$z_{ij} = \begin{cases} 1 & \text{if customer } i \text{ is connected to DC } j \\ 0 & \text{if not} \end{cases}$$

R = additional restoration capacity investment (units/period)

r_{jt}^s = capacity at DC j that is restored in period t of scenario s

U_{jt}^s = cumulative capacity at DC j that is available in period t of scenario s

x_{ijt}^s = proportion of customer i demand that is served from DC j in period t of scenario s

y_{it}^s = proportion of customer i demand that is not met in period t of scenario s

V^s = amount by which costs (systemic impact plus recovery costs) exceed a threshold G in scenario s .

In addition, there are several parameters and data elements used in the model:

T = total number of time periods considered

B_t = available capability for restoration in period t (units)

c_j = resource requirement for restoring one unit of capacity at DC j

F_j = unit cost of initial additional capacity at DC j

K_j = initial capacity of DC j

h = initial investment required for additional restorative capacity (per unit)

q_i = demand at customer i

d_{ij} = distance from DC j to customer i

p^s = probability of scenario s

G = threshold on total cost that defines an extreme event

γ_j^s = fraction of capacity at DC j that remains available immediately after the disruption represented in scenario s

$j^*(i)$ = index of the closest DC to customer i under nominal conditions.

$\lambda, \phi, \beta, \theta, \xi, \mu$ = weighting coefficients.

The optimization model formulation is as follows:

$$\begin{aligned}
\text{Min} \quad & \sum_j F_j w_j + \theta \sum_i \sum_j z_{ij} + hR \\
& + \sum_s p^s \left[\lambda \sum_i \sum_j q_i d_{ij} \sum_t x_{ijt}^s + \phi \sum_i q_i \sum_t y_{it}^s + \mu \sum_j c_j \sum_t r_{jt}^s \right] \\
& + \sum_s p^s \left[\xi \sum_j \sum_t (K_j + w_j - U_{jt}^s) + \beta V^s \right]
\end{aligned} \tag{1}$$

$$\text{s.t.} \quad x_{ijt}^s - z_{ij} \leq 0 \quad \forall i, j, t, s \tag{2}$$

$$\sum_j z_{ij} \leq 2 \quad \forall i \tag{3}$$

$$z_{i, j^*(i)} = 1 \quad \forall i \tag{4}$$

$$\sum_j x_{ijt}^s + y_{it}^s = 1 \quad \forall i, t, s \tag{5}$$

$$\sum_i q_i x_{ijt}^s \leq U_{jt}^s \quad \forall j, t, s \tag{6}$$

$$U_{jt}^s = \gamma_j^s (K_j + w_j) + \sum_{\tau=1}^{t-1} r_{j\tau}^s \quad \forall j, t, s \tag{7}$$

$$U_{jt}^s \leq K_j + w_j \quad \forall j, t < T, s \tag{8}$$

$$U_{jT}^s = K_j + w_j \quad \forall j, s \tag{9}$$

$$V^s \geq \lambda \sum_i \sum_j q_i d_{ij} \sum_t x_{ijt}^s + \phi \sum_i q_i \sum_t y_{it}^s + \sum_j c_j \sum_t r_{jt}^s - G \quad \forall s \tag{10}$$

$$\sum_j c_j r_{jt}^s \leq B_t + R \quad \forall t, s \tag{11}$$

$$z_{ij} \in \{0, 1\} \quad \forall i, j \tag{12}$$

$$w_j, x_{ijt}^s, y_{it}^s, r_{jt}^s, U_{jt}^s, V^s \geq 0 \quad \forall i, j, t, s \tag{13}$$

The objective (1) minimizes the sum of all costs considered in the model. The first line of the objective is the set of design-related costs incurred for capacity expansion (absorptive capacity), establishing back-up connections for customers to secondary DCs (adaptive capacity), and additional

capability for restoring capacity after a disruption (restorative capacity). The second line of the objective reflects the expected costs (across scenarios) of the movements from DCs to customers, plus the unmet demand costs, plus the costs of restoring capacity in damaged DCs. The coefficients λ and ϕ do the unit conversion to equivalent monetary units of the total transportation movements and the unmet demand values. The third line of the objective includes the expected penalties (across scenarios) for un-restored capacity and extreme scenario impacts. Each of these terms also has a weighting coefficient (ξ and β , respectively).

The objective separates the design-related costs (incurred for decisions made before any disruption scenario is experienced, and that are not adjustable within individual scenarios), and the expected value of post-event costs associated with variables that reflect the specifics of each disruption scenario. The first two lines of the objective function can be re-arranged to put it more clearly in the form of a sum of resilience-enhancing investment (*REI*), systemic impact (*SI*) and total recovery expenditure (*TRE*). The systemic impact in scenario s includes the increase in cost for movement of material, penalties for unmet demand (if any), and other costs associated with having un-restored capacity. These costs are measured relative to the cost in the nominal (base) case. In the nominal case, all demand is met, all installed capacity is operational and movements are from the closest DC to each customer, so if there are T periods in total, the total movement cost is: $\lambda T \sum_i q_i d_{i,j^*(i)}$. Thus, in scenario s , we can write the systemic impact, $SI(s)$, as:

$$SI(s) = \lambda \sum_i \sum_j q_i d_{ij} \sum_t x_{ijt}^s + \phi \sum_i q_i \sum_t y_{it}^s + \xi \sum_j \sum_t (K_j + w_j - U_{jt}^s) - \lambda T \sum_i q_i d_{i,j^*(i)} \quad (14)$$

Across the set of scenarios, the expected SI value is then:

$$E[SI] = \sum_s p^s SI(s) \quad (15)$$

The nominal case cost does not depend on the scenario, and since $\sum_s p^s = 1$, equations (14) and (15) can be combined as follows:

$$E[SI] = \sum_s p^s \left[\lambda \sum_i \sum_j q_i d_{ij} \sum_t x_{ijt}^s + \phi \sum_i q_i \sum_t y_{it}^s + \xi \sum_j \sum_t (K_j + w_j - U_{jt}^s) \right] - \lambda T \sum_i q_i d_{i,j^*(i)} \quad (16)$$

The initial investments include additional capacity at the DCs, investments in back-up connections, and investments to increase recovery resources available. These costs are not scenario-dependent and can be written as:

$$REI = \sum_j F_j w_j + \theta \sum_i \sum_j z_{ij} + hR \quad (17)$$

The recovery costs are incurred after the disruption in each scenario. The expectation of these costs is:

$$E[TRE] = \sum_s p^s \left[\mu \sum_j c_j \sum_t r_{jt}^s \right] \quad (18)$$

The sum of equations (16)-(18) corresponds to the first two lines of the objective function (1), minus a constant (the nominal case cost: $\lambda T \sum_i q_i d_{i,j^*(i)}$) which can be ignored in the optimization. The third line of eq. (1) contains some additional penalty terms that are discussed later.

Constraint (2) says that movements from DC j to customer i cannot be made unless that customer is connected to the DC. Constraints (3) and (4) govern the creation of additional connections between DCs and customers. Constraint (4) specifies the primary connections as given, implying that the decisions made in the model are only for the secondary (backup) connections. Constraint (3) allows one back-up connection for each customer, but does not force these connections to be made.

Constraint (5) defines the unmet demand for cases where customer i cannot be served. In this model, partial service to customer i is possible, and that service may be provided by a combination of the primary DC connection and the secondary connection.

Constraints (6)-(9) represent the capacity evolution of DCs over time in each scenario. Figure 4 illustrates what the model is representing for a specific scenario at a given DC. The DC has an initial capacity, K . A decision is made on investment in additional (absorptive) capacity, bringing the total to $K + w$. At an assumed time $t = 0$, the DC capacity is reduced to $\gamma(K + w)$. In the numerical experiments done below, the value of γ is always either 0 or 1, but in general we could use any value $0 \leq \gamma \leq 1$, and Figure 4 is drawn with an intermediate value of γ .




Figure 4. Modeled Evolution of DC Capacity Over Time.

During the first period, the DC operates with capacity $\gamma(K + w)$, but restoration efforts may be undertaken that will increase capacity by a value r_1 at the end of the first period. That capacity is

available during the second period, and further restoration efforts increase capacity by r_2 at the end of the second period. This process continues, with restoration efforts in each period determined within the optimization. Constraint (7) defines the available capacity at DC j during period t in scenario s , denoted by U_{jt}^s . This can be no greater than the initial augmented level $K + w$ [constraint (8)], and the final level of capacity to which the DC is restored must be that level [constraint (9)].

In each period t , the available capacity at each DC (including all restoration undertaken in the first $t - 1$ periods) is used in an optimal way to distribute material to the various customers. Constraint (6) limits the material distributed from that DC to no more than the available capacity.

Figure 5 illustrates the same restoration of DC capacity as in Figure 4, but the final increment of restoration is delayed from period 3 to period 4. When the restoration resources and capability are limited, such delays may be necessary, but if resources are available it is desirable that restoration be completed as early as possible. When the initial capacity in the system is greater than the nominal demand, the optimization could delay some capacity restoration to the end of the model run without incurring unmet demand penalties or forcing any customers to be served from sub-optimal DC locations. Constraint (9) forces the final capacity to be $K + w$, but constraints (8) and (9) don't force the solution to restore all of the capacity as early as possible. However, this type of early recovery is desirable in the solution, so the first term in the third line of the objective function is there to produce that behavior. That term penalizes the difference between the final restored capacity ($K + w$) and the currently available capacity in period t . This difference is illustrated in Figure 5 (for period 2). By summing across time periods, this term is representing the area above the restored capacity step-function and below the end value, $K + w$. By placing a small cost penalty on that area, the model is encouraged to produce the solution shown in Figure 4, rather than the solution shown in Figure 5, if recovery resources are available.

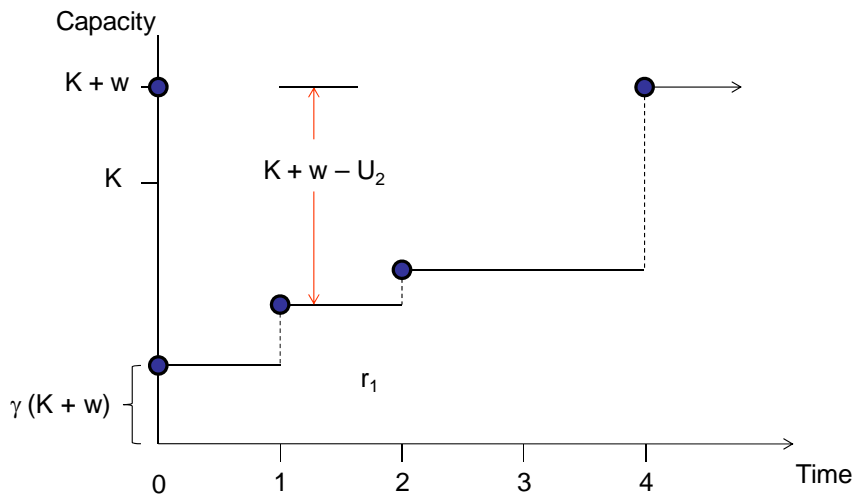


Figure 5. Delayed Restoration of DC Capacity.

Constraint (11) reflects the resource constraint on the restoration activities in each period. A nominal capability to restore disrupted DC capacity (B_t) is available for each period (and may vary across periods). A first-stage design decision can be made to augment that capability by an amount R . The augmentation (purchased at unit cost h , in the objective function) is assumed to be available in all periods. Individual DC facilities may require different amounts of resources to restore one unit of

lost capacity (c_j). Within each period, the available resources ($B_t + R$) can be allocated across various damaged DC's, but the overall level of restoration activity is limited.

Across some set of disruption scenarios, the total cost impact ($SI + TRE$) will vary, and if we attach probabilities to the different scenarios, we can construct a probability distribution of total impact. A conceptual version of such a distribution is shown in Figure 6. The version in Figure 6 is drawn as a continuous probability density function, and this may be appropriate in some circumstances. If we represent uncertainty in the model with discrete scenarios, the distribution of $SI + TRE$ will be discrete, and should be drawn as a probability mass function, but this distinction is unimportant for the moment.




Figure 6. Probability Distribution of Total Impact Across Scenarios, and Definition of V^s .

The purpose of the optimization is to find a set of investment and operational decisions that shift this distribution to the left, resulting in smaller total impacts, and the model measures the expected value of this distribution in the objective function. However, we may also be particularly sensitive to extreme values in the right-hand tail of this distribution, representing a subset of scenarios which produce very large impacts (in general, with small probability).

The variables V^s , and constraints (10) which define their values, create a means of placing special emphasis on reducing the extreme impact values. To implement this mechanism, we define an input parameter, G , that represents a threshold value of total impact for definition of what constitutes an *extreme* scenario. Constraints (10) require that if the system impact plus recovery effort in scenario s exceeds the value G , then V^s is defined as the amount of the difference. The values of V^s are penalized in the objective function (with non-negative weights, βp^s). For each scenario, $V^s \geq 0$ [see constraint (13)], so if the system impact plus recovery effort in a given scenario is less than G , the corresponding $V^s = 0$. However, if the total impact exceeds G , constraint (10) forces V^s to measure the difference. The weights on the V^s terms in the objective function reflect both the probability of occurrence for the extreme scenarios and the overall relative weighting (β) of these large impacts relative to the expected value (which is still computed across all scenarios, extreme or not).

The combination of the input parameters β and G allows us to tune how sensitive the model is to extreme outcomes, by determining which outcomes are counted as extreme (G), and how heavily they are weighted (β).

Mathematically, the optimization model formulated in eqs. (1)-(13) is a mixed-integer linear programming (MILP) problem, because it contains both continuous and discrete decision variables. It is also a particular form of optimization formulation known as a two-stage stochastic programming problem because it contains scenarios whose occurrence is uncertain and represented by probabilities. Some of the variables in the problem are first-stage variables (i.e., determined before the scenario outcome is known), and others are second-stage variables (determined specifically in each scenario). The second-stage variables are sometimes called *recourse* variables, because their values are determined in each scenario, but are conditioned on the choices made in the first stage. A general discussion of stochastic programming models can be found in Kall and Wallace (1994) or Birge and Louveaux (1997).

This page intentionally left blank.

4. COMPUTATIONAL EXPERIMENTS

To explore various aspects of the model, five numerical experiments have been performed using the illustrative setting of four DCs and 39 customers described at the beginning of section 3. All experiments are based on a set of eleven different scenarios, corresponding to loss of all combinations of zero, one or two DC's. Possible scenarios involving three or more DC outages are not considered in these experiments. The presumption in this example is that those scenarios have sufficiently low probability that it is not worth planning for them. Four of the five experiments focus on the scenarios, their relative likelihoods, and the effects of changes on the solutions produced by the optimization model. For these four experiments, the parameter (β) for including penalties on extreme costs has been set to zero, so the model is minimizing expected costs across the set of scenarios. In the fifth experiment, the term relating to extreme costs is introduced, and the effects that has on the model solution are illustrated.

In all five experiments, input parameters in the model have been set as follows:

$$T = 10 \text{ (total number of time periods considered)}$$

$$B_t = 500 \text{ (nominal units of available restoration capability in each period } t \text{)}$$

$$c_j = 1 \text{ (unit requirement for restoring capacity at each DC)}$$

$$h = 300 \text{ (cost of augmenting restoration capability by one unit per period)}$$

$$F_j = 100 \text{ (unit cost of initial additional capacity at each DC)}$$

$$\lambda = 0.1 \text{ (weighting coefficient on ton-miles of movement)}$$

$$\phi = 500 \text{ (penalty coefficient for unmet demand)}$$

$$\theta = 4000 \text{ (fixed cost of additional DC-customer connection)}$$

$$\xi = 10 \text{ (weighting coefficient on un-restored capacity at disrupted DCs)}$$

$$\mu = 50 \text{ (cost conversion coefficient on restoration effort at disrupted DCs).}$$

Computations for all experiments reported here have been done using a commercial MIP solver, Lingo (version 12), (Lindo Systems, 2010). For the modest example solved here, the optimization problem has approximately 22,500 variables and 23,000 constraints, and a typical solution requires about 4 minutes of computation time on a laptop computer. As the problem size increases (larger network, more scenarios, more time periods), the computation times will also increase.

In the first two experiments, the probability of the “no disruption” scenario (i.e., no DCs out of service) is set to zero. Thus, the analysis is based on an assumption that *some* disruption will occur, and the specified probabilities represent the relative likelihood of different types and magnitudes of

disruption. The analysis in these two experiments represents a perspective in which we've committed to preparing for various types of disruptions and want to decide how best to do that. In that instance, the "no disruption" scenario becomes irrelevant (even though there may be a large probability of no disruption occurring over any given planning horizon – e.g., a year). Table 3 shows the assumed scenario probabilities for experiment 1. The single-outage scenarios are assumed to be more likely than the dual outages, and disruption of the various DCs is assumed to be equally likely.

Table 3. Scenarios and Probabilities – Experiment 1

Scenario Number	Inoperable DC's	Probability
1	Nashville	0.175
2	Memphis	0.175
3	Knoxville	0.175
4	Atlanta	0.175
5	Nashville, Memphis	0.05
6	Nashville, Knoxville	0.05
7	Nashville, Atlanta	0.05
8	Memphis, Knoxville	0.05
9	Memphis, Atlanta	0.05
10	Knoxville, Atlanta	0.05
11	No Disruption	0

In experiment 2, the probabilities are adjusted geographically to make outages at the eastern DCs (Atlanta and Knoxville) more likely than disruption at the western DCs (Memphis and Nashville). The scenario probabilities are shown in Table 4.

Table 4. Scenarios and Probabilities – Experiment 2

Scenario Number	Inoperable DC's	Probability
1	Nashville	0.117
2	Memphis	0.117
3	Knoxville	0.233
4	Atlanta	0.233
5	Nashville, Memphis	0.023
6	Nashville, Knoxville	0.046
7	Nashville, Atlanta	0.046
8	Memphis, Knoxville	0.046
9	Memphis, Atlanta	0.046
10	Knoxville, Atlanta	0.093
11	No Disruption	0

Experiments 3 and 4 include the "no disruption" scenario with two different probabilities (0.7 and 0.9) attached to it, representing an analysis that reflects expected value decision making when there is a relatively large probability that nothing happens to disrupt normal operations. In both of these experiments, the assumption of equally likely disruption at the various DC's (from experiment 1) is maintained. Tables 5 and 6 illustrate the scenario probabilities for these experiments.

Table 5. Scenarios and Probabilities – Experiment 3

Scenario Number	Inoperable DC's	Probability
1	Nashville	0.0525

2	Memphis	0.0525
3	Knoxville	0.0525
4	Atlanta	0.0525
5	Nashville, Memphis	0.015
6	Nashville, Knoxville	0.015
7	Nashville, Atlanta	0.015
8	Memphis, Knoxville	0.015
9	Memphis, Atlanta	0.015
10	Knoxville, Atlanta	0.015
11	No Disruption	0.7

Table 6. Scenarios and Probabilities – Experiment 4

Scenario Number	Inoperable DC's	Probability
1	Nashville	0.0175
2	Memphis	0.0175
3	Knoxville	0.0175
4	Atlanta	0.0175
5	Nashville, Memphis	0.005
6	Nashville, Knoxville	0.005
7	Nashville, Atlanta	0.005
8	Memphis, Knoxville	0.005
9	Memphis, Atlanta	0.005
10	Knoxville, Atlanta	0.005
11	No Disruption	0.9

Figures 7-10 summarizes the absorptive capacity solutions (added DC capacity) in the four experiments. In experiment 1 (Figure 7), a total of 2494 units of capacity are added, and the additional capacity is allocated to make three of the four DC's nearly equal in total capacity. The 2494 units of additional capacity represents an increase of about 58%, and allows the system to absorb the capacity disruption in nearly all the scenarios with very little unmet demand. In experiment 2 (Figure 8), the total added capacity is somewhat smaller (2102 units), and the capacity additions are focused on the western DC's (Nashville and Memphis) which have the smaller probability of disruption.

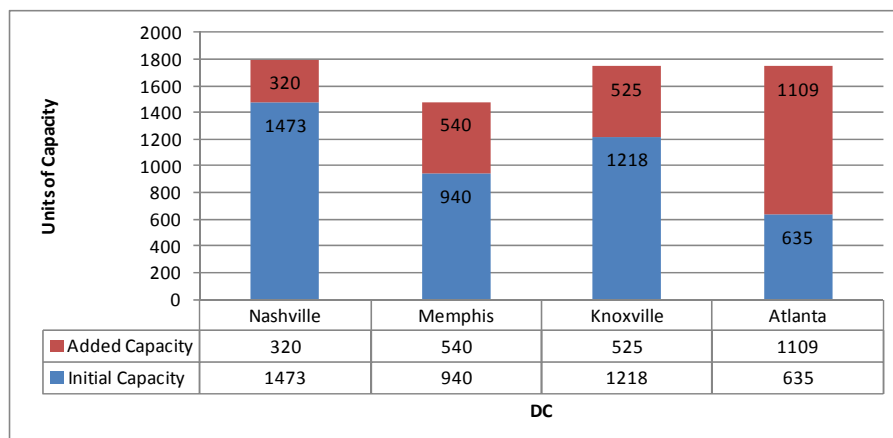


Figure 7. . Summary of DC Capacity Added in Experiment 1.




Figure 8. Summary of DC Capacity Added in Experiment 2.

In experiment 3, where there is a relatively high probability of no disruption at all, the amount of added capacity is substantially smaller, 888 units, and focused in Atlanta and Memphis, as shown in Figure 9. In experiment 4 (Figure 10), where the probability of no disruption is 0.9, the optimal solution is to not add any absorptive capacity at all. The decrease in total added capacity from experiment 1 to experiment 3, and from experiment 1 to experiment 4, mirrors the decreasing likelihood of requiring it to absorb a disruption. In experiment 3, the added capacity is used to make the four DC's have more nearly equal total capacity, as in experiment 1. This reflects the assumption that disruptions, should they occur, are equally likely at all DC's.

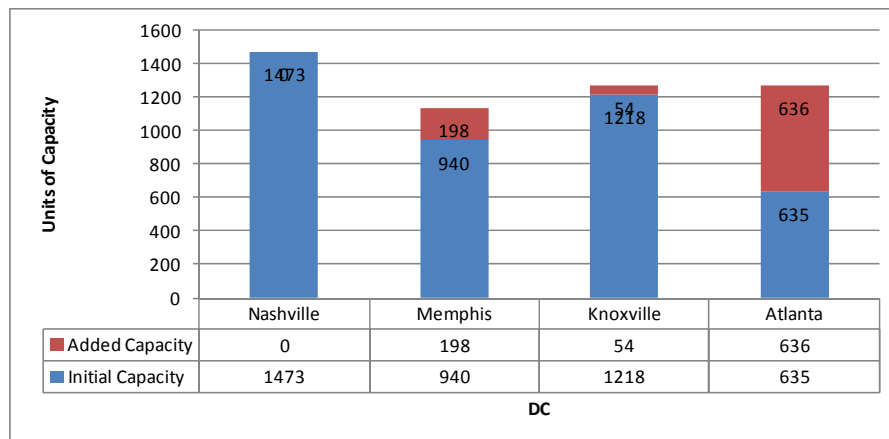


Figure 9. Summary of DC Capacity Added in Experiment 3.




Figure 10. Summary of DC Capacity Added in Experiment 4.

The pattern of secondary connections to customers in the four solutions is summarized in Table 7. For experiments 1 and 2, a total of 32 of the 39 customers have secondary connections established, but the specific connections created are not the same in the two experiments. In experiment 1, there is a clear cutoff based on demand volume – customers whose volume is 45 or higher have a secondary connection, and the seven smallest customers do not. In experiment 2, there is also a strong correlation with demand volume, but a few exceptions are made based on customer location because the probability of disruption is not the same for all DC’s.

In experiment 3, where there is a 70% chance of no disruption at all, the number of back-up connections to customers is only 14, and these connections are generally made for the largest customers. The majority of the back-up connections are to Atlanta, and this is consistent with the investment in absorptive capacity at that DC, noted in Figure 9.

In experiment 4, where no investment in absorptive capacity is made, so the DC’s all have capacity that just meets the nominal assigned demand, there are also no back-up connections made for adaptive capacity because in any disruption scenario the unaffected DC’s would have no additional capacity to handle extra demand. The joint decisions to make no investment in either absorptive capacity or adaptive capacity in this experiment reflect the very high probability of no disruption.

It is also noteworthy that many of the back-up connections made in experiments 1-3 are not to the second closest DC for a given customer. Although using the next closest DC would appear to be a natural decision for each customer when considered independently, there are much more complex interactions of decisions across the network – where to invest in absorptive capacity, and how these decisions affect the creation of adaptive capacity. This illustrates the importance of taking a system-wide perspective on the design of investments to enhance resilience.

Figure 11 shows the amount of restorative capacity investment in each of the experiments. In experiment 1, the addition to capacity to allow more rapid restoration after disruptions is about 8% of the nominal value (250) assumed to exist *a priori*. In experiment 2, there is greater investment, equal to about a 23% increase in restoration capability. This should be viewed in concert with the decision in experiment 2 to invest less in adaptive capacity than in experiment 1. In experiment 2, it is more likely that disruptions will strike in the eastern part of the network (Knoxville and Atlanta),

so absorptive capacity is concentrated in the “safer” western DC’s, and is smaller in total than in experiment 1. However, there is larger investment in restorative capacity to allow disrupted DC’s to be brought back online more quickly and reduce total system impact that way.

Table 7. Summary of Secondary Connections Established

#	City	State	Q	Primary Assignment	2nd Closest DC	Secondary Assignment			
						Exp. 1	Exp. 2	Exp. 3	Exp. 4
1	Opelika	AL	25	Atlanta	Knoxville	--	--	--	--
2	Huntsville	AL	20	Nashville	Atlanta	--	--	--	--
3	Muscle Shoals	AL	110	Nashville	Memphis	Memphis	Memphis	--	--
4	Ormond Beach	FL	137	Atlanta	Knoxville	Nashville	Nashville	--	--
5	Hialeah	FL	82	Atlanta	Knoxville	Memphis	Memphis	--	--
6	Orlando	FL	41	Atlanta	Knoxville	--	Memphis	--	--
7	Hartwell	GA	61	Atlanta	Knoxville	Knoxville	Nashville	--	--
8	Sarepta	LA	173	Memphis	Nashville	Atlanta	Nashville	Atlanta	--
9	Brookhaven	MS	115	Memphis	Nashville	Atlanta	Atlanta	--	--
10	Clarksville	TN	219	Nashville	Memphis	Memphis	Memphis	Memphis	--
11	Athens	TN	45	Knoxville	Atlanta	Memphis	Nashville	--	--
12	Livingston	TN	28	Nashville	Knoxville	--	--	--	--
13	Springfield	TN	184	Nashville	Knoxville	Knoxville	Memphis	Knoxville	--
14	Smyrna	TN	53	Nashville	Knoxville	Atlanta	--	--	--
15	Galatin	TN	32	Nashville	Knoxville	--	--	--	--
16	Clinton	TN	198	Knoxville	Nashville	Nashville	Nashville	Nashville	--
17	Ripley	TN	135	Memphis	Nashville	Knoxville	Nashville	Nashville	--
18	Lebanon	TN	110	Nashville	Knoxville	Knoxville	Atlanta	--	--
19	Jackson	TN	172	Memphis	Nashville	Nashville	Nashville	Atlanta	--
20	Manchester	TN	86	Nashville	Knoxville	Knoxville	Atlanta	--	--
21	Lewisburg	TN	146	Nashville	Knoxville	Atlanta	Memphis	--	--
22	Pulaski	TN	166	Nashville	Memphis	Memphis	Memphis	Atlanta	--
23	Winchester	TN	153	Nashville	Knoxville	Atlanta	Atlanta	Atlanta	--
24	Pine Bluff	AR	122	Memphis	Nashville	Atlanta	Atlanta	--	--
25	Fayetteville	AR	84	Memphis	Nashville	Knoxville	Nashville	--	--
26	Paragould	AR	32	Memphis	Nashville	--	--	--	--
27	Charleston	SC	151	Atlanta	Knoxville	Knoxville	Nashville	Knoxville	--
28	Fort Mill	SC	144	Knoxville	Atlanta	Atlanta	Atlanta	Atlanta	--
29	Spartanburg	SC	66	Knoxville	Atlanta	Atlanta	Atlanta	--	--
30	Stanfield	NC	70	Knoxville	Atlanta	Atlanta	Atlanta	--	--
31	Fletcher	NC	29	Knoxville	Atlanta	--	--	--	--
32	Jefferson	NC	136	Knoxville	Atlanta	Nashville	Nashville	Atlanta	--
33	Burlington	NC	99	Knoxville	Atlanta	Atlanta	Nashville	--	--
34	Hamptonville	NC	145	Knoxville	Atlanta	Atlanta	Nashville	Atlanta	--
35	Newton	NC	162	Knoxville	Atlanta	Atlanta	Nashville	Atlanta	--

36	Nashville	TN	166	Nashville	Knoxville	Atlanta	Memphis	Atlanta	--
37	Memphis	TN	107	Memphis	Nashville	Knoxville	Atlanta	--	--
38	Knoxville	TN	124	Knoxville	Nashville	Nashville	Nashville	--	--
39	Atlanta	GA	138	Atlanta	Knoxville	Nashville	Nashville	--	--




Figure 11. Restorative Capacity Investments in the Four Experiments.

In experiments 3 and 4, when there is a substantial probability of no disruption, there is less investment in additional restorative capacity, and in experiment 4, none at all. This reflects the reduced likelihood that such capacity will be needed (and also the relative costs assumed in these experiments for the various types of resilience-enhancing investments, the costs of unmet demand, etc.).

The results summarized in Figures 7-11, and in Table 7, reflect the design decisions for the network – those investments in the three types of enhancements to system resilience that allow the system to respond to a variety of potential disruption scenarios. In each scenario, there are also adaptation and recovery decisions that are made (within the limits created by the original design) to minimize the system impact and recovery effort. For example, in experiment 1, let us focus on scenario 5 (Nashville and Memphis disrupted). For this experiment, there has been investment in absorptive capacity to increase the capability of all the DCs, so some of the customers normally served by Nashville and Memphis can be shifted to Atlanta and Knoxville. These shifts follow the pattern shown in Table 7, where secondary connections have been established (adaptive capacity). Some customers nominally served by Nashville or Memphis (e.g., customer 2) had no secondary connection established, so their demand is unmet until capacity can be restored at their original DC. Other customers (e.g., customer 3) did have secondary connections established, but to another DC that is disrupted. Their demand is also unmet until some capacity is restored at one or the other of the DCs to which they are connected. However, many of the customers nominally assigned to Nashville or Memphis have secondary connections to Atlanta or Knoxville, and these customers can be served (at least within the capacity established at Atlanta and Knoxville).

Table 8 shows the unmet demand in this scenario of experiment 1, summarized by customer location and time period. The unmet demand is concentrated in eight customer locations, four of which have no back-up connections. In the first period after the disruption, the total unmet demand is 779 units (approximately 18% of the total demand). During the first period, 521 units of capacity are restored (the capability created by the 500 units of nominal restoration capacity plus 21 units of additional restorative capacity in which initial investment was made), and the unmet demand in the second period falls to 258 units. By the third period, sufficient restoration has been accomplished to eliminate the unmet demand in the network.

Figure 12 shows the restoration of capacity at Nashville and Memphis in this scenario. At Nashville, there is modest additional capacity investment initially (see Figure 7). The nominal initial capacity of 1473 is increased by 320 units, to 1793. After the capacity is lost in the disruption at $t = 0$, it is restored over six periods, with the original capacity of 1473 reached after five periods. At Memphis, there is a much larger initial investment in absorptive capacity, increasing the total capacity at the DC from 940 to 1480, an increase of approximately 57%. After the disruption, the restoration of the total capacity requires seven periods, but the original capacity of 940 is reached after four periods. The available recovery resources are divided between the two DCs in each period, although not necessarily equally.

Table 8. Unmet Demand in Scenario 5 of Experiment 1

Customer #	City	State	Q	Unmet Demand by Period		
				1	2	3
2	Huntsville	AL	20	20	20	0
3	Muscle Shoals	AL	110	110	110	0
10	Clarksville	TN	219	219	0	0
12	Livingston	TN	28	28	28	0
15	Gallatin	TN	32	32	0	0
19	Jackson	TN	172	172	100	0
22	Pulaski	TN	166	166	0	0
26	Paragould	AR	32	32	0	0
Total				779	258	0




Figure 12. Restoration of Capacity in Scenario 5 of Experiment 1.

In each experiment, there is a distribution of total impact (*SI* plus *TRE*) values across the various scenarios. For example, Figure 13 shows the distribution for experiment 1. The smaller impact values of (less than \$1 million) are associated with the scenarios that have a single DC disrupted. The larger values correspond to the scenarios that involve two disruptions simultaneously.




Figure 13. Probability Distribution of Total Impacts in Experiment 1.

In these four experiments, no special weight is placed on scenarios that cause large impact values (i.e., $\beta = 0$), but if that were changed and a value of G were specified at \$1.4 million (for example), we could expect some changes in overall policy to reduce the likelihood and magnitude of the largest impacts. Experiment 5 implements values of $G = \$1.4$ million and $\beta = 2$, with all other parameters identical to experiment 1.

One of the principal results from experiment 5 is shown in Figure 14. Comparing Figure 14 with Figure 7, we see that when the extreme impacts are weighted more heavily, there is slightly more absorptive capacity added to the system (2632 units vs. 2494 units), but the pattern of investments is quite similar – bringing the four DCs up to essentially equal capacity.

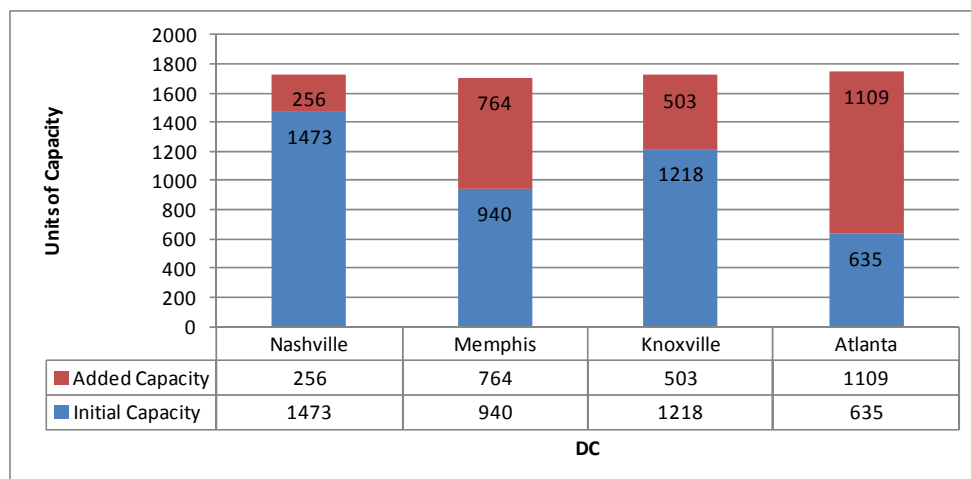


Figure 14. Capacity Additions in Experiment 5.

The total restorative capacity added in experiment 5 is also larger than in experiment 1 (75 units vs. 21 units). The combination of larger absorptive capacity and larger restorative capacity allows faster recovery in the most disruptive scenarios, reducing the total impact cost below the threshold. The distribution of total impact across the scenarios for experiment 5 is shown in Figure 15. The range of impacts for the more severe scenarios (two DCs disrupted) has been reduced very significantly, with all six of those scenarios having nearly equal impacts.




Figure 15. Distribution of Total Impact in Experiment 5.

This page intentionally left blank.

5. CONCLUSIONS

Increasing network resilience involves three related capabilities – providing *absorptive capacity* so that the network can withstand disruptions, providing *adaptive capacity* so that flows through the network can be accommodated via alternate paths, and providing *restorative capacity* so that recovery from a disruptive event can be accomplished quickly and at minimum cost. Resilience-enhancing investments made prior to the occurrence of disruptive events are important complements to effective post-event recovery strategies and a design capability that considers these decisions jointly represents an important advance in tools available to infrastructure planners/managers.

Distribution networks that move or transform materials to meet demands (including electric power, gas distribution, water supply, food production/processing/distribution, manufacturing supply chains, etc.) are an important class of infrastructure networks. The model in this report focuses on an example of such a network to illustrate several core ideas.

This report describes a stochastic optimization model that addresses the design question: What capabilities, resources and/or network elements should be present to best provide network resilience against a variety of potential damage scenarios? This model includes design decisions that represent possible investments in absorptive capacity, adaptive capacity and restorative capacity simultaneously. By including potential investments in all three types of resilience-enhancing capacity, as well as the recovery strategy for a variety of disruption scenarios, we have a useful overall framework for examining design-for-resilience in infrastructure networks.

Mathematically, the optimization model is a mixed-integer linear programming (MILP) problem, because it contains both continuous and discrete decision variables. It has the particular structure of a two-stage stochastic programming problem because it contains scenarios whose occurrence is uncertain and represented by probabilities. Some of the variables in the problem are first-stage variables (i.e., determined before the scenario outcome is known), and others are second-stage variables (determined specifically in each scenario).

A series of computational experiments on a test network with four distribution centers and 39 customer locations has allowed exploration of the behavior of network solutions as important parameters of the problem are varied. This small set of experiments is intended to illustrate basic properties of the solutions, but is not intended to be comprehensive. A wide variety of other numerical experiments could be done with the model.

Computations for all experiments reported here have been done using a commercial MILP solver. For the modest example solved here, the optimization problem has approximately 22,500 variables and 23,000 constraints, and a typical solution requires about 4 minutes of computation time on a laptop computer. As the problem size increases (larger network, more scenarios, more time periods), the computation times will also increase, and this provides motivation for exploring specialized solution methods that might better take advantage of the problem structure.

This page intentionally left blank.

6. REFERENCES

1. Birge, J.R. and F.V. Louveaux (1997). *Introduction to Stochastic Programming*, Springer-Verlag, New York.
2. Bush, G.W. (2002). Homeland Security Presidential Directive-3 (HSPD-3), Washington, DC.
3. Bush, G. W. (2003). Homeland Security Presidential Directive-7 (HSPD-7), Washington, DC.
4. Chen, L. and E. Miller-Hooks (2012). Resilience: An Indicator of Recovery Capability in Intermodal Freight Transport, *Transportation Science*, 46:1, 109-123.
5. Clausen, J., A. Larsen, J. Larsen and N.J. Rezanova (2010). Disruption Management in the Airline Industry – Concepts, Models and Methods, *Computers & Operations Research*, 37:809-821.
6. Clinton, W. (1998). Presidential Decision Directive PDD-63, Protecting America’s Critical Infrastructures, Washington, DC.
7. Fiksel, J. (2003). Designing resilient, sustainable systems, *Environmental Science and Technology* 37(23), 5330 – 5339.
8. Kall, P. and S.W. Wallace (1994). *Stochastic Programming*, Wiley, Chichester, UK.
9. Lindo Systems, Inc. (2010). *Lingo – Optimization Modeling Software for Linear, Nonlinear and Integer Programming*, Chicago, IL.
10. Little, R.G. (2002). Toward More Robust Infrastructure: Observations on Improving the Resilience and Reliability of Critical Systems, in *Proceedings of the 36th Hawaii International Conference on Systems Sciences*, Kauai, Hawaii.
11. Luna, R., N. Balakrishnan and C.H. Dagli (2011). Postearthquake Recovery of a Water Distribution System: Discrete Event Simulation Using Colored Petri Nets, *Journal of Infrastructure Systems*, American Society of Civil Engineers, 17, 25-34.
12. Mansouri, M., R. Nilchiana and A. Mostashari (2010). A Policy Making Framework for Resilient Port Infrastructure Systems, *Marine Policy*, 34:6, 1125-1134.
13. Petersen, K.E. and H. Johansson (2008). Designing Resilient Critical Infrastructure Systems Using Risk and Vulnerability Analysis, *Resilience Engineering Perspectives*, E. Hollnagel, C.P. Nemeth and S. Dekker, eds., Ashgate Publishers, Farnham, Surrey, UK, 160-171.

14. Reagan, R. (1982). Executive Order 13282, National Security Telecommunications Advisory Committee, Washington, DC.
15. U.S. Department of Homeland Security (2009). National Infrastructure Protection Plan, Washington, DC.
16. Vugrin, E. D., Warren, D. E., Ehlen, M. A., and Camphouse, R. C. (2010). "A framework for assessing the resilience of infrastructure and economic systems," Sustainable and resilient critical infrastructure systems: simulation, modeling, and intelligent engineering, K. Gopalakrishnan and S. Peeta, eds., Springer-Verlag, Inc., Berlin, 77-116.
17. Wang, N., A. Fagear and G. Pavlou (2011). Adaptive Post-failure Load Balancing in Fast Reroute Enabled IP Networks, Proceedings of the 12th IFIP/IEEE International Symposium on Integrated Network Management, IM 2011, 470-477.
18. Xu, N., Guikema, S.D., Davidson, R.A., Nozick, L.K., Çağnan, Z., and Vaziri, K. (2007). "Optimizing scheduling of post-earthquake electric power restoration tasks," Earthquake Engineering and Structural Dynamics, 36, 265-284.

DISTRIBUTION

1 MS0899 Technical Library 9536 (electronic copy)

For LDRD reports, add:

1 MS0359 D. Chavez, LDRD Office 1911

This page intentionally left blank.



Sandia National Laboratories