

Interdependent Network Restoration: Modeling Restoration Interdependencies and Evaluating the Value of Information-Sharing

Thomas C. Sharkey^{1,2} Burak Cavdaroglu³ Huy Nguyen² Jonathan Holman⁴
John E. Mitchell⁴ William A. Wallace²

July 11, 2014

Abstract

We consider restoring multiple interdependent infrastructure networks after a disaster damages components in them and disrupts the services provided by them. Our particular focus is on interdependent infrastructure restoration where both the *operations* and the *restoration* of the infrastructures are linked across systems. We provide new mathematical formulations of *restoration interdependencies* in order to incorporate them into an interdependent integrated network design and scheduling (IINDS) problem. The interdependent infrastructure restoration efforts resulting from solving this IINDS problem model a centralized decision-making environment where a single decision-maker controls the resources of all infrastructures. In reality, individual infrastructures often determine their restoration efforts in an independent, decentralized manner with little communication among them. We provide algorithms to model various levels of decentralization in interdependent infrastructure restoration. These algorithms are applied to realistic damage scenarios for interdependent infrastructure systems in order to determine the loss in restoration effectiveness resulting from decentralized decision-making. Our computational tests demonstrate that this loss can be greatly mitigated by having infrastructures *share information* about their planned restoration efforts.

Keywords: OR in societal problem analysis, OR in disaster relief, Interdependent infrastructure restoration

1 Introduction

A disaster is a non-routine event that can have a catastrophic impact on physical, natural, and social systems. Disasters can cause significant damage to residential and commercial structures as well as resulting in immediate and prolonged loss of services from critical civil infrastructure systems. A critical civil infrastructure (CCI) system is defined as an infrastructure with physical components (e.g., transmission lines, cables, pipes) that provide key services to a community. Electrical power systems, transportation, telecommunications, water supply systems, and wastewater systems are examples of CCI systems. While civil

¹Corresponding Author, e-mail: sharkt@rpi.edu.

²Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute, 110 8th St, Troy, NY 12180.

³Industrial Engineering Department, Isik University, Istanbul, Turkey.

⁴Department of Mathematical Sciences, Rensselaer Polytechnic Institute, 110 8th St, Troy, NY 12180.

infrastructures are typically managed by people, the services in these infrastructures are provided by the physical components that encompass the infrastructure itself. One of the basic motivations for managers of infrastructure systems is to enhance the resiliency of communities so that they recover quickly from natural, technological or willful events that may have disastrous impacts on the CCI systems. For example, Hurricane Sandy, which struck the East Coast of the United States in 2012, caused damage estimated at \$65 billion (Hurricane Sandy Rebuilding Task Force [17]) with repair and cleanup costs for New York alone estimated to be \$33 billion (see Kaplan and Hernandez [19]). It had significant effects on multiple infrastructures in these areas including, but not limited to, over 4.5 million customers without power in New York and New Jersey (U.S. Department of Energy, Office of Delivery and Reliability [36]), 57 terminals associated with fuel distribution closed ([36]), all subway lines closed south of 42nd street in New York City with only 80% of the subway operational 5 days after (Kaufman et al. [21]), and over 10 billion gallons of spilled sewage resulting from damage to wastewater treatment plants (Kenward et al. [22]). Therefore, the restoration of services after Hurricane Sandy required significant efforts across the multiple impacted infrastructures.

One of the most effective ways of limiting the impact of the disaster on society is the timely development of a plan to restore the services disrupted by the disaster. In recent years, CCI systems have become more vulnerable to extreme events and the services provided by them are more difficult to restore due to the increasing number of interdependencies among them. The concept of *operational interdependencies* occurs when a component of one infrastructure requires services provided by another infrastructure in order to function properly and has been well-studied (see, e.g., Rinaldi et al. [38], Little [27], and Wallace et al. [43]). These operational interdependencies allow disruptions in one system to spread to others and cause cascading failures across the interdependent infrastructures (see, e.g., McDaniels et al. [31], Lee et al. [24, 25], and McDaniels et al. [32]). For example, if power to a wastewater treatment plant is disrupted and the plant does not have a backup generator, then services in the wastewater system will be disrupted. In addition, the restoration efforts of the different infrastructures are often linked in terms of precedence relationships. Recently, Sharkey et al. [40] introduced the concept of *restoration interdependencies* by cataloging incidents after Hurricane Sandy which linked the restoration efforts of multiple infrastructures. For example, the clearing of downed trees from a street (in the road infrastructure) may not be able to start until a power crew inspects and removes downed power lines from the street. This situation represents a precedence relationship (see Pinedo [37]) between scheduling the restoration tasks in the power and road infrastructures.

Despite these interdependencies, the CCI systems are often controlled and operated independently of one another by both public and private sector decision-makers. This fact complicates the formulation of interdependent infrastructure restoration (IIR) efforts after a disaster since the decision-makers in control of each infrastructure will formulate their individual restoration efforts independently, often with little or no communication with other infrastructures (see, e.g., Comfort [10], Leavitt and Kiefer [23], and McGuire and Schneck [33]). The ideal situation for a community is that IIR efforts are formed in a fashion where there is a *centralized* decision-maker that determines the restoration efforts of *all infrastructures* in order to recover the set of CCI systems. In particular, the centralized efforts will focus on maximizing the total performance of the CCI systems over the course of the restoration horizon after the disaster. Even though centralized decision-making seems to be the best managerial strategy for increasing the resilience of communities after

disasters, it might not always be feasible in practice. Conflicts of interest among infrastructures might be observed during restoration planning because the priorities of each infrastructure might be different from one another. For example, a private sector power company might prefer to restore the services for its higher priority customers first before restoring the power demand of other infrastructures. These conflicts can deter the managers of some infrastructure systems from considering the performance of the restoration over all infrastructure systems and persuade them to make independent decisions which maximize the restoration efforts of their infrastructure. One way of limiting the impact of independent decision-making on restoration of services is governmental interventions. In the hurricane exercises we attended in the Emergency Operating Center (EOC) of New Hanover County, North Carolina in the United States, we observed that the decisions of private sector infrastructures, such as telecommunications and power, are facilitated by the county's local emergency managers. Even though these local emergency managers do not directly intervene in the restoration planning decisions, they facilitate communication among the infrastructures, and advise them on restoration priorities while taking into account the needs of other systems and the community.

Information-sharing during post-disaster restoration planning is an example of collaborative efforts among interdependent infrastructure systems. In order to make effective restoration decisions, managers of an infrastructure system need to consider the restoration plans of the infrastructures that they are dependent upon. Even though infrastructure managers prefer independent decision-making, they can share their restoration plan with the managerial units of the infrastructures that use their services. Sharing information on restoration plans among interdependent systems not only helps infrastructures to improve their individual recovery performance, but also makes the community more resilient. In the United States, the National Infrastructure Coordinating Center (NICC) that operates under the Department of Homeland Security (DHS) is an example of how various infrastructures are coordinated for information-sharing at the national level in case of emergencies. The NICC serves as the information and coordination hub of a national network dedicated to protecting critical infrastructures. The center's primary function in case of a disaster is integration and dissemination of information throughout the critical infrastructure partnership network (U.S. Department of Homeland Security [42]).

The contribution of this paper is to provide an assessment of the improvement in restoration effectiveness that results from information-sharing in the context of decentralized IIR efforts. Although there have been studies concerning collaboration and information-sharing in the context of infrastructure and disaster management (Botterud et al. [6], Dawes [11], Lee and Rao [26], Caruson and MacManus [7], Somers and Svava [41], and Kapucu and Garayev [20]), this paper is the first to quantify the improvement in the restoration efforts that results from information-sharing and coordination among infrastructures. In particular, this paper addresses: (i) the impact of fully decentralized decision-making across infrastructures in terms of the loss of performance of the CCI systems over the restoration horizon, (ii) in a decentralized decision-making environment, protocols for determining the 'best' plan for the restoration efforts of an individual infrastructure, and (iii) the improvement in terms of restoration effectiveness that results in information-sharing among infrastructures in decentralized IIR efforts.

The remainder of the paper is organized as follows. Section 2 surveys the relevant literature related to IIR efforts. Section 3 describes the centralized model for IIR efforts and includes how to incorporate

the new classes of restoration interdependencies described in Sharkey et al. [40] that link the restoration efforts of multiple infrastructures. Section 4 discusses how to algorithmically model the different levels of decentralization and information-sharing among infrastructures in IIR efforts and Section 5 provides a computational analysis of the impact of these levels on IIR efforts. The paper concludes in Section 6.

2 Background and Previous Work

The focus of this section is to provide the background associated with research on infrastructure restoration and related topics in order to build appropriate models and algorithms for interdependent infrastructure restoration (IIR) efforts. Infrastructure restoration efforts are concerned with restoring the disrupted services resulting from damage to the infrastructure(s) caused by a disaster. Therefore, from a modeling perspective, it is important to capture the ‘performance’ of an infrastructure given a set of operational components. This step can be accomplished by modeling the infrastructure as a network where flow in the network correspond to the services provided by the it (see Ahuja et al. [1] for an overview of network flows). In other words, the flow into a demand node models the services provided to that point in the infrastructure and unmet demand can be viewed as a loss of performance.

Lee et al. [24, 25] provide an *interdependent layered network* model to measure the performance of a set of interdependent infrastructure systems given the current operational components. In particular, they capture the flow of services from supply points to demand points in each infrastructure and describe how to model *operational interdependencies* between the infrastructures. An important operational interdependency is an *input interdependency* where a component in infrastructure B requires the services of infrastructure A to function. For example, power is needed for a wastewater treatment plant to ensure proper operations. The infrastructure A is often referred to as a ‘parent’ and infrastructure B as a ‘child’ for this particular input interdependency. In order to model an input interdependency, a binary variable is defined that captures whether the appropriate level of services (i.e., met demand) in infrastructure A is delivered to the node representing the component for infrastructure B . If this variable is zero, then the flow through the component (node or arc) in infrastructure B must be zero. We will apply the interdependent layered network model of Lee et al. [24, 25] in order to capture the performance throughout the set of interdependent infrastructures over the restoration horizon.

The restoration efforts of an infrastructure involve allocating scarce resources, such as work crews, to repair damaged components or install temporary ones in order to re-establish the disrupted services caused by the disaster. The infrastructure is determining the schedule of when components will be repaired or installed or, equivalently, determining the set of operational components in the infrastructure over time. This fact means that models of restoration efforts have both a *network design* aspect (i.e., which components in the network are operational) and a *scheduling* aspect (i.e., when will components become operational). There has been recent research on these *integrated network design and scheduling* (INDS) problems, especially ones that focus on the cumulative performance of the network over the horizon of the problem.

Nurre and Sharkey [35] provide complexity results and dispatching rules for INDS problems whose network performance are measured by solving ‘classic’ network optimization problems including maximum

flow, minimum cost flow, shortest path, and minimum spanning tree problems. This built on the work of Nurre et al. [34] that focused on INDS models to restore infrastructure systems. Elgindy et al. [12] and Kalinowski et al. [18] provide complexity analysis and approximation algorithms for special cases of INDS problems where each component requires a unit processing time to repair or install with the shortest path and maximum flow performance metrics, respectively. Averbakh [3] and Averbakh and Pereira [4] consider a problem that focuses on installing arcs into a network to minimize the *recovery time* of each node, which is defined as the time where a path exists to that node from a hub. Guha et al. [15], Ang [2], Xu et al. [44], and Coffrin et al. [9] examine models that help to restore the power infrastructure. Matisziw et al. [30] examine an INDS model for an infrastructure that is concerned with the connectivity between nodes in the infrastructure. Cavdaroglu et al. [8] consider a model that focuses on determining the restoration efforts of infrastructures while being concerned with the input interdependencies that exist among infrastructures and whose performance focuses on the services in all systems. However, the computational testing of Cavdaroglu et al. [8] only focuses on damage to a single infrastructure and, therefore, does not consider the more complicated disruptive events that cause widespread damage across infrastructures. Our research examines the important issue of how decentralized decision-making impacts restoration of services across interdependent infrastructures.

This work builds upon the model of Cavdaroglu et al. [8] by modeling how the restoration decisions of infrastructures are linked across them in a ‘centralized’ planning model (see Section 3) and further considering damage scenarios that require restoration decisions for all infrastructures. The ‘core’ restoration efforts of each infrastructure are captured in a similar framework as the one proposed in Nurre et al. [34] for single infrastructure restoration. However, we incorporate new constraints that link the restoration efforts of multiple infrastructure systems to capture so-called *restoration interdependencies*. Sharkey et al. [40] introduced this new concept by reporting on incidents observed after Hurricane Sandy that link the restoration efforts of multiple infrastructures. From a scheduling perspective, many of these classes of restoration interdependencies are similar to the concept of precedence constraints (see Pinedo [37]) that restrict the processing of task j until all tasks that have a precedence over j are complete. Table 1 provides an overview of the different precedence-type classes discussed in Sharkey et al. [40].

Class	Definition	Example
Traditional Precedence	A task in infrastructure A cannot begin until a task in infrastructure B is complete.	Debris clearance on a road cannot begin until downed power lines are inspected.
Effectiveness Precedence	A task in infrastructure A can be processed more effectively after a task in infrastructure B is complete.	Floodwaters can be removed quicker from a tunnel after power is restored to pumps in the tunnel.
Options Precedence	A restoration task in an infrastructure can be completed by finishing a task in one of a set of possible infrastructures.	The reopening process of a gas station can be completed by either having power restored to it or bringing a generator to it.
Time-Sensitive Options	A restoration task in infrastructure A must be performed if a restoration task in infrastructure B is not completed by a certain deadline.	A telecommunications work crew must refuel a generator powering a cell tower if power is not restored to the tower within a certain timeframe.

Table 1: Classes and examples of restoration interdependencies.

3 An Interdependent Infrastructure Restoration Model with Restoration Interdependencies

The purpose of this section is to provide a mixed-integer programming formulation that can help to formulate the restoration efforts of multiple interdependent infrastructures while specifically considering both their operational and restoration interdependencies. In particular, we discuss and then formulate the *interdependent integrated network design and scheduling* (IINDS) problem. This problem determines: (1) for each infrastructure, the set of components that will be repaired or installed in the infrastructure network (i.e., the design decisions), (2) the assignment of selected components to available work groups and the time each work group will complete the tasks assigned to them (i.e., the scheduling decisions), and (3) in each time period, the flow of services through the set of interdependent infrastructure networks based on the current operational components in the set of networks (i.e., the interdependent flow decisions). Section 3.1 provides the overview of the model of Cavdaroglu et al. [8] to the problem of determining the restoration efforts in multiple infrastructures while also providing a general overview of the basic IINDS problem. Section 3.2 then discusses how to extend this core formulation to consider the various classes of restoration interdependencies. This IINDS problem can be viewed as modeling the decision-making in a fully-centralized environment in responding to the disruptions caused by the disaster.

3.1 The Base IINDS Problem

The focus of this section is to provide the formulation of the IINDS problem similar to the one examined in Cavdaroglu et al. [8] (where the main difference is the objective function). We provide the details of the formulation here since we will rely on many of the decision variables in both formulating the constraints for the restoration interdependencies and in modeling the various decision-making environments. The objective of this problem is to maximize the level of overall community resilience by restoring the functionality of the set of interdependent infrastructure systems throughout the planning horizon of the problem. In particular, the *performance* of the set of systems is measured as a function of the demand that can be met throughout the systems at time $t = 1, \dots, T$ and the objective is interested in maximizing the cumulative performance of the systems over the T time periods in the horizon. Therefore, our objective implicitly seeks to minimize the services that are disrupted or “down” for long periods of time after an extreme event.

Table 2 provides an overview of the notation, variables, and their definitions for the *infrastructure operations* at time t in the problem. For ease of presentation, we assume that the each damaged component within an infrastructure is represented as an arc and do not include details for networks with multiple commodities (this would require adding another index on supply, demand, and flow). The first assumption is without loss of generality since, if a node i was damaged, we can use standard network expansion techniques to represent this node as two nodes, i' and i'' , and an arc (i', i'') where all incoming arcs into i enter i' and all outgoing arcs from i leave i'' . We make a similar assumption (without loss of generality) in the sense that the input interdependencies occur at the nodes of the child infrastructure - if an arc had an input interdependency, we could split the arc into two arcs with a node in the middle. An input interdependency

between node i in parent infrastructure m and node j in child infrastructure n means that $(i, j) \in F(m, n)$. The binary variable $y_{m,i}^{n,j,t}$ is then equal to 1 if node i has its demand met in infrastructure m and, therefore, node j is operational in infrastructure n . Note further that the terms f_{DF}^m and f_{NR}^m will help to determine the percentage of disrupted services restored at each time period for infrastructure m and will appear in the objective function.

Notation for Infrastructure Operations*	
M	The set of all infrastructures considered.
N^m	The set of all nodes in infrastructure m .
S^m	The set of all supply nodes in infrastructure m .
T^m	The set of all transshipment nodes in infrastructure m .
D^m	The set of all demand nodes in infrastructure m .
E^m	The set of all arcs in infrastructure m initially available.
\bar{E}^m	The set of all arcs in infrastructure m that can be installed into the network.
s_i^m	The amount of supply available at node $i \in S^m$ in infrastructure m .
d_i^m	The amount of demand at node $i \in D^m$ in infrastructure m .
w_i^m	The weight associated with meeting one unit of demand at node $i \in D^m$ in infrastructure m .
u_i^m	The capacity of node i in infrastructure m .
u_{ij}^m	The capacity of arc (i, j) in infrastructure m .
$F(m, n)$	The set of all parent/child node pairs in parent infrastructure m and child infrastructure n .
f_{DF}^m	The total amount of services met in damage-free infrastructure m where all components are operational.
f_{NR}^m	The total amount of services met in infrastructure m with no restoration efforts, i.e., immediately after the disruptive event.
*Note that m and n will be used to denote an infrastructure with n typically used to denote the child infrastructure in an interdependent relationship.	
Variables for Infrastructure Operations	
x_{ijt}^m	The amount of flow on arc $(i, j) \in E^m \cup \bar{E}^m$ in infrastructure m at time t .
v_{it}^m	The amount of demand met at node $i \in D^m$ in infrastructure m at time t .
$y_{m,i}^{n,j,t}$	A binary variable for $(i, j) \in F(m, n)$ representing whether sufficient demand is met at node i in infrastructure m so that node j in infrastructure n is operational in time t .

Table 2: Relevant notation and variables for the infrastructure operations at time t .

Table 3 provides an overview of the notation, variables, and their definitions for the aspects of the IINDS problem that focus on the allocation of work crews to install arcs into the infrastructure networks. The formulation of the IINDS problem will ensure that each available work crew is only processing one arc at a time. The two types of restoration decision variables are those that determine whether an arc was *just completed* by a work group in time period t (α_{kijt}^m) and those that keep track of whether an arc is available in t (β_{ijt}^m) implying that it was completed by a work group at or before t .

Notation for Restoration Scheduling

T	The total number of time periods in the restoration horizon.
ω_t	The weight (or importance) of the network performance in time period t .
K^m	The total number of work crews available to install arcs into infrastructure m .
p_{ij}^m	The processing time for arc $(i, j) \in \bar{E}^m$ in infrastructure m .
Variables for Restoration Scheduling	
α_{kijt}^m	A binary variable that is equal to 1 if arc $(i, j) \in \bar{E}^m$ in infrastructure m is completed by work crew k at time t .
β_{ijt}^m	A binary variable that is equal to 1 if arc $(i, j) \in \bar{E}^m$ in infrastructure m is available at time t .

Table 3: Relevant notation and variables for scheduling the restoration efforts of the infrastructures.

We are now in a position to present the initial IINDS problem (similar to the one presented in Cavdaroglu et al. [8]). In particular, the formulation is:

$$\text{maximize } \sum_{t=1}^T \omega_t \sum_{m \in M} \frac{\sum_{i \in D^m} w_i^m v_{it}^m - f_{NR}^m}{f_{DF}^m - f_{NR}^m}$$

subject to

(IINDS)

$$\sum_{(i,j) \in E^m \cup \bar{E}^m} x_{ijt}^m - \sum_{(j,i) \in E^m \cup \bar{E}^m} x_{jit}^m \leq s_i^m \quad t = 1, \dots, T, \forall i \in S^m, \forall m \in M \quad (1)$$

$$\sum_{(i,j) \in E^m \cup \bar{E}^m} x_{ijt}^m - \sum_{(j,i) \in E^m \cup \bar{E}^m} x_{jit}^m = 0 \quad t = 1, \dots, T, \forall i \in T^m, \forall m \in M \quad (2)$$

$$\sum_{(i,j) \in E^m \cup \bar{E}^m} x_{ijt}^m - \sum_{(j,i) \in E^m \cup \bar{E}^m} x_{jit}^m = -v_{it}^m \quad t = 1, \dots, T, \forall i \in D^m, \forall m \in M \quad (3)$$

$$0 \leq v_{it}^m \leq d_i^m \quad t = 1, \dots, T, \forall i \in D^m, \forall m \in M \quad (4)$$

$$0 \leq \sum_{(j,i) \in E^m \cup \bar{E}^m} x_{jit}^m \leq u_i^m \quad t = 1, \dots, T, \forall i \in T^m, \forall m \in M \quad (5)$$

$$0 \leq x_{ijt}^m \leq u_{ij}^m \quad t = 1, \dots, T, \forall m \in M, \forall (i, j) \in E^m \quad (6)$$

$$0 \leq x_{ijt}^m \leq u_{ij}^m \beta_{ijt}^m \quad t = 1, \dots, T, \forall m \in M, \forall (i, j) \in \bar{E}^m \quad (7)$$

$$d_i^m - v_{it}^m \leq (1 - y_{m,i}^{n,j,t})(d_i^m) \quad t = 1, \dots, T, \forall (i, j) \in F(m, n) \text{ with } j \in N^n \text{ and } i \in D^m \quad (8)$$

$$\sum_{(j,h) \in E^m \cup \bar{E}^m} x_{jht}^n \leq s_j^n y_{m,i}^{n,j,t} \quad t = 1, \dots, T, \forall (i, j) \in F(m, n) \text{ with } j \in S^n \text{ and } i \in D^m \quad (9)$$

$$\sum_{(h,j) \in E^m \cup \bar{E}^m} x_{hjt}^n \leq d_j^n y_{m,i}^{n,j,t} \quad t = 1, \dots, T, \forall (i, j) \in F(m, n) \text{ with } j \in D^n \text{ and } i \in D^m \quad (10)$$

$$\sum_{(h,j) \in E^m \cup \bar{E}^m} x_{hjt}^n \leq u_j^n y_{m,i}^{n,j,t} \quad t = 1, \dots, T, \forall (i,j) \in F(m,n) \text{ with} \quad (11)$$

$$j \in T^n \text{ and } i \in D^m$$

$$\sum_{(i,j) \in \bar{E}^m} \sum_{s=t}^{\min\{T, t+p_{ij}^m-1\}} \alpha_{kij}^m \leq 1 \quad t = 1, \dots, T, \forall m \in M, k = 1, \dots, K^m \quad (12)$$

$$\beta_{ijt}^m - \beta_{ij(t-1)}^m = \sum_{k=1}^{K^m} \alpha_{kijt}^m \quad t = 2, \dots, T, \forall m \in M, \forall (i,j) \in \bar{E}^m \quad (13)$$

$$\alpha_{kijt}^m, \beta_{ijt}^m \in \{0, 1\} \quad t = 1, \dots, T, \forall m \in M, \forall (i,j) \in \bar{E}^m \quad (14)$$

The objective function of the IINDS problem focuses on the performance of the set of interdependent infrastructure networks as measured by the average percentage of disrupted services restored across infrastructures. For infrastructure m in time period t , the numerator of the term in the objective measures the current amount of restored (weighted) demand while the denominator provides the best possible amount of restored (weighted) demand. This objective essentially eliminates different scales and/or units of measurement of the services of an infrastructure (e.g., kW/unit time for power or gallons/unit time for wastewater and water systems) and thus captures the overall performance across infrastructures. The IINDS problem then focuses on maximizing the total performance of the infrastructures over the restoration horizon. Constraints (1), (2), (3) represent constraints on the flow in and out of supply, transshipment, and demand nodes, respectively. Constraints (4) ensure that we do not deliver more than the required demand at a node. Constraints (5), (6), and (7) are node and arc capacity constraints. Note that constraint (7) enforces that an arc in \bar{E}^m can only have flow on it if it has installed into the network by time t . Constraints (8)-(11) capture the interdependency of nodes $(i,j) \in F(m,n)$. In particular, if the flow into node i is less than the required demand, then $y_{m,i}^{n,j,t}$ will be zero and no flow will be allowed to enter and/or leave node j in infrastructure n . Constraints (12) ensure that work crew k is not processing more than one arc at a time. Constraints (13) ensure that an arc becomes available when a work crew finishes processing it.

Our focus is on analyzing and understanding the impact of the interdependencies that exist between infrastructures and, therefore, the IINDS problem captures those aspects at the expense of sacrificing some of the details of the operations of each individual infrastructure. For example, we have assumed a linear (network flow) representation of each infrastructure where the power infrastructure operates according to the laws of physics and could be more accurately modeled using a nonlinear formulation (see, for example, the discussion in Bienstock and Mattia [5] and the work of Hijazi et al. [16]). Despite this limitation, Nurre et al. [34] observed that this type of linear model does provide a good approximation for determining the *restoration planning* decisions. Further, based on observations of Nurre et al. [34], we have represented the scheduling decisions in the IINDS problem using a time-extended formulation.

It is important to note that once the $y_{m,i}^{n,j,t}$ decisions are fixed for infrastructure n , i.e., we know the ‘outages’ in the network in each time period due to the lack of services in parent infrastructures, then we can determine the restoration plan for infrastructure n . In other words, the knowledge of the $y_{m,i}^{n,j,t}$ results in infrastructure n being able to solve a single network version of the IINDS problem (similar to the INDS

problems studied in Nurre et al. [34] and Nurre and Sharkey [35]) to determine their best possible restoration plan given their interdependencies. In this problem, the infrastructure seeks to maximize the services it provides over the restoration horizon. This observation is important since it allows us to model decentralized decision-making in the IINDS problem, especially problems *without any restoration interdependencies*.

3.2 The IINDS Problem with Restoration Interdependencies

The purpose of this section is to describe how to model different classes of *restoration interdependencies* that were identified by Sharkey et al. [40] within the IINDS problem. For each class, we will provide a motivating example and then discuss any new notation and/or variables that are necessary to appropriately capture an incident of that class. In addition, new or modified constraints will be presented that should be incorporated into the IINDS problem to represent incidents of those restoration interdependencies.

3.2.1 Traditional Precedence

A traditional precedence relationship exists between restoration tasks in two infrastructures when a ‘task’ in infrastructure m must be completed before the task in infrastructure n can start its processing. The term ‘task’ in infrastructure m is used to capture both of the following situations: (i) a specific arc in \bar{E}^m must be restored before the arc in infrastructure n can begin and (ii) services in infrastructure m must be restored to a node in order to allow for the processing of the arc in infrastructure n . Examples of situation (i) are when a power crew must inspect a downed power line prior to debris removal on a road and when debris must be removed from a road prior to repairing a power line running along side of it. An example of situation (ii) is when power must be restored to a node that feeds a subway line prior to running test trains to test the quality of the repairs to the subway line. We refer to situation (i) as an *arc-based* precedence constraint and (ii) as a *node-based* precedence constraint. Table 4 provides the new notation and variables necessary to capture these precedence constraints.

Notation for Traditional Precedence	
ATP	The set of arc-based precedence constraints. For $((i, j), m), ((a, b), n) \in ATP$, arc (i, j) in infrastructure m must be completed before arc (a, b) in infrastructure n can be processed.
NTP	The set of node-based precedence constraints. For $((i, m), ((a, b), n) \in NTP$, node i in infrastructure m must have restoration demand met while arc (a, b) in infrastructure n is being processed.
μ_{nab}^{mi}	The amount of demand required at node i in infrastructure m in a time period for the purposes of restoring arc (a, b) in infrastructure n .
Variables for Traditional Precedence	
ν_{nabt}^{mi}	The amount of demand met in time t at node i in infrastructure m for the purposes of restoring arc (a, b) in infrastructure n .
ξ_{nabt}^{mi}	A binary variable that is equal to 1 if node i in infrastructure m at time t has sufficient demand met for the restoration of arc (a, b) in infrastructure n .

Table 4: Relevant notation and variables for the two classes of traditional precedence constraints.

Traditional precedence relations are modeled in constraints (15)-(18). The arc-based precedence constraints are modeled by ensuring that task (i, j) is completed at least p_{ab}^n (i.e., the processing time of arc (a, b) in infrastructure n) units before arc (a, b) (see constraints (15)). The node-based precedence constraints require us to replace constraints (3) with constraints (16) which include the restoration demands of tasks that have a traditional precedence relationship with node i . Constraints (17) ensure that the restoration demand is met at node i for task (a, b) if the binary variable ξ_{nabt}^{mi} is equal to 1. Constraints (18) then ensure that the restoration demand for task (a, b) is met during each time period it is being processed.

$$\sum_{k=1}^K \alpha_{kabt}^n \leq \beta_{ij,t-p_{ab}^n}^m \quad t = p_{ab}^n + 1, \dots, T, \quad ((i, j), m), ((a, b), n) \in ATP \quad (15)$$

$$\sum_{(i,j) \in E^m \cup \bar{E}^m} x_{ijt}^m - \sum_{(j,i) \in E^m \cup \bar{E}^m} x_{ijt}^m = -v_{it}^m - \sum_{\substack{((i,m), ((a,b),n)) \\ \in NTP}} v_{nabt}^{mi} \quad t = 1, \dots, T, \forall i \in D^m, \quad \forall m \in M \quad (16)$$

$$\xi_{mabt}^{ni} \leq \frac{1}{\mu_{nab}^{mi}} v_{nabt}^{mi} \quad t = 1, \dots, T, \quad \forall ((i, m), ((a, b), n)) \in NTP \quad (17)$$

$$\sum_{k=1}^{K^n} \alpha_{kabt}^n \leq \frac{1}{p_{ab}^n} \sum_{s=t-p_{ab}^n+1}^t \xi_{nabs}^{mi} \quad t = p_{ab}^n, \dots, T, \quad \forall ((i, m), ((a, b), n)) \in NTP \quad (18)$$

3.2.2 Effectiveness Precedence

An effectiveness precedence relationship exists between restoration tasks in two infrastructures if the failure to complete a ‘task’ in infrastructure m before a task in infrastructure n starts its processing results in the task in n being processed at a slower rate than normal. An example of such a relationship would be a situation in which if power is not restored to pumps near a road, then clearing floodwater would require a longer time. We again distinguish the situations between arc-based and node-based effectiveness precedence relationships. Table 5 provides the relevant definitions and notation for this class of restoration interdependency.

In order to capture the effectiveness precedence relationships, we need to modify constraints (12)-(13) in order to incorporate the possibility that arc (i, j) may become operational by being processed at its extended speed if it is a part of a relationship in AEP or NEP . Constraints (19)-(20) present this modification. The key in modeling the effectiveness precedence relationships is to realize that the ‘normal speed’ processing task of arc (a, b) has a *traditional* precedence constraint with the task in the other infrastructure. For arc-based effectiveness precedence relationships, similar constraints as constraints (15) are incorporated. For node-based effectiveness precedence relationships, we need to include the v_{nabt}^{mi} into the right hand side of node i 's flow balance constraint, which is similar to what was done for NTP in constraints (16). We then

Notation for Effectiveness Precedence

AEP	Set of arc-based effectiveness precedence relationships. For $((i, j), m), ((a, b), n)) \in AEP$, arc (i, j) in infrastructure m must be completed prior to processing (a, b) in infrastructure n at its normal speed.
NTP	Set of node-based effectiveness precedence relationships. For $((i, m), ((a, b), n)) \in NEP$, node i in infrastructure m must have restoration demand met while arc (a, b) in infrastructure n is being processed at its normal speed.
e_{ab}^n	The extended processing time for arc (a, b) in infrastructure n .

Variables for Effectiveness Precedence

α_{kabt}^{ne}	A binary variable that is equal to 1 if arc $(a, b) \in \bar{E}^n$ in infrastructure n is completed by work crew k at time t while being processed at its extended processing time.
ν_{nabt}^{mi}	The amount of demand met in time t at node i in infrastructure m for the purpose of restoring arc (a, b) in infrastructure n at its normal speed.
ξ_{mabt}^{mi}	A binary variable that is equal to 1 if node i in infrastructure m at time t has sufficient demand met for the restoration of arc (a, b) at its normal speed in infrastructure n .

Table 5: Relevant notation and variables for the two classes of effectiveness precedence constraints.

include constraints similar to constraints (17)-(18) to relate the normal processing speed of arc (a, b) with meeting the restoration demand at node i in infrastructure m .

$$\sum_{(a,b) \in \bar{E}^n} \sum_{s=t}^{\min\{T, t+p_{ab}^n-1\}} \alpha_{kabs}^n \quad k = 1, \dots, K^n, t = 1, \dots, T, \quad (19)$$

$$+ \sum_{\substack{((i,j),m),((a,b),n) \in AEP, \\ ((i,m),((a,b),n)) \in NEP}} \sum_{s=t}^{\min\{T, t+e_{ab}^n-1\}} \alpha_{kabs}^{ne} \leq 1 \quad \forall n \in M$$

$$\beta_{abt}^n - \beta_{ab(t-1)}^n = \sum_{k=1}^{K^n} \alpha_{kabt}^n + \sum_{k=1}^{K^n} \alpha_{kabt}^{ne} \quad t = 2, \dots, T, \forall n \in M, \quad (20)$$

$$\forall (a, b) \in \bar{E}^n$$

3.2.3 Options Precedence

An options precedence relationship exists for restoration task (a, b) in infrastructure n if there is a set of restoration tasks for which one task from this set must be completed prior to the processing of arc (a, b) . In other words, for task $(a, b) \in \bar{E}^n$, there is an arc set $((i, j), m) \in E_{abn}$ and a node set $(i, m) \in N_{abn}$ that represents the set of precedence options and at least one of these options (across both sets) must be completed prior to the processing of arc (a, b) . As an example, the reopening process of a gas station (or grocery store) can begin after either power is restored to the station or a portable generator is brought in and installed at the station. In order to capture an options precedence relationship, we need to capture that some option was completed prior to starting task (a, b) . For example, if all of the options precedence relationships for (a, b) were arc-based (so $N_{abn} = \emptyset$), constraints (21) establish this connection. For nodes in N_{abn} , we

can introduce a new binary variable at time t that captures whether they have sufficient demand met for the processing of task (i, j) for its processing time (which essentially captures the right hand side of constraints (18)) and then alter constraints (21) to sum over the β variables for the arc-based options and the new variables for the node-based options. We omit the detailed mathematical formulation of this combination since our computational testing does not contain options precedence relationships.

$$\sum_{k=1}^{K^n} \alpha_{kabt}^n \leq \sum_{((i,j),m) \in E_{abn}} \beta_{ij,t-p_{ab}^n}^m \quad t = p_{ab}^n + 1, \dots, T \quad (21)$$

3.2.4 Time-Sensitive Options

A time-sensitive options relationship exists between two restoration tasks in different infrastructures when one task must be completed before a certain deadline *or the other task* must be completed. Examples of time-sensitive options relationships include: (i) when power must be restored to a wastewater treatment plant by a certain point or a clean-up task must be done at the plant and (ii) when power must be restored to a cell tower prior to its back-up generator running out of fuel or a refueling task must be accomplished. The approach to modeling time-sensitive options relationships is to include binary variables (or leverage existing ones) to determine whether the initial task is completed by its deadline and, if it is not, then force the other restoration task to be completed should the associated component be operational in its network. We focus on presenting the modeling approaches for the time-sensitive clean-up tasks and time-sensitive refueling tasks. Table 6 presents the relevant notation and variables for this class of interdependencies.

Notation for Time-Sensitive Options Relationships	
<i>CLN</i>	The set of time-sensitive clean-up tasks. For $((a, m), ((i, j), n)) \in CLN$, node a in infrastructure m must have its demand met before some deadline or a clean-up task must be performed on (i, j) .
<i>REF</i>	The set of time-sensitive refueling tasks. For $(j, n) \in REF$, node j in infrastructure n may be operational through the use of a generator rather than having its demand met in the power infrastructure.
θ_{ij}^n	The last period in which node a in infrastructure m can have its demand met before its associated clean-up task (i, j) in infrastructure n must be completed.
ψ_j^n	The number of time periods node j in infrastructure n with a generator can function before needing to be refueled.
Variables for Time-Sensitive Options Relationships	
α_{jt}^n	A binary variable indicating whether node j in infrastructure n has its power demand met with an operational generator at time t .
ρ_{kjt}^n	A binary variable indicating whether node j in infrastructure n with a generator has been refueled by work crew k in time t .

Table 6: Relevant notation and variables for the two classes of time-sensitive options relationships.

We first describe the necessary modifications to capture the clean-up tasks. Note that we have assumed that for $((a, m), ((i, j), n)) \in CLN$ that $(a, i) \in F(m, n)$, i.e., node a is the node in the parent infrastructure m in order to have node i operational in the child infrastructure n . Recall that we will apply an expansion technique so that all flow through i will travel on arc (i, j) . We replace constraints (7) with constraints (22)-(23). Constraints (22) ensure that that arc (i, j) is operational only if node i is operational in infrastructure m up until the restoration deadline for services to node i (or, equivalently, parent node a). Constraints (23) then imply that arc (i, j) is operational beyond the deadline if either the clean-up task is complete ($\beta_{ijt}^n = 1$) or power was restored by the deadline ($y_{m,a}^{n,j,\theta_{ij}^n} = 1$). Constraints (24)-(25) ensure that clean-up tasks cannot begin until after the deadline and after power is restored.

$$0 \leq x_{ijt}^n \leq u_{ij}^n y_{m,a}^{n,i,t} \quad t = 1, \dots, \theta_{ij}^n, \forall ((a, m), ((i, j), n)) \in CLN \quad (22)$$

$$0 \leq x_{ijt}^n \leq u_{ij}^n (\beta_{ijt}^n + y_{m,a}^{n,j,\theta_{ij}^n}) \quad t = \theta_{ij}^n + 1, \dots, T, \forall ((a, m), ((i, j), n)) \in CLN \quad (23)$$

$$\alpha_{ijt}^n = 0 \quad \forall ((a, m), ((i, j), n)) \in CLN, t = 1, \dots, p_{ij}^n + \theta_{ij}^n \quad (24)$$

$$\beta_{ijt}^n \leq y_{m,a}^{n,j,(t-p_{ij}^n)} \quad t = \theta_{ij}^n + p_{ij}^n, \dots, T, ((a, m), ((i, j), n)) \in CLN \quad (25)$$

The modifications to capture refueling tasks include utilizing a binary variable to determine whether node j in infrastructure n is operational (o_{jt}^n), either through having its power met at node i in infrastructure m or having an operational generator. We then replace $y_{m,i}^{n,j,t}$ in the appropriate constraint describing the operations of node j in infrastructure n , i.e., constraints (9)-(11), with o_{jt}^n . The relationship between this binary variable and the interdependency and refueling decisions is captured in constraints (26) where, in order for $o_{jt}^n = 1$, power demand needs to be met at node i in infrastructure m or the generator needs to have been refueled recently. Constraints (27) would then extend constraints (12) by including the potential to allocate work group k to processing a refueling task.

$$o_{jt}^n \leq y_{m,i}^{n,j,t} + \sum_{k=1}^{K^m} \sum_{s=\max\{0, t-\psi_j^n+1\}}^t \rho_{kjt}^n \quad t = 1, \dots, T, k = 1, \dots, K^n, \forall (j, n) \in REF \quad (26)$$

$$\sum_{\substack{i \in N^n \\ (j,n) \in REF}} \rho_{kjt}^n + \sum_{(i,j) \in \bar{E}^n} \sum_{s=t}^{\min\{T, t+p_{ij}^n-1\}} \alpha_{kijst}^n \leq 1 \quad t = 1, \dots, T, \forall n \in M, k = 1, \dots, K^n \quad (27)$$

4 Modeling the Different Decision-Making Environments of Interdependent Infrastructure Restoration

The purpose of this section is to describe and model three distinct decision-making environments that may arise during interdependent infrastructure restoration. The resulting models and algorithms capturing each of the decision-making environments will then be tested to determine the level of effectiveness of them. In particular, we consider *centralized*, *decentralized*, and *information-sharing* decision-making environments.

4.1 Centralized Decision-Making Environment

The centralized decision-making environment is the ‘ideal’ situation for the IINDS problem in the sense that there is a single centralized authority that plans the restoration efforts of all infrastructures to maximize the overall restoration effectiveness. This environment can be captured by directly solving the IINDS problem (which would include the base model and any additional constraints required to capture the restoration interdependencies). The centralized environment, therefore, represents the best possible solution to the IINDS problem and will serve as a benchmark for the other environments. From a practical perspective, this environment would arise when there is a local emergency management team that includes representatives from all responding parties and whom have all agreed to fully cooperate with one another for the benefit of the area (potentially at the sacrifice of their own organization’s objective).

4.2 Decentralized Decision-Making Environment

The decentralized decision-making environment represents a situation in which no information or cooperation exists between the infrastructure networks in responding to the disaster. This situation then implies that each infrastructure will form its restoration efforts *independently* of one another. From a computational perspective, this decision-making environment can be modeled by having infrastructure n solve a version of the IINDS problem where we have ‘isolated’ n by fixing the binary interdependency variables $y_{m,i}^{n,j,t}$ and then focus on maximizing the services in the infrastructure. The infrastructure further assumes that they cannot process any tasks that are the ‘child’ of a traditional precedence relationship in planning their restoration efforts since they do not communicate with the other infrastructures (although the infrastructure can dynamically change their remaining restoration activities when new information is obtained about the completion of the parent precedence task). Essentially, in this scenario an infrastructure n makes assumptions about which components with interdependencies (both operational and restoration) will be available over the restoration horizon and then optimizes its restoration plan based on these assumptions. We explore both *optimistic* and *pessimistic* assumptions about the availability of these interdependent components.

We note that it is possible for infrastructure n to restore all services in its network without repairing all components in \bar{E}^n . This situation arises due to potential redundancies in the system. In reality, however, the infrastructure would still continue to install arcs (or make repairs) over the restoration horizon. Therefore, in the decentralized decision-making environment, we assume that the infrastructure optimizes its restoration plans and then applies another decision rule to ‘fill in’ the activities of its work groups for the remainder

of the horizon. In particular, we assume that the infrastructure will install (of those that are not already scheduled) arcs in the order of how close they are to supply points in the network. This assumption is consistent with the restoration guidelines of certain power companies (see, for example, Nurre et al. [34]). This step becomes especially important since certain infrastructures may not have an incentive (in terms of their IINDS problem) to process tasks that have precedence relationships with tasks in other infrastructures but we need to model the fact that, in reality, the infrastructure would continue to process tasks.

An infrastructure n is *optimistic* in terms of planning their restoration activities if they assume that the only reason components in their infrastructure will not be operational at time t is due to damage to the component itself. In other words, infrastructure n optimistically assumes that services in all other infrastructures will be restored immediately after the disaster so that $y_{m,i}^{n,j,t} = 1$ for all nodes j in infrastructure n and all time periods t . It is interesting to observe that the restoration plan of infrastructure n will not change over the restoration horizon for infrastructures with only *operational interdependencies* even if new observations are made about the availability of components with interdependencies. For instance, if some pump stations in the wastewater infrastructure are without power initially, an optimistic wastewater restoration plan would assume that these pump stations have power starting at time $t = 1$. If the manager of the wastewater system observed these pump stations were without power at time $t = 5$ and remained optimistic, there would be no incentive to change the restoration plan since the plan for the remainder of the horizon is optimal given the implemented restoration decisions from time $t = 1$ to $t = 5$. However, for an infrastructure with restoration interdependencies coming into its tasks (for example, a task in a different infrastructure must be completed prior to starting a task in the infrastructure), then there may be an incentive to change at time $t = 5$ since the new information may indicate that an unscheduled task can now be processed.

An infrastructure n is *pessimistic* in terms of planning their restoration activities if they assume that disrupted services from other infrastructures into the infrastructure will *never* be restored. In other words, in determining the restoration activities for the entire horizon initially, we assume that $y_{m,i}^{n,j,t} = y_{m,i}^{n,j,0}$ for all nodes j in infrastructure n and time periods t . For instance, if a pump station in the wastewater infrastructure is without power initially, a pessimistic wastewater restoration plan will assume power is disrupted to this pump station for the entire horizon. In this situation, it is important to note that the restoration plan of infrastructure n may change when new observations are made about the availability of components with interdependencies. If the manager of the wastewater system observed that at time $t = 0$ a pump station was without power but it had power at time $t = 5$, then we should update the restoration plan over the remainder of the horizon where we assume that the interdependency variable for that pump station is now ‘on’ from time $t = 5$ to T .

In terms of modeling the decentralized decision-making environment, it is necessary to capture the fact that the restoration plan for the *remainder of the horizon* may change when new information about other infrastructures’ restoration activities is obtained. In particular, we assume that each infrastructure n can observe at time t : (1) the set of completed tasks of all other infrastructures by t (so we observe β_{ijs}^m for $s \leq t$ for all $(i, j) \in \bar{E}^m$ and $m \in M$) and the set of current interdependency variables into infrastructure n or (2) the set of started tasks of all other infrastructures and the set of current interdependency variables. We refer to situation (1) as the infrastructure receiving information at the end of processing restoration tasks

and situation (2) as them receiving information at the start of processing restoration tasks. At this point, the infrastructure updates its restoration plan by solving the IINDS problem for the remainder of the horizon where any restoration decision implemented before t cannot be altered. In other words, we are in a non-preemptive environment in which the infrastructure must finish any task it had begun processing before t . This means that, if we have current restoration decisions $\alpha_{kij_s}^{m'}$ and are updating our restoration plan at time t , we enforce that $\alpha_{kij_s}^m = \alpha_{kij_s}^{m'}$ for $s = 1, \dots, t + p_{ij}^m - 1$. This captures that if we started processing arc (i, j) at or before time period t , we must complete it.

In order to determine the decentralized IIR plan (either under the optimistic or pessimistic settings), we would first determine the restoration plans for each individual infrastructure $n \in M$ by solving the reduced IINDS problem for just that infrastructure where we have fixed the $y_{m,i}^{n,j,t}$ variables according to the assumed settings. This determines the scheduling decision variables for each infrastructure n (the α_{kijt}^n and β_{ijt}^n) for the restoration horizon. We then begin to increase t and allow for the updating of the scheduling decision variables over the remainder of the horizon for each t when new information is obtained about the restoration efforts of the other infrastructures. When t reaches the end of the horizon, then we have determined the implemented restoration plan for each infrastructure n (we now know the α_{kijt}^n and β_{ijt}^n for all t). We then view these scheduling decision variables as fixed in the IINDS problem and apply the IINDS problem to evaluate the cumulative network performance of this plan.

4.3 Information-Sharing Decision-Making Environment

The information-sharing decision-making environment represents a situation in which infrastructures actively *share* their planned restoration efforts (and the operations within their own network) but do not necessarily coordinate or cooperate with one another. This environment is appropriate for situations when there is an emergency manager in the area that facilitates discussions among the decision-makers in the interdependent infrastructures but does not influence their decisions. In this environment, the restoration planning decisions are made at the beginning of the horizon (based on the announced plans of the other infrastructures) and there is no reason to alter these decisions since the known information will remain the same over the horizon of the problem.

The framework for this environment is that all infrastructures initially determine their restoration plans independently and then announce their plans to the other infrastructures. Each infrastructure n now has an updated view of the other infrastructures' activities over the restoration horizon and can update their planned restoration activities by solving another IINDS problem. However, at the same time, the other infrastructures are updating their planned restoration activities. Therefore, the information upon which the restoration activities of infrastructure n are based do not necessarily represent the true restoration plans of the other infrastructures. The announcement of the current restoration plans and then the updating of them based on the newly announced plans can continue for any number of 'rounds.' Therefore, we experiment with the number of times which infrastructures will announce their plans and update them to determine if the restoration plans tend to converge to a stable solution (i.e., one where no infrastructure has any incentive to change their planned restoration activities).

From a modeling perspective, each infrastructure n can plan their restoration efforts in a similar way as was done for the decentralized decision-making environment (i.e., we isolate the decisions of the infrastructure by fixing the interdependency variables and scheduling variables of all other infrastructures and solving the IINDS problem) but we now have that the $y_{m,i}^{n,j,t}$ will vary over t during the time horizon. We further note that infrastructure n will ‘fill in’ the activities of their work groups in a similar fashion as the decentralized environment should they reach a point where there are unscheduled tasks and idle work groups. The interesting aspect of the information-sharing environment is that the infrastructures may continue to adapt their restoration plans based on any changes to other infrastructures’ restoration plans.

5 Computational Results: The Price of Independent Decision-Making and the Value of Information-Sharing

This section presents a computational analysis to determine the empirical price of independent decision-making in the context of interdependent infrastructure restoration and the resulting improvement (or value) of information-sharing. This price is related to the concept of the *price of anarchy* (see, for example, Roughgarden [39]) but does not necessarily force the resulting restoration plans of the independent infrastructures to form stable solutions. Instead, we assume that the independent decision-makers determine their restoration plan according to the protocols described in Section 4.2. This computational analysis will be done by examining realistic data sets of infrastructure systems in New Hanover County, North Carolina in the United States (Section 5.1) and infrastructure systems of a realistically-constructed artificial community (Section 5.2). The former data set does not contain any restoration interdependencies in the damage scenario and, further, only has operational interdependencies from the power infrastructure to other infrastructures. This means that one can view the power system as the ‘lead’ infrastructure since the only uncertainty faced by other infrastructures in their restoration efforts is when components in their systems will have power restored. The Customizable ARTificial Community (CLARC) county data set (see Loggins et al. [28]) is more general in the sense that there is no longer a ‘lead’ infrastructure since power has restoration interdependencies both coming from it and going into it from other infrastructures. All tests were conducted on a laptop computer using IBM ILOG CPLEX Optimization Studio 12.2 to implement the algorithms and integer programming models discussed in Section 4. Although not the focus of this paper, we do note that all integer programs required less than 30 minutes of computational time to determine the optimal solution to the respective problem.

5.1 New Hanover County Analysis: Value of Information-Sharing with a Lead Infrastructure

The focus of this section is on examining the value of information-sharing and price of decentralized decision-making for the data set associated with New Hanover County, North Carolina in the United States. New Hanover County is a coastal county in southern North Carolina that includes the city of Wilmington and the Cape Fear beaches. Nurre et al. [34] use this data set in their work on single infrastructure restoration

efforts. This data set includes four main infrastructures: (i) the power infrastructure, (ii) the landline and mobile telecommunications infrastructures, (iii) the wastewater infrastructure, and (iv) the (potable) water infrastructure. Further, the data on each of these infrastructures was gathered through extensive collaborations with the managers of the infrastructure systems and the emergency manager of the county.

Table 7 provides the characteristics of each of these four infrastructures. It is important to note that the power network is composed of the *transmission* network and coarsely modeled distribution networks. In particular, the transmission network is composed of 39 nodes and 46 arcs while each demand node has an arc directly connecting a substation node in the transmission network to it (so we essentially have modeled the distribution network to this demand node with a single arc). Therefore, there is more redundancy in the transmission network than appears in the summary statistics. The telecommunications network is multi-commodity and contains a commodity for each origin/destination pair. The water infrastructure focuses on moving potable water from treatment facilities to demand points. There is, typically, no redundancy in a wastewater infrastructure and that is captured in this data set by the fact that the network is composed of five trees, one for each treatment plant in the infrastructure. Note that, in reality, flow moves from customers to treatment plants; however, for modeling purposes we view each customer as ‘demand’ and the treatment plants as supply points. This allows the model to capture which customers receive services in the infrastructure and appropriately weigh these services based on the customers.

Infrastructure	Nodes	Arcs	Supply Nodes	Demand Nodes
Power	472	489	8	443
Telecommunications	73	344	38	38
Water	305	398	58	206
Wastewater	543	538	5	322

Table 7: Data describing the infrastructures of New Hanover County.

Table 8 provides an overview of the operational interdependencies present in the data set. It should be noted that we only have input interdependencies from power to the other infrastructures, i.e., there are nodes and arcs in each of the three other infrastructures who need power in order to function in their own infrastructure. Given that these are the only interdependencies (including restoration interdependencies) in the damage scenario considered for this data set, we are in a situation where the power infrastructure can be viewed as the ‘lead’ infrastructure. In other words, the only situation linking the restoration efforts of these infrastructures is that certain components of the other three infrastructures need power to function and, therefore, once the power restoration plan is known, the decision-making can occur independently across the other three systems without sacrificing restoration effectiveness. Essentially, for the centralized decision-making situation, the optimal restoration efforts can be determined by solving the base IINDS model from Section 3.1.

We consider a damage scenario based upon a Category 3 hurricane that causes up to 120 miles per hour gust speed through the region (referred to as Cat-3 Scenario) and whose damage was determined using HAZUS (see the United States Federal Emergency Management Agency [13, 14]). The purpose of focusing on this type of scenario is that the damage is *realistic* in the sense that it is based on the vulnerability

Infrastructure	Telecommunications	Water	Wastewater
Power	35	49	153

Table 8: Size of $F(m, n)$, i.e., the number of input (operational) interdependences by pair of infrastructures where $m = \text{Power}$.

of different components within the infrastructure to an event that can hit the county as opposed to simply randomly generating damage scenarios. Therefore, the conclusions from the computational testing are based on a damage scenario that is realistic and could be faced by decision-makers. Table 9 provides an overview of the damage scenario.

Infrastructure	Damaged Nodes	Damaged Arcs	E^m
Power	0	19	19
Telecommunications	13	0	13
Water	12	0	12
Wastewater	25	0	25

Table 9: Data describing the Cat-3 Scenario for New Hanover County.

We are now in position to describe the computational analysis performed on the New Hanover County infrastructures and damage scenarios. The purpose of our testing is to determine the impact of decentralized decision-making in the context of the IINDS problem and how to reduce this impact by implementing the information-sharing mechanisms. In addition, we are interested in determining for the decentralized situation whether the optimistic or pessimistic planning assumptions result in stronger performance. Based on the objective function of the IINDS problem, each infrastructure’s restoration efforts are essentially ‘rated’ on a scale of zero to one where this rating provides the comparison of the actual restoration efforts with the ideal upper bound on the restoration efforts (i.e., we restore all services instantaneously and operate the network at ‘peak’ performance for the entire horizon). To be precise, we examine problems where the rating is equal to the *percentage* of restored services delivered over the restoration horizon compared to this best possible total of restored services. Therefore, our objective function of the IINDS problem will be the sum of these four ratings (one per infrastructure) and vary from zero to four.

The centralized decision-making environment is captured by solving the IINDS problem directly and represents the best situation for the recovery of the set of infrastructure systems from the disruptive event. For the decentralized decision-making environment, we consider both the optimistic view in creating infrastructure restoration efforts (Optimistic) and the pessimistic view (Pessimistic - End) where the activities of the other infrastructures are not known until they are completed (so, we do not know which arcs the other infrastructures are currently processing). For the pessimistic view, we also consider the impact of infrastructures sharing which tasks they are currently processing, which is equivalent to them announcing that they are processing an arc once it is started (we refer to this situation as Pessimistic - Start). For the information-sharing environment, recall that the only interdependencies linking the infrastructures and their restoration efforts are situations where an infrastructure component requires power to function properly. This situ-

ation means the information-sharing environment is essentially accomplished by the power infrastructure announcing their restoration plan. Once the power infrastructure determines and announces their restoration plan, each of the three other infrastructures can optimize their own restoration efforts independently of one another.

Table 10 provides the percentage sacrifice in restoration effectiveness resulting from each decision-making environment for this damage scenario. We consider this damage scenario with $K^m = 2$ work groups available to process restoration tasks per infrastructure and $T = 30$ time periods. This number of time periods is meant to represent approximately one week, where each time period represents a six-hour interval. The percentage sacrifice is defined essentially as the optimality gap of the decisions that result from each environment for that instance, i.e., it is calculated as

$$\frac{\text{Obj(Central)} - \text{Obj(Env)}}{\text{Obj(Central)}}$$

and thus provides a measure of the performance of that decision-making environment. For the decentralized decision-making environment, it is clear that formulating the restoration efforts under the optimistic assumption is better than the pessimistic restoration efforts, although the difference is small should the power infrastructure be willing to share information when it begins processing tasks. However, there is a significant improvement (decreases sacrifice in restoration effectiveness from 15.24% to 9.47% or a savings of over 37%) in moving from a decentralized decision-making environment to an information-sharing environment. This improvement demonstrates the need for emergency managers to facilitate discussions among infrastructure managers in restoration efforts after an event causing wide-spread damage.

Scenario	Centralized	Optimistic	Pessimistic-End	Pessimistic-Start	Info-Sharing
Cat-3 Scenario	0%	15.24%	19.72%	16.61%	9.47%

Table 10: The percentage sacrifice, in terms of restoration effectiveness, from the centralized decision-making environment for the Cat-3 damage scenario in New Hanover County.

5.2 CLARC County: Value of Information-Sharing with Restoration Interdependencies

The CustomizabLe ARTificial Community (CLARC) county was created to mimic a coastal county with a population of approximately 500,000 people. It was designed based on protocols observed in the infrastructure systems of New Hanover County, North Carolina and can be publicly shared, upon request, since it does not contain any sensitive information (see Loggins et al. [28]). The damage scenarios considered for CLARC county contain restoration interdependencies and, therefore, do not contain a ‘lead’ infrastructure like the scenario considered for New Hanover County.

Table 11 provides an overview of the power, telecommunications, transportation, and wastewater infrastructures of CLARC county. Similar to New Hanover County, the power network is composed of the *transmission* network and coarsely modeled distribution networks. In particular, the transmission network is composed of 62 nodes and 73 arcs while each demand node has an arc directly connecting a substation

node in the transmission network to it. The telecommunications network is multi-commodity and contains a commodity for each origin/destination pair. The transportation network focuses on ensuring that emergency services can move from their supply points (police stations, fire stations, and hospitals) and arrive at demand nodes representing population centers. There is no redundancy in the wastewater infrastructure of CLARC county - it is composed of eight trees, one for each treatment plant in the infrastructure. Table 12 provides the number of operational input interdependencies between each pair of infrastructures in the CLARC county data set.

Infrastructure	Nodes	Arcs	Supply Nodes	Demand Nodes
Power	838	849	4	776
Telecommunications	29	72	19	10
Transportation	698	1317	94	367
Wastewater	718	710	8	345

Table 11: Data describing the infrastructures of CLARC county.

Infrastructure	Telecommunications	Transportation	Wastewater
Power	29	373	224

Table 12: Size of $F(m, n)$, i.e., the number of input (operational) interdependencies by pair of infrastructures where $m = \text{Power}$.

We consider two distinct damage scenarios to the infrastructures of CLARC county in our computational testing. The damage to the components of the infrastructures was created using a vulnerability analysis based on HAZUS (see Loggins and Wallace [29]). We refer to the first damage scenario as ‘medium-scale’ and the second damage scenario as ‘large-scale’ based on the number of damaged components in each scenario compared to the expected number. Table 13 provides a breakdown on the number of damaged nodes and arcs in each damage scenario. Note that these are a subset of the nodes and arcs in the infrastructure and adding them together yields the number of arcs in \bar{E}^m (since we use a network expansion technique to translate a damaged node to a damaged arc). Therefore, for each damage scenario, the number of arcs in ‘ E^m ’ for the scenario is less than the amount listed in Table 11 since we remove the arcs in \bar{E}^m .

Table 14 discusses the number of restoration interdependencies for each damage scenario. The traditional precedence relationships for the power and transportation infrastructures were created by examining whether damaged arcs in the power and transportation infrastructures were close to each other (in particular, whether the power line was next to the road). For a pair of damaged arcs in the power and transportation infrastructures that were close to each other, we assumed that this represented a situation in which trees and other debris brought down a power line. Therefore, a power crew must inspect the line to make sure it is not ‘live’ (a task in the power infrastructure) before the trees and debris can be removed from the road (a task in the transportation infrastructure). This represents the (Transportation, Power) entries in Table 14. Further, the trees and debris must be cleared prior to repairing the power line, which represents the entries of (Power, Transportation). For the time-sensitive relationships, we note that these relationships are not

necessarily a function of the damage of an infrastructure because they have more to do with how long services are disrupted. For example, a cell tower’s generator will need to be refueled if the tower was without traditional power for a certain number of time periods. This is a function of both the damage done to the power infrastructure and its restoration process. Therefore, the number of relationships is constant across damage scenarios.

Infrastructure	Medium-Scale Scenario			Large-Scale Scenario		
	Damaged Nodes	Damaged Arcs	E^m	Damaged Nodes	Damaged Arcs	E^m
Power	5	3	8	14	14	28
Telecommunications	3	0	3	7	0	7
Transportation	0	183	183	0	276	276
Wastewater	41	0	41	76	0	76

Table 13: Data describing the damage scenarios for CLARC County.

Type	Infrastructures	Medium-Scale	Large-Scale
Traditional Precedence Relationships	Power, Transportation	1	10
	Transportation, Power	1	10
Time-Sensitive Cleanup Relationships	Power, Wastewater	52	52
Time-Sensitive Refueling Relationships	Power, Telecommunications	19	19

Table 14: Number of restoration interdependencies by type and infrastructures.

We are now in position to provide a computational analysis for the CLARC County data set and damage scenarios with restoration interdependencies similar to what we did for the New Hanover County data set and damage scenario. The main difference in terms of our analysis pertains to the information-sharing environment since each infrastructure (including power) has restoration interdependencies or operational interdependencies coming into it. As noted in Section 4.3, each infrastructure will adapt their restoration plans based on the announced restoration plans of other infrastructures. A *round* of information-sharing would involve each infrastructure announcing their restoration plan and then adapting it based on the restoration plans of the other infrastructures. In our testing, we will consider different number of information-sharing rounds, where $\text{InfoS}(\ell)$ refers to the restoration efforts after ℓ rounds.

Table 15 provides the percentage sacrifice in restoration effectiveness resulting from each decision-making environment. We consider both damage scenarios under 2 and 3 work groups available to process restoration tasks per infrastructure and $T = 30$ time periods in the restoration horizon. For the decentralized decision-making environment, it is clear that the best strategy is to have each infrastructure formulate *optimistic* restoration efforts. In each instance, the optimistic decentralized environment is better than the pessimistic decentralized environment (even when you have more information about when tasks are started). The information-sharing environment improves, somewhat significantly, upon the optimistic decentralized environment (in 3 of the cases, decreasing the sacrifice by over 40%). This observation is important in the sense that each infrastructure is still autonomous (i.e., it still makes its final restoration plan) but moves towards the centralized solution.

Scenario	K^m	Centralized	Optimistic	Pessimistic-End	Pessimistic-Start	InfoS(1)	InfoS(5)
Medium-Scale	2	0%	10.50%	15.98%	15.72%	3.13%	3.07%
	3	0%	14.20%	16.97%	16.95%	3.03%	2.91%
Large-Scale	2	0%	27.85%	55.81%	27.86%	17.71%	16.34%
	3	0%	29.41%	54.37%	31.36%	15.32%	16.81%

Table 15: The percentage sacrifice, in terms of restoration effectiveness, from the centralized decision-making environment for each damage scenario and number of work groups.

Recall that the information-sharing environment allows the infrastructures to announce their restoration plans and then adapt them based on the announced plans of the other infrastructures. We can allow any number of ‘rounds’ of this process but the interest lies in whether we have reached a stable solution, i.e., no infrastructure changes their solution between rounds. It turns out that only one of the four considered instances reaches a stable solution in 7 rounds. Table 16 provides the evolution of the percentage sacrifice of the information-sharing environment as a function of the number of rounds. It is clear that, while the individual restoration plans of an infrastructure may improve their objective by adapting to the announced plans, the fact that the other plans are changing too may result in a worse overall objective function. For the case of $K^m = 2$ and the medium-scale damage scenario, the restoration plans essentially stabilize after 3 rounds of information-sharing. Typically, infrastructure restoration efforts will need to be determined quickly after the disruptive event and, therefore, the number of information-sharing rounds will typically be small. Therefore, the traditional ‘stopping rule’ for the information-sharing environment will be determined by the number of rounds in IIR efforts. The results in Table 16 demonstrate that a small number of rounds does not necessarily imply that a stable solution across infrastructures in their restoration plans will be reached. An alternative stopping rule could be when a stable solution has been reached (which can be guaranteed if no infrastructure adapts their restoration plans between two consecutive rounds) or when the number of rounds reaches its limit.

Scenario	K^m	InfoS(1)	InfoS(2)	InfoS(3)	InfoS(4)	InfoS(5)	InfoS(6)	InfoS(7)
Medium-Scale	2	3.13%	3.10%	3.07%	3.07%	3.07%	3.07%	3.07%
	3	3.03%	3.00%	2.84%	2.82%	2.91%	2.93%	2.84%
Large-Scale	2	17.71%	17.72%	18.31%	17.91%	16.34%	16.45%	17.23%
	3	15.32%	19.44%	14.82%	15.97%	16.81%	14.81%	18.03%

Table 16: The percentage sacrifice, in terms of restoration effectiveness, across the number of rounds in an information-sharing environment

6 Conclusions

This paper focuses on examining the issues of decentralized decision-making and information-sharing in the restoration of interdependent networks after a disaster damages components and disrupts services in them. We present the interdependent integrated network design and scheduling (IINDS) problem that is

based on the work of Cavdaroglu et al. [8] and extend it to incorporate the new concept of *restoration interdependencies* that link the restoration efforts of multiple infrastructure networks. These restoration interdependencies are especially important in the context of interdependent infrastructure restoration since their impact may prevent scheduled work to begin at its planned time. Therefore, their impact in terms of restoration effectiveness may be greater in decentralized decision-making environments. We propose several algorithmic models to capture different decision-making environments and apply them to realistic damage scenarios for the case studies of New Hanover County and the artificial CLARC county. We found that the empirical ‘price’ (in terms of restoration effectiveness) of decentralized decision-making is high (at least 10% in all cases but upwards of 30% for many cases) and significant gains (cutting the loss in restoration effectiveness by 30-50%) can be made by having infrastructures simply share their planned restoration efforts with one another. The information-sharing environment is desirable since it allows the independent infrastructures to remain autonomous but allows them to gain more information about other infrastructures’ restoration efforts.

In terms of future research, it would be interesting to investigate how to measure the theoretical price of anarchy (see, for example, Roughgarden [39]) for interdependent network restoration. In this setting, each infrastructure would be viewed as an independent decision-maker in forming their restoration efforts. We would be interested in examining the ratio of the restoration effectiveness of any stable solution (or Nash equilibrium, i.e., a solution where no infrastructure has an incentive to change their restoration efforts) and the optimal restoration effectiveness resulting from centralized restoration efforts.

Acknowledgements

The work of Thomas Sharkey was supported in part by the U.S. National Science Foundation under grant number 1254258. The work of Thomas Sharkey, John Mitchell, and William Wallace was supported in part by the U.S. National Science Foundation under grant number 1314350. The work of Burak Cavdaroglu and William Wallace was supported by the U.S. Department of Homeland Security under Award Number: 2008-ST-061-ND 0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

References

- [1] R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network flows: Theory, algorithms, and applications*. Prentice-Hall, Englewood Cliffs, New Jersey, 1993.
- [2] C.C. Ang. Optimized recovery of damaged electrical power grids. Master’s thesis, Naval Postgraduate School, Monterrey, California, 2006.
- [3] I. Averbakh. Emergency path restoration problems. *Discrete Optimization*, 9(1):58–64, 2012.

- [4] I. Averbakh and J. Pereira. The flowtime network construction problem. *IIE Transactions*, 44(8):681–694, 2012.
- [5] D. Bienstock and S. Mattia. Using mixed-integer programming to solve power grid blackout problems. *Discrete Optimization*, 4:115–141, 2007.
- [6] A. Botterud, M.D. Ilic, and I. Wangensteen. Optimal investments in power generation under centralized and decentralized decision making. *IEEE Transactions on Power Systems*, 1:254–263, 20.
- [7] K. Caruson and S.A. MacManus. Disaster vulnerabilities: How strong a push toward regionalism and intergovernmental cooperation. *The American Review of Public Administration*, 38:286–306, 2008.
- [8] B. Cavdaroglu, E. Hammel, J.E. Mitchell, T.C. Sharkey, and W.A. Wallace. Integrating restoration and scheduling decisions for disrupted interdependent infrastructure systems. *Annals of Operations Research*, 203(1):279–294, 2013.
- [9] C. Coffrin, P. Van Hentenryck, and R. Bent. Strategic planning for power system restoration. In *Proceedings of the International Conference on Vulnerability and Risk Analysis and Management*, pages 180–187, 2011.
- [10] L.K. Comfort. Rethinking security: Organizational fragility in extreme events. *Public Administration Review*, 62:98–107, 2002.
- [11] S.S. Dawes. Interagency information-sharing: Expected benefits, managerial risks. *Policy Analysis and Management*, 15(3):377–394, 2007.
- [12] T. Elgindy, A. T. Ernst, M. Baxter, M. W. P. Savelsbergh, and T. Kalinowski. Incremental network design with shortest paths. Forthcoming in *European Journal of Operational Research*.
- [13] Federal Emergency Management Agency Mitigation Division. HAZUS-MH MR4 Flood Model technical manual. 2010.
- [14] Federal Emergency Management Agency Mitigation Division. HAZUS-MH MR4 Hurricane Model technical manual. 2010.
- [15] S. Guha, A. Moss, J.S. Naor, and B. Schieber. Efficient recovery from power outage. In *Proceedings of the thirty-first annual ACM symposium on Theory of Computing (STOC)*, pages 574–582, 1999.
- [16] H.L. Hijazi, C. Coffrin, and P. Van Hentenryck. Convex quadratic relaxations of mixed-integer non-linear programs in power systems. Technical report, NICTA, Canberra ACT, Australia, 2013.
- [17] Hurricane Sandy Rebuilding Task Force. Hurricane Sandy rebuilding strategy. U.S. Department of Housing and Urban Development., Washington, DC, 2013.
- [18] T. Kalinowski, D. Matsypura, and M.W.P. Savelsbergh. Incremental network design with maximum flows. Technical report, University of Newcastle, Australia.

- [19] T. Kaplan and R. Hernandez. Cuomo, in aid appeal, cites broad reach of storm, 2012. November 26, 2012. Accessed at http://www.nytimes.com/2012/11/27/nyregion/governor-cuomo-says-hurricane-sandy-was-worse-than-katrina.html?_r=0 on February 6, 2014.
- [20] N. Kapucu and V. Garayev. Designing, managing, and sustaining functionally collaborative emergency management networks. *The American Review of Public Administration*, 43:312–330, 2012.
- [21] S. Kaufman, C. Qing, N. Levenson, and M. Hanson. Transportation during and after Hurricane Sandy. Technical report, Rudin Center for Transportation, NYU Wagner Graduate School of Public Service, 2012.
- [22] A. Kenward, D. Yawitz, and U. Raja. Sewage overflows from Hurricane Sandy. Technical report, Climate Central, Princeton, New Jersey, 2013.
- [23] W.M. Leavitt and J.J. Kiefer. Infrastructure interdependency and the creation of a normal disaster: The case of Hurricane Katrina and the city of New Orleans. *Public Works Management & Policy*, 10:306–314, 2006.
- [24] E.E. Lee, J.E. Mitchell, and W.A. Wallace. Restoration of services in interdependent infrastructure systems: A network flows approach. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 37(6):1303–1317, 2007.
- [25] E.E. Lee, J.E. Mitchell, and W.A. Wallace. Network flow approaches for analyzing and managing disruptions to interdependent infrastructure systems. In John G. Voeller, editor, *Wiley Handbook of Science and Technology for Homeland Security*, volume 2, pages 1419–1428. John Wiley, 2009.
- [26] J.K. Lee and H. Rao. Exploring the causes and effects of inter-agency information sharing systems adoption in the anti/counter-terrorism and disaster management domains. In *Proceedings of the 8th Annual International Conference on Digital Government Research: Bridging Disciplines & Domains*, 2007.
- [27] R. Little. Controlling cascading failure: Understanding the vulnerabilities of interconnected infrastructures. *Journal of Urban Technology*, 9(1):109–123, 2002.
- [28] R. Loggins, J. Taylor, S. Brehm, and W.A. Wallace. *CLARC County*. Technical report, Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute, 2013.
- [29] R.A. Loggins and W.A. Wallace. Rapid assessment of hurricane damage disruptions to interdependent civil infrastructure systems. Technical report, Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute, 2014.
- [30] T.C. Matisziw, A.T. Murray, and T.H. Grubestic. Strategic network restoration. *Networks and Spatial Economics*, 10:345–361, 2010.

- [31] T. McDaniels, S. Chang, K. Peterson, J. Mikawoz, and D. Reed. Empirical framework for characterizing infrastructure failure interdependencies. *Journal of Infrastructure Systems*, 13:175–184, 2007.
- [32] T. McDaniels, S. Chang, and D. Reed. Characterizing infrastructure failure interdependencies to inform system risk. In John G. Voeller, editor, *Wiley Handbook of Science and Technology for Homeland Security*, volume 2, pages 1–16. John Wiley, 2009.
- [33] M. McGuire and D. Schneck. What if Hurricane Katrina hit in 2020? The need for strategic management of disasters. *Public Administration Review*, 70:201–207, 2010.
- [34] S.G. Nurre, B. Cavdaroglu, J.E. Mitchell, T.C. Sharkey, and W.A. Wallace. Restoring infrastructure systems: An integrated network design and scheduling problem. *European Journal of Operational Research*, 223(3):794–806, 2012.
- [35] S.G. Nurre and T.C. Sharkey. Integrated network design and scheduling problems with parallel identical machines: Complexity analysis and dispatching rules. Forthcoming in *Networks*, 2014.
- [36] Office of Electricity Delivery United States Department of Energy and Energy Reliability. Hurricane Sandy situation reports, 2012.
- [37] M.L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Springer, New York, New York, 2012.
- [38] S. M. Rinaldi, J. P. Peerenboom, and T. K. Kelly. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine*, 21(6):11–25, 2001.
- [39] T. Roughgarden. *Selfish routing and the price of anarchy*. MIT Press, Boston, MA, 2005.
- [40] T.C. Sharkey, S.G. Nurre, H. Nguyen, J.H. Chow, J.E. Mitchell, and W.A. Wallace. Identification and classification of restoration interdependencies in the wake of Hurricane Sandy. Technical report, Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute, 2014. Submitted to *Journal of Infrastructure Systems*.
- [41] S. Somers and J.H. Svara. Assessing and managing environmental risk: Connecting local government management with emergency management. *Public Administration Review*, 69:181–193, 2009.
- [42] United States Department of Homeland Security. National infrastructure coordinating center. <http://www.dhs.gov/national-infrastructure-coordinating-center>.
- [43] W.A. Wallace, D. Mendonca, E. Lee, J.E. Mitchell, and J. Chow. Managing disruptions to critical interdependent infrastructures in the context of the 2001 World Trade Center attack. *Beyond September 11th: An Account of Post-Disaster Research*, pages 165–198, 2003.
- [44] N. Xu, S.D. Guikema, R.A. Davidson, L.K. Nozick, Z. Cagnan, and K. Vaziri. Optimizing scheduling of post-earthquake electric power restoration tasks. *Earthquake Engineering and Structural Dynamics*, 36:265–284, 2007.