Last-Mile Restoration for Multiple Interdependent Infrastructures

Carleton Coffrin
Brown University
Providence, RI 02912
cjc@cs.brown.edu

Pascal Van Hentenryck
Optimization Research Group, NICTA
University of Melbourne, Australia
pvh@nicta.com.au

Russell Bent
Los Alamos National Laboratory
Los Alamos, NM 87545
rbent@lanl.gov

Abstract

This paper considers the restoration of multiple interdependent infrastructures after a man-made or natural disaster. Modern infrastructures feature complex cyclic interdependencies and require a holistic restoration process. This paper presents the first scalable approach for the last-mile restoration of the joint electrical power and gas infrastructures. It builds on an earlier three-stage decomposition for restoring the power network that decouples the restoration ordering and the routing aspects. The key contributions of the paper are (1) mixed-integer programming models for finding a minimal restoration set and a restoration ordering and (2) a randomized adaptive decomposition to obtain high-quality solutions within the required time constraints. The approach is validated on a large selection of benchmarks based on the United States infrastructures and state-of-the-art weather and fragility simulation tools. The results show significant improvements over current field practices.

Background and Motivation

Restoring critical infrastructure after a significant disruption (e.g., a natural or man-made disaster) is an important task with consequences on both human and economic welfare. Damaged components must be prioritized and repaired, to restore service as quickly as possible without causing additional instability. Last-mile restoration considers infrastructure damages at the city or the state scale and is particularly complex as it amounts to solving a pickup and delivery routing problem, whose objective function minimizes loss of service over time in an interdependent infrastructure. It contrasts with humanitarian relief efforts which are more concerned with effectively establishing one-time supply chains.

Last-mile restoration has attracted increased attention in recent years but the majority of the research is devoted to single infrastructures, e.g., the power network or potable water supply. However, modern infrastructures exhibit multiple, often cyclic, interdependencies. For instance, the gas network may fuel an electric generator or a gas compressor may consume electricity to increase the pressure in

pipelines. Therefore, it is critical to restore these infrastructures jointly to maximize the level of service over time.

This paper proposes the first last-mile restoration approach of multiple complex interdependent infrastructures. It uses mixed-integer programs (MIP) for modeling interdependent power and gas networks, combining the linearized DC model for the power network and a flow model for the gas network. The models are then integrated into the multistage last-mile restoration approach proposed in (Van Hentenryck, Coffrin, and Bent 2011) for the power network, which is used to advise federal agencies when hurricanes of category 3 or above approach the United States. The infrastructure interdependencies induce computational difficulties for MIP solvers in the prioritization step, which we address by using a randomized adaptive decomposition (RAD) approach. The RAD approach iteratively improves a restoration order by selecting smaller restoration subproblems which are solved independently. The proposed approach was evaluated systematically on a large collection of benchmarks generated with state-of-the-art hazard and fragility simulation tools on the infrastructure of the United States. The results demonstrate the scalability of the approach, which finds very high-quality solutions to large last-mile restoration problems and brings significant improvements over current field practices.

The rest of the paper describes the modeling of multiple interdependent infrastructures and our approach for last-mile restoration of such infrastructures. It presents the experimental results and concludes with a discussion of related work in restoration of interdependent infrastructures.

Infrastructure Modeling

Power and gas infrastructures can be modeled and optimized at various levels of abstraction. Linear approximations are typically used for applications involving topological changes, a design choice followed by this paper as well. This section presents a demand maximization model for interdependent power and gas infrastructures, which is a key building block for the restoration models.

The Power Infrastructure The power infrastructure is modeled in terms of the Linearized DC Model (LDCM), a standard tool in power systems (e.g., (Murillo-Sánchez and

Model 1 Power System Demand Maximization.

Inputs:
$\quad \mathcal{PN} = \langle B, L, s \rangle$ - the power network

Variables:
$\quad \theta_i \in (-\Re, \Re)$ - phase angle on bus i (rad)
$\quad E_i^g \in (0, \hat{E}_i^g)$ - power injected by generator i
$\quad E_i^o \in (0, \hat{E}_i^o)$ - power consumed by load i
$\quad E_i^l \in (-\hat{E}_i^l, \hat{E}_i^l)$ - power flow on line i

Maximize
$$\sum_{i \in B^o} E_i^o \tag{M1.1}$$

Subject to:
$$\theta_s = 0 \tag{M1.2}$$
$$\sum_{j \in B_i^o} E_j^o = \sum_{j \in B_i^g} E_j^g + \sum_{j \in LI_i} E_j^l - \sum_{j \in LO_i} E_j^l \quad \forall i \in B \tag{M1.3}$$
$$E_i^l = b_i(\theta_{L_i^+} - \theta_{L_i^-}) \quad \forall i \in L \tag{M1.4}$$

Gan 1997; Knight 1972; Wood and Wollenberg 1996; Powell 2004; Gomez-Exposito, Conejo, and Canizares 2008)). In the LDCM, a *power network* $\mathcal{PN}$ is represented by a collection of buses B and a collection of lines L connecting the buses. Each bus $i \in B$ may contain multiple generation units B_i^g and multiple loads B_i^o and $B^g = \bigcup_{i \in B} B_i^g$ and $B^o = \bigcup_{i \in B} B_i^o$ are used to denote the generators and loads across all buses. Each generator $j \in B^g$ has a maximum generation value $\hat{E}_j^g$ and each load $k \in B^o$ has a maximum consumption value $\hat{E}_k^o$. Each line $i \in L$ is assigned a *from* and *to* bus denoted by L_i^+ and L_i^- respectively and is characterized by two parameters: a maximum capacity $\hat{E}_i^l$ and a susceptance b_i. LO_j and LI_j denote all the lines oriented *from* or *to* a given bus j respectively. Lastly, one bus s is selected arbitrarily as the *slack-bus* to remove numerical symmetries. Model 1 presents a LDCM for maximizing the load of a power network $\mathcal{PN} = \langle B, L, s \rangle$. The decision variables are: (1) the phase angles of the buses θ; (2) the production level of each generator E^g; (3) the consumption level of each load E^o; (4) the flow on each line E^l which can be negative to model a flow in the reverse direction. The objective (M1.1) maximizes the total load served. Constraint (M1.2) fixes the phase angle of the slack bus. Constraint (M1.3) ensures flow conservation (i.e., Kirchhoff's Current Law) at each bus, and constraint (M1.4) ensures the line flows are defined by line susceptances.

The Gas Infrastructure   We use a network flow model for the gas system, which is also common in practice (e.g., (Carvalho et al. 2009; Monforti and Szikszai 2010)). The gas model is similar to the power model. A *gas network* $\mathcal{GN}$ is represented by a collection of junctions J and a collection of pipelines P connecting the junctions. Each junction $i \in J$ may contain multiple generation units J_i^g (aka well fields) and multiple loads J_i^o (aka city gates) and we define J^g and J^o as in the power system. Each generator $j \in J^g$ has a maximum generation value $\hat{G}_j^g$ and each load $k \in J^o$ has a maximum consumption value $\hat{G}_k^o$. Each pipeline $i \in P$ is associated with a *from* and *to* junction which are denoted by P_i^+

Model 2 Gas System Demand Maximization.

Inputs:
$\quad \mathcal{GN} = \langle J, P \rangle$ - the gas network

Variables:
$\quad G_i^g \in (0, \hat{G}_i^g)$ - gas injected by well field i
$\quad G_i^o \in (0, \hat{G}_i^o)$ - gas consumed by city gate i
$\quad G_i^p \in (-\hat{G}_i^p, \hat{G}_i^p)$ - gas flow on pipeline p

Maximize
$$\sum_{i \in J^o} G_i^o \tag{M2.1}$$

Subject to:
$$\sum_{j \in J_i^o} G_j^o = \sum_{j \in J_i^g} G_j^g + \sum_{j \in PI_i} G_j^p - \sum_{j \in PO_i} G_j^p \quad \forall i \in J \tag{M2.2}$$

and P_i^- respectively and a flow limit of $\hat{G}_i^p$. The sets PO_j and PI_j are defined as in the power system. Gas networks also have compressors which are denoted by the set PC. It is convenient to refer to the set of all pipelines attached to a compressor $i \in PC$ as P_i^c. The effects of compressors is only significant for the interdependent model and are a perfect example of why modeling interdependencies are so critical for finding high-quality restoration plans. Model 2 presents a linear program for maximizing the demand in a gas network. The inputs are a gas network $\mathcal{GN} = \langle J, P \rangle$ and the decision variables are: (1) the production level of each generator G^g; (2) the consumption level of each load G^o; (3) the flow on each pipeline G^p which can be negative as well. The objective (M2.1) maximizes the total loads served. Constraint (M2.2) ensures flow conservation at each junction. Independently, both models are linear programs (LP).

The Interdependent Power and Gas Infrastructure   The power and gas networks have different types of interdependencies. *Sink-source* connections are common. For example, a gas city gate G^o can fuel a gas turbine engine which is an electric generator E^g. *Sink-sink* connections also appear. For example, a city gate G^o requires some energy from a load E^o to regulate its valves. All of these interdependencies can be modeled in terms of implications $a \rightarrow c$ which indicate that consequent c is not operational whenever antecedent a is not served at full capacity. *Pipeline compressors* also induce fundamental interdependencies. Indeed, compressors consume electricity from a load E^o to increase the pressure on a pipeline P, as sufficient line pressure is a feasibility requirement for the gas network. This dependency is modeled as a capacity reduction, since pressure is not captured explicitly in the linear gas model.

An interdependent model is inherently multi-objective. In practice however, policy makers typically think of infrastructure restoration in terms of financial or energy losses. Both cases are naturally modeled as a linear combination of the power and gas objectives. The objectives only consider the set of loads in the networks which are not antecedent to a dependency. We use $T^e \subseteq B^o$ and $T^g \subseteq J^o$ to denote the filtered loads for the power and gas networks. If W^e and W^g are the weights of the infrastructures, then the joint objective is $W^e \sum_{i \in T^e} E_i^o + W^g \sum_{i \in T^g} G_i^o$. The maximal demand

Model 3 Interdependent Demand Maximization.

Inputs:

$\mathcal{PN} = \langle B, L, s \rangle$	- the power network
$\mathcal{GN} = \langle J, P \rangle$	- the gas network
A, A^c, C	- the interdependencies
T^e, T^g	- the demand points
W^e, W^g	- the demand weights

Variables:

$y_i \in \{0, 1\}$	- item i is activated
$z_i \in \{0, 1\}$	- item i is operational
$fl_i \in \{0, 1\}$	- all of item i's load is satisfied
$\theta_i \in (-\Re, \Re)$	- phase angle on bus i (rad)
$E_i^g \in (0, \hat{E}_i^g)$	- power injected by generator i
$E_i^o \in (0, \hat{E}_i^o)$	- power consumed by load i
$E_i^l \in (-\hat{E}_i^l, \hat{E}_i^l)$	- power flow on line i
$G_i^g \in (0, \hat{G}_i^g)$	- gas injected by well field i
$G_i^o \in (0, \hat{G}_i^o)$	- gas consumed by city gate i
$G_i^p \in (-\hat{G}_i^p, \hat{G}_i^p)$	- gas flow on pipeline p

Maximize

$$W^e \sum_{d \in T^e} E_d^o + W^g \sum_{d \in T^g} G_d^o \quad \text{(M3.1)}$$

Subject to:

$$y_i = 1 \quad \forall i \in N \setminus C \quad \text{(M3.2.1)}$$
$$fl_i \Leftrightarrow \hat{I}_i^o \le I_i^o \quad \forall i \in A \quad \text{(M3.2.2)}$$
$$y_i = \bigwedge_{j \in A_i^c} fl_j \quad \forall i \in C \quad \text{(M3.2.3)}$$

$$z_i = y_i \quad \forall i \in B \quad \text{(M3.3.1)}$$
$$z_i = y_i \wedge y_j \quad \forall j \in B, \forall i \in B_j^g \cup B_j^o \quad \text{(M3.3.2)}$$
$$z_i = y_i \wedge y_{L_i^+} \wedge y_{L_i^-} \quad \forall i \in L \quad \text{(M3.3.3)}$$
$$\theta_s = 0 \quad \text{(M3.3.4)}$$
$$\sum_{j \in B_i^o} E_j^o = \sum_{j \in B_i^g} E_j^g + \sum_{j \in LI_i} E_j^l - \sum_{j \in LO_i} E_j^l \quad \forall i \in B \quad \text{(M3.3.5)}$$
$$\neg z_i \rightarrow E_i^g = 0 \quad \forall i \in B^g \quad \text{(M3.3.6)}$$
$$\neg z_i \rightarrow E_i^o = 0 \quad \forall i \in B^o \quad \text{(M3.3.7)}$$
$$\neg z_i \rightarrow E_i^l = 0 \quad \forall i \in L \quad \text{(M3.3.8)}$$
$$z_i \rightarrow E_i^l = B_i(\theta_{L_i^+} - \theta_{L_i^-}) \quad \forall i \in L \quad \text{(M3.3.9)}$$

$$z_i = y_i \quad \forall i \in J \cup PC \quad \text{(M3.4.1)}$$
$$z_i = y_i \wedge y_j \quad \forall j \in J, \forall i \in J_j^g \cup J_j^o \quad \text{(M3.4.2)}$$
$$z_i = y_i \wedge y_{P_i^+} \wedge y_{P_i^-} \quad \forall i \in P \quad \text{(M3.4.3)}$$
$$\sum_{j \in J_i^o} G_j^o = \sum_{j \in J_i^g} G_j^g + \sum_{j \in PI_i} G_j^p - \sum_{j \in PO_i} G_j^p \quad \forall i \in J \quad \text{(M3.4.4)}$$
$$\neg z_i \rightarrow G_i^g = 0 \quad \forall i \in J^g \quad \text{(M3.4.5)}$$
$$\neg z_i \rightarrow G_i^o = 0 \quad \forall i \in J^o \quad \text{(M3.4.6)}$$
$$\neg z_i \rightarrow G_i^p = 0 \quad \forall i \in P \quad \text{(M3.4.7)}$$
$$\neg z_i \rightarrow -\tilde{G}_j^p \le G_j^p \le \tilde{G}_j^p \quad \forall j \in P_i^c \quad \forall i \in PC \quad \text{(M3.4.8)}$$

satisfaction of each network is often useful, we use M^e and M^g to refer to the maximum power and gas demand satisfaction respectively.

We are almost in a position to present the interdependent model. The missing piece of information is the recognition that, whenever a component is not active, it may induce other components to be non-operational as well. For example, if a bus is inactive, then all of the components connected to that bus (e.g., lines, generators, loads) become non-operational. These intra-network dependencies, which are modeled in terms of logical constraints, are not present in the demand maximization model for a single infrastructure. Computationally, they imply that demand maximization of interdependent infrastructures becomes a MIP model, instead of a LP. The complete demand maximization model for the interdependent power and gas infrastructure is presented in Model 3. For clarity, we use the logical constraints, not their linearizations which can be obtained through standard transformations. The inputs are specified in terms of following additional notations: N is the collection of all of the infrastructure components, i.e., $N = B \cup B^g \cup B^o \cup L \cup J \cup J^g \cup J^o \cup P \cup PC$; The sink-sink and sink-source interdependencies are specified by antecedent and consequent relations. The set A is the collection of all antecedent items and C is the set of all consequent items; for each consequent $i \in C$ the set $A_i^c \subseteq A$ denotes all antecedents of i. The collection of all load points in both infrastructures is $I^o = G^o \cup E^o$, and $\hat{I}_i^o$ is the maximum load of a resource $i \in I^o$. The model inputs are then given by the network $\mathcal{IN} = \langle \mathcal{PN}, \mathcal{GN}, A, A^c, C, T^e, T^g, W^e, W^g \rangle$. The variables include those described in Models 1 and 2 and the objective function (M3.1) was described earlier.

To model the effect of the interdependencies on the network topologies, a binary variable y_i is associated with each component $i \in N$ and denotes whether the component is active. Another variable z_i is associated with component i to denote whether component i is operational. Most of the y_i variables are set to one: only those affected by interdependencies may be zero, as per Constraint (M3.2.1). The antecedent i of a dependency is always a load point and is only operational when its load is at full capacity which is captured by binary variable fl_i and Constraint (M3.2.2). That is $fl_i = 1$ if and only if $\hat{I}_i^o \le I_i^o$. Constraint (M3.2.3) specifies that each consequent $i \in C$ is active if all of of its antecedents A_i^c are at full capacity. Constraint (M3.4.8) specifies the capacity reduction of a compressor-dependent pipeline $j \in P_i^c$ when its compressor $i \in PC$ is not operational, i.e., $z_i = 0$. Note that the regular operating capacity of pipeline j is $\hat{G}_j^p$, while its reduced capacity is $\tilde{G}_j^p$.

Constraints (M3.3.1–M3.3.9) model the power system. Constraints (M3.3.1–M3.3.3) describe which components are operational following the operational rules sketched out previously. Constraints (M3.3.4) and (M3.3.5) are from Model 1. Constraints (M3.3.6–M3.3.9) imposes restrictions on power flow, consumption, and production depending on the operational state: They ensure that a non-operational generator, load, or line cannot produce, consume, or transmit power. Constraints (M3.4.1–M3.4.8) model the gas system. The principles are the same as the power system, except for constraints (M3.4.8) which models the effects of non-operational compressors which were discussed previously.

Joint Infrastructure Repair and Restoration

The joint repair and restoration of an interdependent infrastructure is extremely challenging computationally. It is a multiple pickup and delivery vehicle routing problem, whose objective function is defined in terms of a series of demand maximization problems, one for each repair action. Each of these demand maximizations is a MIP, which leads

MULTI-STAGE-IRVRP($Network$ $\mathcal{IN}$, $IRVRP$ G)
1 $\mathcal{R} \leftarrow MinimumRestorationSetProblem(G, \mathcal{IN})$
2 $\mathcal{O} \leftarrow RestorationOrderProblem(\mathcal{IN}, \mathcal{R})$
3 **return** $PrecedenceRoutingProblem(G, \mathcal{O})$

Figure 1: The Multi-Stage IRVRP Algorithm.

to an overall intractable formulation. Indeed, even for a single infrastructure, where the demand maximization is a LP, tackling the problem globally is beyond the scope of existing MIP solvers. For this reason, we follow the multi-stage approach proposed in (Van Hentenryck, Coffrin, and Bent 2011), which was shown to produce high-quality solutions to the joint repair and restoration of the power system, even for large instances.

The multi-stage approach consists of three steps and is depicted in Figure 1. As inputs, the Infrastructure Restoration Vehicle Routing Problem (IRVRP) requires an infrastructure network $\mathcal{IN}$ and an IRVRP instance G, which contains the network damage information and other data necessary for constructing the vehicle routing problem. The first step is a minimum restoration set problem which determines the smallest set of items to restore the infrastructure to full capacity. The second set is a restoration order problem, which produces the order in which the components must be repaired. This order produces precedence constraints which are injected into the pickup and delivery routing problem to produce the restoration plan. Only the first two steps are affected when an interdependent power and gas infrastructure is considered and this paper only studies these two steps.

The Minimum Restoration Set Problem

The Minimum Restoration Set Problem (MRSP) determines a smallest set of items needed to restore the network to full capacity (Model 4). The optimization heavily builds on Model 3 but it has four significant changes. First, additional inputs are necessary, i.e., the set of damaged components $D \subseteq N$. Second, the objective (M4.1) now minimizes the number of repairs. Third, constraints (M4.2 and M4.3) ensure that the network will operate at full capacity. Fourth, constraint (M4.4) ensures that only undamaged items are activated. The remaining constraints are identical to (M3.2.2–M3.4.8) in Model 3.

The Restoration Ordering Problem

Once a set $R \subseteq D$ of items to repair is obtained, the Restoration Ordering Problem (ROP) determines the best order in which to repair the items. The ROP ignores the routing aspects and the duration to move from one location to another, which would couple the routing and demand maximization aspects. Instead, it views the restoration as a sequence of discrete steps and chooses which item to restore at each step. Model 5 depicts the ROP model for interdependent infrastructures. The ROP essentially duplicates Model 3 $|R|$ times, where R is the set of selected items to repair. These models are linked through the decision variables y_{ki} which specify whether item i is repaired at step k. Constraint (M5.2) ensures that undamaged items are activated, constraint (M5.3)

Model 4 The MRSP for Interdependent Infrastructures.

Inputs:
    M^e, M^g - the maximum demands in undamaged networks
    D - the set of damaged items
    All inputs from Model 3
Variables:
    Identical to Model 3
Minimize
$$\sum_{i \in D} y_i \tag{M4.1}$$
Subject to:
$$\sum_{d \in T^e} E_d^o \geq M^e \tag{M4.2}$$
$$\sum_{d \in T^g} G_d^o \geq M^g \tag{M4.3}$$
$$y_i = 1 \quad \forall i \in N \setminus (C \cup D) \tag{M4.4}$$
Constraints (M3.2.2–M3.4.8) from Model 3

Model 5 The ROP Model for Interdependent Infrastructures.

Inputs:
    R - the set of items to restore
    D - the set of damaged items
    All inputs from Model 3
Variables:
    Variables of Model 3 replicated $|R|$ times
Maximize
$$\sum_{k=1}^{|R|} W^e \sum_{d \in T^e} E_{kd}^o + W^g \sum_{d \in T^g} G_{kd}^o \tag{M5.1}$$
Subject to: $(1 \leq k \leq |R|)$
$$y_{ki} = 1 \quad \forall i \in N \setminus (C \cup D) \tag{M5.2}$$
$$\sum_{i \in R} y_{ki} = k \tag{M5.3}$$
$$y_{(k-1)i} \leq y_{ki} \quad \forall i \in R \tag{M5.4}$$
$|R|$ replicates of constraints (M3.2.2-M3.4.8) from Model 3

---

makes sure that at most one item is repaired at each step, and constraint (M5.4) ensures that an item remains repaired in subsequent steps. The objective (M5.1) maximizes the satisfied demands at each step. The remaining model constraints are identical to (M3.2.2–M3.4.8) in Model 3 but are replicated for each of the k models.

The ROP model is significantly more challenging for interdependent infrastructures because the demand maximization problem is now a MIP instead of a LP, which is the case for a single infrastructure. MIP solvers have significant scalability issues, mainly because the ROP generalizes the transmission switching problem which is known to be extremely challenging for state-of-the-art MIP solvers (e.g., (Fisher, O'Neill, and Ferris 2008)).

Randomized Adaptive Decompositions

To overcome these computational difficulties, we use a Randomized Adaptive Decomposition (RAD) scheme. RAD schemes have been found useful in a variety of applications in logistics (Bent and Van Hentenryck 2007), scheduling (Pacino and Hentenryck 2011), and disaster management (Simon, Coffrin, and Hentenryck 2012).

```
ROP-RAD(R, D, [s..S])
1   O ← ROP-UTIL(R, D);
2   while ¬stoppingCriteria()
3   do ⟨S_1, ..., S_l⟩ ← RandomPartition(O, [s..S]);
4       O ← PROP(⟨S_1, ..., S_l⟩, D);
5   return O;
```

Figure 2: The RAD algorithm for the ROP.

Informal Presentation First observe that the ROP can be viewed as a function $ROP : R \times D \rightarrow O$ that, given a set R of components to repair and a set of damage components D $(R \subseteq D)$, produces an ordering O of R maximizing the satisfied demands over time. The RAD scheme repeats the following two steps:

1. Partition the sequence O into the subsequences $S_1, ..., S_l$, i.e., $O = S_1 :: S_2 :: ... :: S_l$ where :: denotes sequence concatenation.

2. Solve an ROP problem, called the Priority Restoration Order Problem (PROP), in which the items in S_j must be scheduled before the items in S_{j+1} $(1 \leq j < l)$.

Obviously, the PROP produces a lower bound to the ROP. However, it enjoys a nice computational property: It can be solved by solving a sequence of smaller decoupled ROPs defined as

$$ROP(S_1, D)$$
$$...$$
$$ROP(S_i, D \setminus (S_1 \cup ... \cup S_{i-1}))$$
$$...$$
$$ROP(S_l, D \setminus (S_1 \cup ... \cup S_{l-1})).$$

The RAD scheme then starts from a solution O_0 obtained by a standard utilization heuristic. At iteration i, the scheme has a solution O_i which is partitioned to obtain a PROP P_i which is solved by exploiting the decoupling to obtain a solution O_{i+1}. The successive solutions satisfy

$$O_0 \leq O_1 \leq ... \leq O_i \leq ...$$

The RAD scheme also ensures that the random partition of a solution σ into $S_1 :: S_2 :: ... :: S_l$ produces subsequences of length between two parameters s and S in order to generate ROPs that are non-trivial and computationally tractable.

Formalization The RAD algorithm for the ROP is depicted in Figure 2. Observe that the partition uses the current solution O and that the PROP never degrades the solution quality since O is a solution to the PROP. The algorithm could be easily generalized to a variable neighborhood search (Hansen and Mladenovic 1998) by increasing the sequence size, e.g.,

$$S = (1 + \alpha)S$$

when no improvement to the solution is found after several iterations. This was not necessary to obtain high-quality solutions on our benchmarks however. There are many possibilities for the *stoppingCriteria()* function. We found that a combination of a fixed time limit and a limit on the number of iterations without improvement saved time on easier problems. The PROP is formally defined as follows.

Definition 1 (PROP). *Given* $S_1 \cup ... \cup S_l \subseteq D$, *the Priority Restoration Order Problem* $PROP(\langle S_1, ..., S_l \rangle, D)$ *is a ROP problem* $ROP(S_1 \cup ... \cup S_l, D)$ *with the following additional constraints* $(1 \leq j \leq l)$:

$$\forall i \in S_j : y_{ti} = 1 \text{ where } t = \sum_{n=1}^{j} |S_n| \qquad (1)$$

Observe that a consequence of these constraints is that all items in $S_1, ..., S_j$ are repaired before the items in S_{j+1} $(1 \leq j < l)$. We now show that the PROP can be decomposed into a set of independent ROPs.

Theorem 1. *A Priority Restoration Ordering Problem* $P = PROP(\langle S_1, ..., S_l \rangle, D)$ *can be solved optimally by solving* l *independent ROPs:*

$$\mathcal{R}_1 = ROP(S_1, D)$$
$$...$$
$$\mathcal{R}_i = ROP(S_i, D \setminus (S_1 \cup ... \cup S_{i-1}))$$
$$...$$
$$\mathcal{R}_l = ROP(S_l, D \setminus (S_1 \cup ... \cup S_{l-1})).$$

Proof. It is sufficient to show that the union of the objectives and constraints of the l independent ROPs is equivalent to the original PROP P. The objective equivalence follows from the fact that the sum of the objective functions of $\mathcal{R}_1, ..., \mathcal{R}_l$ is the objective function of P. The system of constraints is more interesting. The additional constraints of the PROP produce four properties. Consider a subsequence S_j and let $s_j = 1 + \sum_{n=1}^{j-1} |S_n|$ and $t_j = \sum_{n=1}^{j} |S_n|$: The following properties hold:

$$y_{si} = 1 \qquad \forall i \in (S_1 ... \cup S_{j-1})$$
$$y_{si} = 0 \qquad \forall i \in (S_j ... \cup S_l)$$
$$y_{ti} = 1 \qquad \forall i \in (S_1 ... \cup S_j)$$
$$y_{ti} = 0 \qquad \forall i \in (S_{j+1} ... \cup S_l)$$

These follow from the PROP Constraints (1) and constraints (M5.4) and are enforced in the l independent ROPs through the selection of the restoration and damage sets, i.e.,

$$\mathcal{R}_j = ROP(S_j, D \setminus (S_1 \cup ... \cup S_{j-1})).$$

Substituting in the ROP model, Constraints (M5.2) yields

$$y_{ki} = 1 \ \forall i \in N \setminus (C \cup D \setminus (S_1 \cup ... \cup S_{j-1}))$$

which ensures all the y variables satisfy the first PROP property at time s. By definition, the ROP will only restore the items in the restoration set. Assigning the restoration set to S_j ensures the remaining PROP properties hold.

Constraints (M5.4) in the PROP ensure that, once an item is repaired, it remains repaired. The key observation for this constraint is to look at the y_{ki} variables in s-t intervals, i.e., $[s_1..t_1][s_2..t_2] ... [s_l..t_l]$. For one of these intervals $[s_i..t_i]$ the ROPs enforce the precedence constraints

$$y_{(k-1)e} \leq y_{ke} \ \forall e \in S_i \ k \in [s_i + 1..t_i]$$

We now show that the remaining inequalities in the PROP can be removed when the four PROP properties are enforced. First, we know that all elements in S_i are repaired after time t_i, i.e.,

$$y_{ke} = 1 \ \forall e \in S_i , k \geq t_i$$

Hence, all subsequent inequalities for S_i are guaranteed to be satisfied and can be removed. We also know that all elements in S_i were not repaired before time s_i, i.e.,

$$y_{ke} = 0 \ \forall e \in S_i \ , k < s_i$$

Hence, all the previous inequalities for S_i are guaranteed to be satisfied and can be removed. Applying these simplifications for all of the intervals $1 \leq j \leq l$ reveals that the ROPs enforce all of the relevant constraints.

Constraints (M5.3) in the PROP ensures that at most one item is repaired at every time step. Using arithmetic transformations the original constraint

$$\sum_{i \in R} y_{ki} = k \ \forall k \in [1..|R|]$$

can be rewritten in terms of the subsequences

$$\sum_{j=1}^{l} \sum_{i \in S_j} y_{ki} = k \ \forall k \in [1..|R|].$$

By the PROP properties, for any subsequence S_j, we know all of the elements in $S_1, \ldots, S_{j-1}$ have been set to 1 and all of the elements in $S_{j+1}, \ldots, S_l$ have been set to 0. That is,

$$\sum_{m=1}^{j-1} \sum_{i \in S_m} y_{ki} = \sum_{m=1}^{j-1} |S_m|$$

$$\sum_{m=j+1}^{l} \sum_{i \in S_m} y_{ki} = 0.$$

The range $1..|R|$ can be partitioned into l s-t intervals, $[s_1..t_1], [s_2..t_2], \ldots, [s_l..t_l]$ and combined with the previous formulas, then the PROP Constraints (M5.3) for subsequence S_j become

$$\sum_{m=1}^{j-1} |S_m| + \sum_{i \in S_j} y_{ki} + \sum_{m=j+1}^{l} 0 = k \ \forall k \in [s_j..t_j].$$

After expanding the definition of $[s_j..t_j]$ the constant term $\sum_{m=1}^{j-1} |S_m|$ can be removed by changing the interval range to obtain

$$\sum_{i \in S_j} y_{ki} = k \ \forall k \in [1..|S_j|].$$

Hence, constraints (M5.3) becomes j disjoint constraints in the PROP model which are enforced in the ROPs.

The optimal solution of $\mathcal{P}$ can thus be obtained by concatenating the optimal solutions of the subproblems $\mathcal{R}_1, \ldots, \mathcal{R}_l$. □

The Utilization Heuristic The utilization heuristic used by the RAD algorithm (Figure 2) is designed to approximate current best-practices for prioritizing repairs. In existing best-practices, each infrastructure provider works independently and prioritizes their repairs based on the percentage of network flow that each element uses under normal operating conditions. This measure is called the utilization of

the element. Because each utility works independently, each infrastructure system is solved independently using Models 1 and 2. Given a flow on the power network f_e or gas network f_g, the utilization of these components is f_e/M^e and f_g/M^g respectively. Each infrastructure provider prioritizes repairs based on the greatest utilization values. Given that the utilization value is unit-less, these restoration priorities may be extended to the multi-infrastructure domain by using the weighting factors W^e and W^g. This greedy heuristic serves both as a seed for our hybrid optimization approach and as the basis for comparison. The experimental section will demonstrate that optimization brings significant improvements over this current best-practice.

Computational Considerations The RAD approach should be contrasted with a local search approach that would swap items in the current ordering. Such a local search is computationaly expensive, since a swap between items in positions i and j requires the solving of $(j - i + 1)$ Model 3 instances, which are MIP models. Moreover, the complexity of these MIP models makes it hard to determine which moves are attractive in the local search and thus forces the local search to examine a large number of costly moves. In contrast, the RAD scheme exploits temporal locality, the subsequences are small, and the MIP solver uses the linear relaxation to guide the large neighborhood exploration.

Practical Considerations In practice, even some ROP problems with fewer than 10 items can be challenging to solve optimally and may take several minutes. For this reason, our RAD scheme uses a time limit on the subproblems and does not always solve the ROPs optimally. It is also useful to point out that, in practice, all the decoupled ROPs can be solved in parallel. This feature was not used in our implementation but would be highly beneficial in practice.

Experimental Results

The benchmarks were produced by Los Alamos National Laboratory and are based on the power and gas infrastructures of the US. The disaster scenarios were generated using state-of-the-art hurricane simulation tools used by the National Hurricane Center (FEMA 2010; Reed 2008). The power network has 326 components and the gas network has 93 components. Network damages range from 10 to 120 components. The experiments were run on Intel Xeon 2.8GHz processors on Debian Linux. The algorithms were implemented in the COMET system using SCIP as a MIP solver. The execution times were limited to 1 hour to be compatible with the disaster recovery context. The weighting parameters W^e, W^g were selected to balance the demands of the networks in percentage ($W^e = 0.5/M^e, W^g = 0.5/M^g$), these results are consistent for other weightings. The subsequences in the decomposition are of sizes between 4 and 8. Our approach is compared to the *utilization heuristic* which approximates the current best-practices in multiple infrastructure restoration. The experiments only focus on the ROP problem which is the bottleneck of the approach.

Table 1: Multiplicative Effects of ROP Algorithms

BM	$|D|$	ROP MIP	RAD	$|R|$	MRSP+ROP MIP	RAD
1	10	1.42*	1.42	3	1.42*	1.42
2	12	4.594*	4.594	6	4.594*	4.594
3	14	2.277*	2.277	7	2.277*	2.277
4	14	2.967*	2.967	6	2.967*	2.967
5	14	2.928*	2.928	6	2.852*	2.852
6	14	2.856*	2.856	6	2.856*	2.856
7	15	2.056*	2.056	8	2.039*	2.039
8	17	1.536*	1.536	12	1.536*	1.536
9	17	1.567	1.57	11	1.57*	1.57
10	19	–	1.739	9	–	1.739
11	20	–	1.058	13	–	1.058
12	23	1.818	1.882	10	1.849*	1.849
13	23	5.191	6.724	6	6.721*	6.721
14	24	7.392	7.392	4	7.392*	7.392
15	26	–	1.345	13	1.333*	1.333
16	26	1.66	1.956	10	1.953*	1.953
17	27	–	2.507	13	–	2.486
18	29	1.355	2.342	12	2.379*	2.379
19	31	–	1.595	18	1.595	1.595
20	31	2.155	3.52	8	3.585*	3.585
21	33	–	2.333	17	–	2.136
22	36	1.678	1.796	13	1.711	1.799
23	43	–	2.283	24	–	2.356
24	46	–	2.49	28	2.163	2.558
25	49	–	2.029	22	2.076*	2.076
26	56	–	2.596	28	2.495	2.593
27	57	–	2.622	29	–	2.49
28	61	–	2.38	25	1.788	2.45
29	73	–	2.628	40	–	2.57
30	79	–	2.234	39	–	3.189
31	92	–	1.761	55	–	2.015
32	92	–	1.968	52	–	2.255
33	120	–	1.536	73	–	1.722
MIP-μ		2.716	2.988		2.981	2.987
MSRP MIP-μ			2.721		2.689	2.719
RAD-μ			2.513			2.558
Large RAD-μ			2.23			2.389

Table 2: ROP Algorithm Runtimes (seconds)

BM	$|D|$	ROP MIP	RAD	$|R|$	MRSP+ROP MIP	RAD
1	10	178.8	147.6	3	5.192	26.47
2	12	503.5	538.2	6	46.11	454.6
3	14	214.7	160.1	7	33.65	59.94
4	14	153.8	134.1	6	13.32	38.17
5	14	402.2	127.5	6	41.43	45.71
6	14	570.7	468.4	6	37.03	233.3
7	15	1101	850.4	8	135.7	243.5
8	17	496.9	250.4	12	129	142.9
9	17	3614	1224	11	502.1	428.3
10	19	–	880.7	9	–	264.9
11	20	–	885.9	13	–	577.6
12	23	3619	1142	10	408.5	354.1
13	23	3619	1276	6	22.54	141.5
14	24	3615	435.3	4	7.029	57.8
15	26	–	864	13	763.5	389
16	26	3618	740.2	10	180.9	469.9
17	27	–	2007	13	–	529.1
18	29	3620	1739	12	1102	506.7
19	31	–	1782	18	3615	713.7
20	31	3626	1047	8	52.22	84.65
21	33	–	1249	17	–	466.1
22	36	3632	901.1	13	3615	201
23	43	–	3646	24	–	2481
24	46	–	3708	28	3618	1983
25	49	–	3446	22	1425	484.4
26	56	–	3693	28	3618	1874
27	57	–	3647	29	–	1784
28	61	–	2611	25	3620	517.9
29	73	–	3691	40	–	3685
30	79	–	3722	39	–	2965
31	92	–	3730	55	–	3682
32	92	–	3679	52	–	3679
33	120	–	3758	73	–	3709
Proved	8			17		
ROP Proof Time μ	452.7			46.21		
Early Finish		24				29
Finish Time μ		1038				256.3

As mentioned earlier, the final routing is not affected by considering multiple interdependent infrastructures. Tables 1 and 2 present the quality and runtime data from the various ROP algorithms on 33 damage scenarios. The results are first grouped into ROP and MRSP+ROP, to show the benefits of including the MRSP stage. Column $|D|$ is the number of damaged items, *MIP* is the restoration result using Model 5, *RAD* is the restoration result using the decomposition from Figure 2, and $|R|$ is the restoration set size after using the MRSP. The values in the MIP and RAD columns indicate the multiplicative improvement over the utilization heuristic. For example, a value of 2.0 indicates that the optimization algorithm doubled the amount of satisfied demands over the time of the restoration (e.g., reduces the size of the power and gas "blackout" by 2). An asterisk indicates a proof of optimality. Entries are omitted for the MIP when no solution was found within the time constraints. The aggregate statistics at the bottom of the Table 1 summarize the results. Due to the incomplete MIP data, several subsets are of inter-

est: "MIP-μ" is the set of instances that the MIP can solve; "MRSP MIP-μ" is the set of instances that the MIP can solve when the MRSP is used; "RAD-μ" is the set of all instances; and "Large RAD-μ" is the set of instances where $|D| \geq 40$. To provide an intuition behind the numbers reported in Table 1, Figure 3 depicts the detailed restoration plans on Benchmark 20 for the utilization heuristic, the MIP approach, and the RAD algorithm. The figure shows the significant benefits provided by optimization technologies in general and the RAD approach in particular.

Overall, the results indicate that the RAD approach significantly improves the practice in the field, and more than doubles the level of service within the time constraints. The instances without the MRSP stage are particularly interesting, since they illustrate the scalability issues better. The statistics indicate that the RAD algorithm consistently outperforms the MIP approach, improving the solution quality from 2.716 to 2.988 on average. The MIP approach also has severe difficulties on the larger instances. The detailed data
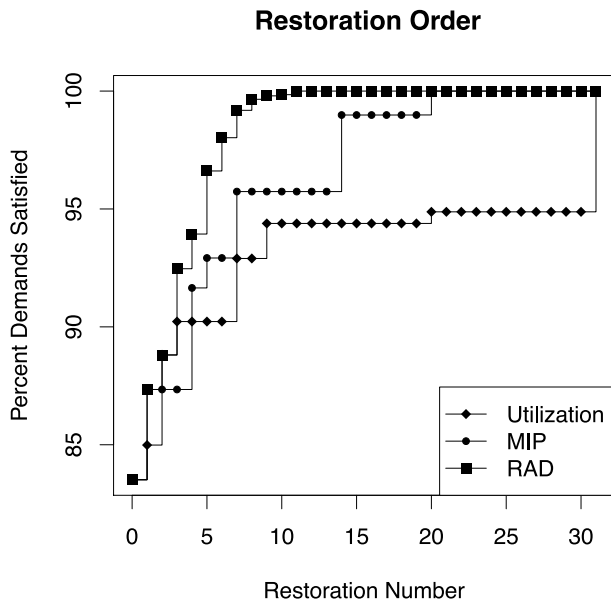
Restoration Order

Figure 3: Restoration Plans on Benchmark 20.

reveals that the RAD algorithm always matches the optimal solutions and, in cases where optimal solutions are not obtained, it often improves over the (suboptimal) MIP solution. The results also show that the RAD algorithm is significantly faster than the MIP approach.

The MRSP stage significantly reduces the damage set $|D|$ (close to a factor of 2). The results with the MRSP (i.e., the last three columns of Tables 1 and 2) indicate the MRSP brings significant improvements to the MIP model. The average quality on the smaller benchmarks is improved from 2.716 to 2.981 and 6 more benchmarks become feasible. The runtime benefits to the MIP model are also significant, as 9 more instances can be proven optimal and the proof runtimes are reduced by a factor of 10. The quality improvements of the MRSP to the RAD algorithm are negligible, except for the largest benchmarks. For the largest benchmarks, the MRSP increases solution quality from 2.23 to 2.389. The MRSP also produces runtime benefits, as the RAD algorithm terminates early on 29 benchmarks and the average early completion time is reduced by a factor of 5, to less than 5 minutes.

The decoupling of the ROP problem into the MRSP+ROP problems may remove the optimal ROP solution as benchmark 5 indicates. However such effects become insignificant as the damage size grows and the challenge of finding a high quality ROP solution increases. The decoupling is thus valuable both in terms of solution quality and runtimes.

Related Work

Disaster management and, in particular, the restoration of interdependent infrastructures are inherently interdisciplinary as they span the fields of reliability engineering, vulnerability management, artificial intelligence, and modeling of complex systems. The importance of interdependent infrastructure restoration was recognized soon after the 2001

World Trade Center Attack (Wallace et al. 2003) and this recognition continued to spread across several areas in the past decade (Cho 2007; Dueas-Osorio and Vemuru 2009; Buldyrev et al. 2010). Interdependence studies in the reliability engineering area (Dueas-Osorio and Vemuru 2009; Ouyang and Dueas-Osorio 2011) have primarily focused on topological properties such as *betweenness* and *connectivity loss*. In the area of artificial intelligence, to the best of our knowledge, restoration of interdependent infrastructures have not been studied. However, power system restoration has been considered (Bertoli et al. 2002; Hadzic and Andersen 2005; Bell et al. 2009). Although these methods are an excellent application of planning, configuration, and diagnosis techniques, they also use *connectivity* as the primary power model. These topological metrics provide some sufficient conditions for infrastructure operations. However, their fidelity is insufficient to incorporate line capacity constraints which are critical to model the pipeline compressor interdependencies studied here. Furthermore, the accuracy of topological metrics for models of infrastructure systems has recently been questioned (Hines, Cotilla-Sanchez, and Blumsack 2010) and the benefits of flow-based models of infrastructure systems is increasingly recognized by the reliability engineering community (Dueas-Osorio and Hernandez-Fajardo 2008).

References (Lee, Mitchell, and Wallace 2004; 2007; Gong et al. 2009; Cavdaroglu et al. 2011) are the closest related work and warrant a detailed review. Our approach fundamentally differs from these earlier studies: It uses the more accurate LDCM for power systems, it scales to large instances, and it models cyclic interdependencies. Reference (Lee, Mitchell, and Wallace 2004) provides a good background paper on the nature and classification of various interdependencies. Early work focused on solving the MRSP (Lee, Mitchell, and Wallace 2007) and considered the power, telephone, and subway infrastructure in New York City, but focused only on restoring the power infrastructure. (Gong et al. 2009) assumed that restoration tasks have predefined due dates and developed a logic-based benders decomposition for a weighted sum of different competing objectives. The impact on the actual infrastructure is not taken into account. (Cavdaroglu et al. 2011) tried to jointly solve the multi-machine model of (Gong et al. 2009) and the MRSP (Lee, Mitchell, and Wallace 2007). Although their model incorporated interdependencies, only damage and restoration of the power grid was studied. Computation times are between 3 and 18 hours, with optimality gaps of 0.4% and 3.0% respectively. They report that using the MRSP instead of the full damage decreases the quality of the solution by 4.5%. The worst-case effect in our formulation is similar. However, in damage scenarios for which the optimal solution is known, our MRSP/ROP decoupling rarely cuts off the optimal solution.

Conclusion

This paper considered the restoration of multiple interdependent infrastructures after a man-made or natural disaster. It presented the first scalable approach for the last-mile restoration of the joint electrical power and gas infrastruc-

tures, which features complex cyclic interdependencies. The underlying algorithms build on an earlier three-stage decomposition for restoring the power network that decouples the restoration ordering and the routing aspects. At the technical level, the key contributions of the paper are mixed-integer programming models for finding a minimal restoration set, restoration ordering, and a randomized adaptive decomposition scheme that obtains high-quality solutions within the required time limits. The approach is validated on a large selection of benchmarks based on the United States infrastructures and state-of-the-art weather and fragility simulation tools. The results show significant improvements over current field practices.

References

Bell, K.; Coles, A.; Coles, A.; Fox, M.; and Long, D. 2009. The role of ai planning as a decision support tool in power substation management. *AI Commun.* 22(1):37–57.

Bent, R., and Van Hentenryck, P. 2007. Randomized Adaptive Spatial Decoupling For Large-Scale Vehicle Routing with Time Windows. In *Proceedings of the 22th National Conference on Artificial Intelligence (AAAI'07)*. AAAI Press.

Bertoli, P.; Cimatti, R.; Slaney, J.; and Thibaux, S. 2002. Solving power supply restoration problems with planning via symbolic model checking. In *In ECAI*, 576–580. Wiley.

Buldyrev, S. V.; Parshani, R.; Paul, G.; Stanley, H. E.; and Havlin, S. 2010. Catastrophic cascade of failures in interdependent networks. *Nature* 464(7291):1025–1028.

Carvalho, R.; Buzna, L.; Bono, F.; Gutiérrez, E.; Just, W.; and Arrowsmith, D. 2009. Robustness of trans-european gas networks. *Phys. Rev. E* 80:016106.

Cavdaroglu, B.; Hammel, E.; Mitchell, J.; Sharkey, T.; and Wallace, W. 2011. Integrating restoration and scheduling decisions for disrupted interdependent infrastructure systems. *Annals of Operations Research (to appear)*.

Cho, N. 2007. Critical Infrastructure Rebuild Prioritization using Simulation Optimization. Master's thesis, Air Force Institute of Technology.

Dueas-Osorio, L., and Hernandez-Fajardo, I. 2008. Flow-based reliability assessment of infrastructure systems. In *14th World Conference on Earthquake Engineering (14WCEE)*, volume 31, 157 – 167.

Dueas-Osorio, L., and Vemuru, S. M. 2009. Cascading failures in complex infrastructure systems. *Structural Safety* 31:157 – 167. Risk Acceptance and Risk Communication.

FEMA. 2010. Fema hazus overview. Available online at http://www.fema.gov/plan/prevent/hazus/.

Fisher, E.; O'Neill, R.; and Ferris, M. 2008. Optimal transmission switching. *IEEE Transactions on Power Systems* 23(3):1346 –1355.

Gomez-Exposito, A.; Conejo, A. J.; and Canizares, C. 2008. *Electric Energy Systems: Analysis and Operation (Electric Power Engineering Series)*. CRC Press.

Gong, J.; Lee, E. E.; Mitchell, J. E.; and Wallace, W. A. 2009. Logic-based multiobjective optimization for restoration planning. In Chaovalitwongse, W.; Furman, K. C.; and Pardalos, P. M., eds., *Optimization and Logistics Challenges in the Enterprise*, volume 30 of *Springer Optimization and Its Applications*. Springer US. 305–324.

Hadzic, T., and Andersen, H. R. 2005. Interactive reconfiguration in power supply restoration. In *In 11th International Conference on Principles and Practice of Constraint Programming CP05*. Springer-Verlag GmbH.

Hansen, P., and Mladenovic, N. 1998. An Introduction to Variable Neighborhood Search. In Voss, S.; Martello, S.; Osman, I. H.; and Roucairol, C., eds., *Meta-Heuristics: Advances and Trends in Local Search Paradigms for Optimization*. Kluwer Academic Publishers, Norwell, MA. 433–458.

Hines, P.; Cotilla-Sanchez, E.; and Blumsack, S. 2010. Do topological models provide good information about vulnerability in electric power networks? *ArXiv e-prints*.

Knight, U. G. 1972. *Power systems engineering and mathematics, by U. G. Knight*. Pergamon Press Oxford, New York,.

Lee, E.E., I.; Mitchell, J.; and Wallace, W. 2004. Assessing vulnerability of proposed designs for interdependent infrastructure systems. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*, 8 pp.

Lee, E.; Mitchell, J.; and Wallace, W. 2007. 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.

Monforti, F., and Szikszai, A. 2010. A montecarlo approach for assessing the adequacy of the european gas transmission system under supply crisis conditions. *Energy Policy* 38(5):2486 – 2498.

Murillo-Sánchez, C., and Gan, D. 1997. Matpower. *Ithaca*.

Ouyang, M., and Dueas-Osorio, L. 2011. An approach to design interface topologies across interdependent urban infrastructure systems. *Reliability Engineering and System Safety* 96(11):1462 – 1473.

Pacino, D., and Hentenryck, P. V. 2011. Large neighborhood search and adaptive randomized decompositions for flexible job-shop scheduling. In Walsh, T., ed., *IJCAI*, 1997–2002. IJCAI/AAAI.

Powell, L. 2004. *Power System Load Flow Analysis (Professional Engineering)*. McGraw-Hill Professional.

Reed, D. A. 2008. Electric utility distribution analysis for extreme winds. *Journal of Wind Engineering and Industrial Aerodynamics* 96(1):123–140.

Simon, B.; Coffrin, C.; and Hentenryck, P. V. 2012. Randomized adaptive vehicle decomposition for large-scale power restoration. In Beldiceanu, N.; Jussien, N.; and Pinson, E., eds., *9th International Conference on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems (CPAIOR'12)*, volume 7298 of *Lectures Notes in Computer Science*, 379–394. Nantes, France: Springer Verlag.

Van Hentenryck, P.; Coffrin, C.; and Bent, R. 2011. Vehicle routing for the last mile of power system restoration. *Proceedings of the 17th Power Systems Computation Conference (PSCC'11), Stockholm, Sweden*.

Wallace, W. A.; Mendona, D.; Lee, E.; Mitchell, J.; and Chow, J. 2003. Managing disruptions to critical interdependent infrastructures in the context of the 2001 world trade center attack. In *Center, University of Colorado*, 165–198.

Wood, A. J., and Wollenberg, B. F. 1996. *Power Generation, Operation, and Control*. Wiley-Interscience.