Resilient Upgrade of Electrical Distribution Grids

Emre Yamangil
Rutgers University
New Brunswick, NJ, 08901
yamangil@rutgers.edu

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

Scott Backhaus
Los Alamos National Laboratory
Los Alamos, NM, 87545
backhaus@lanl.gov

Abstract

Modern society is critically dependent on the services provided by engineered infrastructure networks. When natural disasters (e.g. Hurricane Sandy) occur, the ability of these networks to provide service is often degraded because of physical damage to network components. One of the most critical of these networks is the electrical distribution grid, with medium voltage circuits often suffering the most severe damage. However, well-placed upgrades to these distribution grids can greatly improve post-event network performance. We formulate an optimal electrical distribution grid design problem as a two-stage, stochastic mixed-integer program with damage scenarios from natural disasters modeled as a set of stochastic events. We develop and investigate the tractability of an exact and several heuristic algorithms based on decompositions that are hybrids of techniques developed by the AI and operations research communities. We provide computational evidence that these algorithms have significant benefits when compared with commercial, mixed-integer programming software.

Introduction

Natural disasters such as earthquakes, hurricanes, and other extreme weather pose serious risks to modern critical infrastructure including electrical distribution grids. At the peak of Hurricane Sandy, 65% of New Jersey's customers lost power (Mansfield and Linzey 2013). Recent U.S. government sources (Executive Office of the President 2013; US Department of Energy 2013) suggest that new methodologies for improving system resilience to these events is necessary. Here, we focus on developing methods for designing and upgrading distribution grids to better withstand and recover from these threats that are inspired by techniques developed in the artificial intelligence and operations research communities. Our approach minimizes the upgrade budget while meeting a minimum standard of service by selecting from a set of potential upgrades, e.g. adding redundant lines, adding distributed (microgrid) generation (i.e. wind, solar, and combined heat and power), hardening existing components, etc.

We formulate our approach, i.e. Optimal Resilient Distribution Grid Design (ORDGDP), as a two-stage mixed-integer program. The first (investment) stage selects from the set of potential upgrades to the network. The second (operations) stage evaluates the network performance benefit of the upgrades against a set of damage scenarios sampled from a stochastic distribution. We first develop an exact solution method that exploits decomposition across the sampled scenarios. We also develop a metaheuristic that we call Scenario Based Variable Neighborhood Decomposition Search (SB-VNDS) that is a hybrid of Variable Neighborhood Search (Lazic et al. 2010) and the exact method. We present numerical evidence that our exact method is more efficient than out-of-the-box commercial mixed-integer programming solvers, and that our heuristic achieves near-optimal results in a fraction of the time required by exact methods.

Literature Review Network design problems and their variations are generally NP-complete (Tomaszewski, Pióro, and Żotkiewicz 2010; Nace et al. 2013; Johnson, Lenstra, and Kan 1978). However, recent work by (Bent, Berscheid, and Toole 2010) demonstrates that AI-based methods can lead to substantial improvement for realistic applications.

While the specific problem of designing resilient distribution systems is novel, a number of related problems exist. The flow of electric power in tree-like distribution networks is related to multi-commodity network flows making our problem similar to the design of multi-commodity flow networks with stochastic link and edge failures (Santoso et al. 2003; Garg and Smith 2008). However, the second stage of our formulation requires binary variables making our problem considerably more difficult than typical second-stage flow problems. The interdiction literature includes related max-min or min-max problems where the goal is to operate or design a system to make it as resilient as possible to an adversary who can damage up to k elements. Such models are similar to ours if a k is chosen that bounds the worst-case disaster (Chen and Phillips 2013; Chen et al. 2014; Salmeron, Wood, and Baldick 2009; Delgadillo, Arroyo, and Alguacil 2010). Binary variables at all stages make these models computationally challenging and solvable only for small k. Here, we exploit the probabilistic nature of our adversary to increase the size of tractable problems (eliminates a stage of binary variables).

In power engineering, papers have primarily focused on resilient system operation (Golari, Fan, and Wang 2014; Li et al. 2014; Khushalani, Solanki, and Schulz 2007) using controls such as line switching. The ORDGDP is a fundamental generalization of the resilient operations problem because 1) this problem is embedded in our second stage and 2) minimizing the number of switch actions (Li et al. 2014) can be thought of as a design problem for a single scenario. Finally, there is also a general power grid expansion planning problem for stochastic events (Jabr 2013) that is a variant of the single commodity flow problem, with the twist that flows are not directly controllable. Like stochastic multi-commodity flow, the second-stage variables are not binary.

The key contributions of the paper include:

- Computationally efficient algorithms for solving stochastic network design problems with discrete variables at each stage. The algorithms are based on hybrid optimization methods similar to recent work that combines Bender's Decomposition with heuristic master solutions (Raidl, Baumhauer, and Hu 2014).
- Introduction of a problem of critical importance to energy problems where AI researchers can make significant contributions. AI has made many recent significant contributions to energy problems (Hentenryck, Gillani, and Coffrin 2012; Reddy and Veloso 2012; Coffrin, Hentenryck, and Bent 2012; Garg, Jayram, and Narayanaswamy 2013; Reddy and Veloso 2013; Jain, Narayanaswamy, and Narahari 2014; Reddy and Veloso 2011; Shann and Seuken 2013; Thibaux et al. 2013).

Problem Description

Nomenclature

Parameters

$\mathcal{N}$	set of nodes (buses).
$\mathcal{E}$	set of edges (lines and transformers).
$\mathcal{S}$	set of disaster scenarios.
$\mathcal{D}_s$	set of edges that are inoperable during $s \in S$.
$\mathcal{D}'_s$	set of edges that are inoperable even though they are hardened during disaster $s \in S$.
c_{ij}	cost to build a line between bus i and j. 0 if line already exists.
κ_{ij}	cost to build a switch on a line between bus i and j.
ψ_{ij}	cost to harden a line between bus i and j.
$\zeta_{i,k}$	cost of generation capacity on phase k at bus i.
α_i	cost to build a generation facility at node i.
Q_{ijk}	line capacity between bus i and bus j on phase k.
$\mathcal{P}_{ij}$	set of phases for the line between bus i and bus j.
$\mathcal{P}_i$	set of phases allowed to consume or inject at bus i.
β_{ij}	parameter for controlling how much variation in flow between the phases is allowed.
$d_{i,k}$	demand for power at bus i for phase k.
$G_{i,k}$	existing generation capacity on phase k at node i.
$Z_{i,k}$	maximum amount of generation capacity on phase k that can be built at node i.
$\mathcal{C}$	the set of sets of nodes that includes a cycle.
λ	fraction of critical load that must be served.
γ	fraction of all load that must be served.
$\mathcal{L}$	set of buses whose load is critical.

Variables

x_{ij}	determines if line i, j is built.
τ_{ij}	determines if line i, j has a switch.
t_{ij}	determines if line i, j is hardened.
$z_{i,k}$	determines the capacity for generation on phase k at node i.
u_i	determines the generation capacity built at node i.
x^s_{ij}	determines if line i, j is used during disaster s.
τ^s_{ij}	determines if switch i, j is used during disaster s.
t^s_{ij}	determines if line i, j is hardened during disaster s.
$z^s_{i,k}$	determines the capacity for generation on phase k at bus i during disaster s.
u^s_i	indicates if the generation capacity is used at node i during disaster s.
$g^s_{i,k}$	generation produced for bus i on phase k during disaster s.
$l^s_{i,k}$	load delivered at bus i on phase k during disaster s.
$y^s_{i_j}$	determines if the jth load at bus i is served or not during disaster s.
$f^s_{ij,k}$	flow between bus i and bus j on phase k during disaster s.
$\bar{x}^s_{ij}$	determines if at least one edge between i and j is used during disaster s.
$\bar{\tau}^s_{ij}$	determines if at least one switch between i and j is used during disaster s.
$x^s_{ij,0}$	determines if there exists flow on line i, j from j to i, during disaster s.
$x^s_{ij,1}$	determines if there exists flow on line i, j from i to j, during disaster s.

Distribution Grid Modeling A distribution network is modeled as a graph with nodes $\mathcal{N}$ (buses) and edges $\mathcal{E}$ (power lines and transformers). In the physical system, each edge is composed of one, two, or three circuits or "phases" and the electrical loads at the nodes are connected to and consume power from specific phases (Garcia et al. 2000) ($\mathcal{P}$). In many papers, multiple phases are approximated as a single phase with a single edge flow. However, under the damaged and stressed conditions considered in this work, the flows on the phases are often unbalanced, i.e. unequal, making it important to model all phases to accurately evaluate flow constraints on each phase. The phase flows are *not* directly controllable, but are related to nodal voltages and power injections by non-convex, physics-based equations (Garcia et al. 2000). Incorporation of these equations into the current formulation increases the complexity, however, the structure of distribution networks enables a simplification.

The design of protection systems for the vast majority of distribution circuits is based on the these circuits having a tree-like structure. Therefore, although distribution grids are often designed to contain many possible loops, switches are used to ensure that these grids are operated in a tree or forest topology. While, the switches introduce binary variables that increase the complexity of the ORDGDP, a linearized version of the electrical power flow equations (i.e. DC power flow) on the resulting trees is equivalent to a commodity flow model. We use a multi-commodity flow model that models each phase separately (Fig. 1).

The linearization of the power flow equations assumes uniform voltage magnitude at all nodes and ignores reactive power flows. In practice, we expect these are reasonable approximations because, prior to being upgraded, the distribution grid is already feasible with respect to voltage and reactive power flows. By adding lines or distributed power sources, we put loads closer to generation thereby reducing voltage variability and reactive power flow and the potential for violating unmodeled constraints. In principle, it is possible to construct solutions where this is not the case, but the solutions to ORDGDP found by our algorithms has not resulted in these situations. However, this is an important area of future work, and we are developing methods to eliminate solutions that violate voltage or reactive power flow limits.

Damage Modeling The ORDGDP is also defined by a set of scenarios, $\mathcal{S}$. These scenarios are provided by a user or are drawn from a probabilistic damage model (the case here). Each scenario is defined by the lines of the network that are damaged and are inoperable. Many networks consist of mix of line and pole types, i.e. overhead and underground lines. In our model, this is reflected by different damage probabilities. For the purposes of this paper, we assume a static (peak) demand profile for each scenario. However, we note that multiple load patterns can be included by creating scenarios representing the same damage set with different demand profiles.

Upgrade Options We focus on four user-definable design options in distribution networks: 1) Hardening existing lines to lower the probability of damage, 2) Build new lines to add redundancy, 3) Build switches, to add operating flexibility, and 4) building distributed generation (sources of power). While deregulation has split network operation from generation ownership in transmission systems, in distribution systems (the focus here), this split varies from locale to locale and is our motivation for including generation as a design option. For example, Central Hudson has recently added generators for resilience and reliability (Central Hudson Gas and Electric 2014).

Optimization model Given a disaster $s \in S$, $\mathcal{Q}(s)$ in Fig. 1 defines the set of feasible distribution networks. The constraints of $\mathcal{Q}(s)$ involve a number of well-known constraints in the combinatorial optimization literature, including knapsacks, multi commodity flows, and tree constraints. In this model, Eq. 1 is a capacity constraint on phase flows. When the line is not built the flow is forced to 0 by x^s. Eq. 2 forces all phases to flow in the same direction, an engineering constraint. Eq. 3 states that the flow on a line is 0 when the switch is open. Eq. 4 limits the fractional flow imbalance between the phases to a value smaller than β_{ij}. Imbalance between phases cannot be extreme otherwise equipment may be damaged. Here, we use β_{ij} = 0.15 for transformers, and 1.0 otherwise. Eq. 5 removes components in the damage set

$$\mathcal{Q}(s) = \{x^s, \tau^s, t^s, z^s, u^s :$$

$$-x^s_{ij,0} Q_{ijk} \leq f^s_{ij,k} \leq x^s_{ij,1} Q_{ijk} \qquad \forall ij \in \mathcal{E}, k \in \mathcal{P}_{ij} \tag{1}$$

$$x^s_{ij,0} + x^s_{ij,1} \leq x^s_{ij} \qquad \forall ij \in \mathcal{E} \tag{2}$$

$$(\tau^s_{ij} - 1) Q_{ijk} \leq f^s_{ij,k} \leq (1 - \tau^s_{ij}) Q_{ijk} \quad \forall ij \in \mathcal{E}, k \in \mathcal{P}_{ij} \tag{3}$$

$$\frac{\sum_{k \in \mathcal{P}_{ij}} f_{ij,k}}{\frac{|\mathcal{P}_{ij}|}{(1-\beta_{ij})}} \leq f^s_{ij,k'} \leq \frac{\sum_{k \in \mathcal{P}_{ij}} f_{ij,k}}{\frac{|\mathcal{P}_{ij}|}{(1+\beta_{ij})}} \qquad \forall ij \in \mathcal{E}, k' \in \mathcal{P}_{ij} \tag{4}$$

$$x^s_{ij} = t^s_{ij} \leq \begin{cases} 0 & \text{if } ij \in \mathcal{D}'_s \\ 1 & \text{else} \end{cases} \qquad \forall ij \in \mathcal{D}_s \tag{5}$$

$$l^s_{i,k} = \sum_{j=0}^{n_i} y^s_{i_j} d_{i_j,k} \qquad \forall i \in \mathcal{N}, k \in \mathcal{P}_i \tag{6}$$

$$0 \leq g^s_{i,k} \leq z^s_{i,k} + G_{i,k} \qquad \forall i \in \mathcal{N}, k \in \mathcal{P}_i \tag{7}$$

$$g^s_{i,k} - l^s_{i,k} - \sum_{j \in \mathcal{N}} f^s_{ij,k} = 0 \qquad \forall i \in \mathcal{N}, k \in \mathcal{P}_i \tag{8}$$

$$0 \leq z^s_{i,k} \leq Z_{i,k} u_i \qquad \forall i \in \mathcal{N}, k \in \mathcal{P}_i \tag{9}$$

$$\sum_{ij \in \mathcal{E}(C)} (x^s_{ij} - \tau^s_{ij}) \leq |V| - 1 \qquad \forall C \in \mathcal{C} \tag{10}$$

$$\tau^s_{ij} \leq x^s_{ij} \qquad \forall ij \in \mathcal{E} \tag{11}$$

$$\sum_{i \in \mathcal{L}, k \in \mathcal{P}_i} l^s_{i,k} \geq \lambda \sum_{i \in \mathcal{L}, k \in \mathcal{P}_i} d_{i,k} \tag{12}$$

$$\sum_{i \in \mathcal{N} \setminus \mathcal{L}, k \in \mathcal{P}_i} l^s_{i,k} \geq \gamma \sum_{i \in \mathcal{N} \setminus \mathcal{L}, k \in \mathcal{P}_i} d_{i,k} \tag{13}$$

$$x^s, y^s, \tau^s, u^s, t^s \in \{0, 1\}\} \tag{14}$$

Figure 1: Set of feasible distribution networks

from the network by linking the two damage sets with the hardening variables. Eq. 6 requires that all or none of the load at a bus is served. Once again, this an engineering limitation of most networks. Eq. 7 limits the distributed generation output by the generation capacity. Eq. 8 ensures flow balance at the nodes for all phases. Eq. 9 caps the generation capacity installed at the nodes. Eq. 10 eliminates network cycles, forcing a tree or forest topology. Eq. 11 states a switch is used only if the line exists. Eq. 12 ensures a minimum fraction λ of critical load is served. Here, we generally require λ = 0.98. Eq. 13 ensures that a minimum fraction of load is served. Here, $\gamma = 0.5$. Eqs. 12 and 13 are the resilience criteria that must be met by $\mathcal{Q}(s)$ and are similar to the $n - k - \epsilon$ criteria of (Chen et al. 2014). Eq. 14 states which variables are discrete.

One of the more difficult constraints in this formulation is Eq. 10 due to possible combinatorics. There are different ways to implement cycle constraints, and we use the formulation in Fig. 2.

$$\sum_{ij \in \mathcal{E}(C)} (\bar{x}^s_{ij} - \bar{\tau}^s_{ij}) \leq |V| - 1 \qquad \forall C \in \mathcal{C} \tag{15}$$

$$x^s_{ij} \leq \bar{x}^s_{ij} \qquad \forall ij \in \mathcal{E} \tag{16}$$

$$\tau^s_{ij} \geq x^s_{ij} + \bar{\tau}^s_{ij} - 1 \qquad \forall ij \in \mathcal{E} \tag{17}$$

Figure 2: Cycle constraints

While the multi-graph structure introduces a large number of cycles, there is a relatively small number of cycles when the multi-edges are reduced to one edge. Thus, we introduce binary variables (linear number) for the edges of the corresponding single-edge graph and enumerate the possible cycles in that graph (Eq. 15). Then, Eqs 16 and 17 are used to pass information between artificial cycle variables and the actual line and switch variables.

For each $s \in \mathcal{S}$, $\mathcal{Q}(s)$ determines the set of feasible distribution networks. There are some redundant variables in this formulation that improve the separability of the problem. The ORDGDP is the minimum cost design that falls in the intersection of all the $\mathcal{Q}(s)$ (Fig. 3).

$$\min \quad \sum_{ij \in \mathcal{E}} c_{ij} x_{ij} + \sum_{ij \in \mathcal{E}} \kappa_{ij} \tau_{ij} + \sum_{ij \in \mathcal{E}} \psi_{ij} t_{ij} + \sum_{i \in \mathcal{N}} \alpha_i u_i + \sum_{i \in \mathcal{N}, k \in \mathcal{P}_i} \zeta_{i,k} z_{i,k} \tag{18}$$

s.t.

$$x_{ij}^s \le x_{ij} \qquad \forall ij \in \mathcal{E}, s \in \mathcal{S} \tag{19}$$

$$\tau_{ij}^s \le \tau_{ij} \qquad \forall ij \in \mathcal{E}, s \in \mathcal{S} \tag{20}$$

$$t_{ij}^s \le t_{ij} \qquad \forall ij \in \mathcal{E}, s \in \mathcal{S} \tag{21}$$

$$z_{i,k}^s \le z_{i,k} \qquad \forall i \in \mathcal{N}, k \in \mathcal{P}_i, s \in \mathcal{S} \tag{22}$$

$$u_i^s \le u_i \qquad \forall i \in \mathcal{N}, s \in \mathcal{S} \tag{23}$$

$$z_{i,k} \le M_{i,k} u_i \qquad \forall i \in \mathcal{N}, k \in \mathcal{P}_i \tag{24}$$

$$(x^s, \tau^s, t^s, z^s, u^s) \in \mathcal{Q}(s) \qquad \forall s \in \mathcal{S} \tag{25}$$

$$x, \tau, t, u \in \{0, 1\} \tag{26}$$

Figure 3: Optimal Resilient Distribution Grid Design

Eq. 18 minimizes the cost of building lines and switches, hardening lines, and building facilities and generation. For notational simplicity, existing lines, switches, and generation are included as variables in the objective with 0 cost, however in practice these enter the formulation as constants. Eqs. 19 through 24 tie the first stage (construction) decisions with second stage variables ($\mathcal{Q}(s)$). Eq. 25 states that the mixed-integer vector $(x^s, \tau^s, t^s, z^s, u^s)$ constitutes a feasible distribution network for scenario s.

Generalizations Without loss of generality, the formulation in Fig. 3 assumes the x_{ij}^s variables are treated as constants if the lines exist and are not in $\mathcal{D}_s$. Furthermore, Fig. 3 also assumes that hardened lines and new lines are built with switches. This is reflective of current industry practices and arises from the observation that switch costs are negligible when compared with the cost of the line itself. However, this assumption can be eliminated by modifying constraints 19 and 21 as follows:

$$\begin{aligned} x_{ij}^s &= x_{ij} \quad \forall ij \notin \mathcal{D}_s \\ x_{ij}^s &\le x_{ij} \quad \forall ij \in \mathcal{D}_s \\ t_{ij}^s &= t_{ij} \quad \forall ij \notin \mathcal{D}'_s \\ t_{ij}^s &\le t_{ij} \quad \forall ij \in \mathcal{D}'_s \end{aligned} \tag{27}$$

Finally, for notational simplicity, the formulation of Fig. 3 also assumes $ij \notin \mathcal{D}_s, ij \in \mathcal{D}'_s$ never occurs. However, if necessary this assumption can be relaxed by introducing auxiliary variables and additional constraints.

Chance Constraints For some networks, a very small number of scenarios in $\mathcal{S}$ may drive the total cost in Eq. 18. In real-world applications, the designer of the network may lower the total investment cost by accepting some risk of not always satisfying the resiliency criteria. In these situations, we can relax Eqs. 12 and 13 to a set of chance constraints:

$$P\left(\begin{array}{ll} \sum_{i \in \mathcal{L}, k \in \mathcal{P}_i} l_{i,k}^s \ge \lambda \sum_{i \in \mathcal{L}, k \in \mathcal{P}_i} d_{i,k} & \forall s \in \mathcal{S} \\ \sum_{i \in \mathcal{N}, k \in \mathcal{P}_i} l_{i,k}^s \ge \gamma \sum_{i \in \mathcal{N}, k \in \mathcal{P}_i} d_{i,k} & \forall s \in \mathcal{S} \end{array}\right) \ge 1 - \epsilon \tag{28}$$

When assuming the scenarios follow a uniform distribution, this is equivalent to stating that these constraints are violated in $\epsilon|S|$ of the scenarios. Thus, we can restate these constraints as:

$$\begin{array}{ll} \sum_{i \in \mathcal{L}, k \in \mathcal{P}_i} l_{i,k}^s \ge \lambda \sum_{i \in \mathcal{L}, k \in \mathcal{P}_i} d_{i,k}(1 - v_s) & \forall s \in \mathcal{S} \\ \sum_{i \in \mathcal{N}, k \in \mathcal{P}_i} l_{i,k}^s \ge \gamma \sum_{i \in \mathcal{N}, k \in \mathcal{P}_i} d_{i,k}(1 - v_s) & \forall s \in \mathcal{S} \\ \sum_{s \in S} v_s \le \epsilon |S| & \end{array} \tag{29}$$

Algorithms

In this section we discuss the algorithms we developed for solving the ORDGDP. ORDGDP is a two-stage mixed integer programming (MIP) problem with a block diagonal structure that includes coupling variables between the blocks. We developed an exact algorithm that is vastly more efficient than a commercial state-of-the-art MIP solver. We then used the exact algorithm to develop a hybrid with variable neighborhood search that is competitive with the exact solver and is better than a heuristic used by the industry.

Scenario-Based Decomposition (SBD) Decomposition is often used for solving two-stage stochastic MIPs (Vanderbeck and Wolsey 2010), and it can be applied to ORDGDP after the following key observation:

Observation 0.1 *The second stage variables do not appear in the objective function. Therefore any optimal first stage solution based on a subset of the second stage subproblems that is feasible for the remaining scenarios, is an optimal solution for the original problem.*

Based on this observation, we can apply SBD to solve the ORDGDP. At high level, Algorithm 1 solves problems with iteratively larger sets of scenarios until a solution is obtained that is feasible for all scenarios. The algorithm takes as input the set of disasters (scenarios) and an initial scenario to consider, S'. Line 2 solves ORDGDP on S', where $P(S')$ and σ^* are used to denote the problem and solution respectively. Line 3 then evaluates σ^* on the remaining scenarios in $S \backslash S'$. The function $l : P'(s, \sigma^*) \rightarrow \mathbb{R}_+$, is an infeasibility measure that is 0 if the problem is feasible, positive otherwise. This is implemented by maximizing the reliability constraints, i.e. total and critical demand satisfied. It measures the gap between the delivered and the required demand (the right hand side of the Eqs. 12 and 13). This function prices the current solution over $s \in S \backslash S'$. If all prices are 0, then the algorithm terminates with solution σ^* (lines 4-5). Otherwise, the algorithm adds the scenario with the worst infeasabilty measure to S' (line 7).

We also tested other decomposition strategies such as Benders and Dantzig-Wolfe, however, their performance was tempered by the ORDGDP structure. The ORDGDP has MIP formulations at both stages of the problem and does not contain optimality conditions in the second stage (only feasability conditions). These approaches rarely out performed the commerical MIP solver.

Algorithm 1: Scenario Based Decomposition

input: A set of disasters S and let $S' = S_0$;
1 **while** $S \setminus S' \neq \emptyset$ **do**
2 $\quad \sigma^* \leftarrow$ Solve $P(S')$;
3 $\quad I \leftarrow \langle s_1, s_2 \ldots s_{|S \setminus S'|} \rangle\, s \in S \setminus S'$: $l(P'(s_i, \sigma^*)) \geq l(P'(s_{i+1}, \sigma^*))$;
4 $\quad$ **if** $l(P'(I(0), \sigma^*)) \leq 0$ **then**
5 $\quad\quad$ **return** σ^*;
6 $\quad$ **else**
7 $\quad\quad S' \leftarrow S' \cup I(0)$;
8 $\quad$ **return** σ^*

Greedy Algorithm A computationally efficient way of generating feasible solutions to the ORDGDP relaxes the coupling first stage variables and solves each scenario $s \in \mathcal{S}$ individually. The solutions are combined by taking the maximum of each construction variable ($\mathcal{X} = x \cup \tau \cup t \cup z \cup u$) over all scenarios (Algorithm 2). The switch construction cost is determined by switches that are needed to reduce the network into a tree for every scenario (line 4). Although the Greedy Algorithm is simple and fast, it rarely results in an optimal investment decision. However, it is representative of the types of heuristics used by the industry: see Reference (Munoz et al. 2014) for a survey.

Algorithm 2: Greedy

input: A set of disasters S;
1 **for** $s \in S$ **do**
2 $\quad \sigma^s \leftarrow Solve(P'(s))$;
3 $\sigma^*(x) = \max\{\sigma^s(x) | \forall s \in S\},\ \forall x \in \mathcal{X}$;
4 Update $\sigma^*(x_i)$ with switches to preserve feasibility;
5 **return** σ^*

Variable Neighborhood Search To overcome the limitations of greedy heuristics like Algorithm 2, we developed an approach based on Variable Neighborhood Decomposition (VNS) Search (Lazic et al. 2010). The algorithm fixes a subset of first stage variables to their current value and searching the remaining variables for a better solution. If all the first stage variables are fixed, the problem decomposes into $|S|$ separate problems that are easily solved and provide heuristic justification for focusing on first stage variables. More formally, $P(\sigma, J)$ denotes the problem with first stage variables, $J \in \mathcal{X}$,fixed to σ, i.e. $x_j = \sigma(x_j)$, and P^{LP} is the LP relaxation of problem P.

Algorithm 3 describes the VNS procedure. Line 1 computes the solution to the LP relaxation of the ORDGDP, (σ^{LP}). Line 4 counts the number of variable assignments that are different between the solution to LP relaxation (σ^{LP}) and the best known solution σ^* ($\sigma(x)$ denotes the variable assignment of x in solution σ). Line 5 orders the variables of $\mathcal{X}$ by the difference between their assignments in σ^* and σ^{LP}. Heuristically, those variables whose assignments are furthest from their LP assignment represent good opportunities to improve σ^*. The algorithm updates the rate at which the neighborhood size is increased ($step$) based on whether or not the algorithm is in a restart situation (lines 8 and 11). If the algorithm is in a restart, the ordering of the variables is also randomized (line 9). Line 13 computes the best solution in the neighborhood of σ where the first k elements of J are fixed. If the resulting solution is better, then the algorithm proceeds with a new σ^* (lines 15-18)–f is used as shorthand for Eq. 18. Otherwise, the size of the neighborhood is increased (lines 20-23). The iterations terminate when the maximum number of restarts is reached (line 2), the maximum number of neighborhood resizings is reached (line 12), or a time limit is reached. In this paper, MAXRESTARTS $= 10$, MAXITERATIONS $= 4$, MAXTIME $= 48$ CPU hours, and $d = 2$.

Algorithm 3: Variable Neighborhood Search

input: σ', MAXTIME, MAXRESTARTS and MAXITERATIONS;
1 Let $\sigma^{LP} \leftarrow Solve(P^{LP})$, $\sigma^* \leftarrow \sigma'$, $restart \leftarrow false$;
2 **while** $t <$ MAXTIME *and* $i <$ MAXRESTARTS **do**
3 $\quad j \leftarrow 0$;
4 $\quad n \leftarrow |x \in \mathcal{X} : |\sigma^*(x) - \sigma^{LP}(x)| \neq 0|$;
5 $\quad J \leftarrow \langle \pi_1, \pi_2 \ldots \pi_{|J|} \rangle \in \mathcal{X}$: $|\sigma^*(\pi_i) - \sigma^{LP}(\pi_i)| \leq |\sigma^*(\pi_{i+1}) - \sigma^{LP}(\pi_{i+1})|$;
6 $\quad$ **if** $restart$ **then**
7 $\quad\quad i \leftarrow i + 1$;
8 $\quad\quad step \leftarrow \frac{4n}{d}$, $k = |\mathcal{X}| - step$;
9 $\quad\quad$ *shuffle*(J)
10 $\quad$ **else**
11 $\quad\quad step \leftarrow \frac{n}{d}$, $k = |\mathcal{X}| - step$;
12 $\quad$ **while** $t <$ MAXTIME *and* $j \leq$ MAXITERATIONS **do**
13 $\quad\quad \sigma' \leftarrow Solve(P(\sigma^*, J(1, \ldots, k)))$;
14 $\quad\quad$ **if** $f(\sigma') < f(\sigma^*)$ **then**
15 $\quad\quad\quad \sigma^* \leftarrow \sigma'$;
16 $\quad\quad\quad i \leftarrow 0$;
17 $\quad\quad\quad restart \leftarrow false$;
18 $\quad\quad\quad j \leftarrow$ MAXITERATIONS;
19 $\quad\quad$ **else**
20 $\quad\quad\quad j \leftarrow j + 1$;
21 $\quad\quad\quad k = k - \frac{step}{2}$;
22 $\quad\quad\quad$ **if** $j >$ MAXITERATIONS **then**
23 $\quad\quad\quad\quad restart \leftarrow true$;
24 $\quad$ **return** σ^*

In our experimentation, VNS outperformed other popular random walk heuristics, such as Simulated Annealing (SA). We conjecture that this is because the ORDGDP does not appear to have a concise neighborhood structure, which is generally a prerequisite for successful SA implementations. Here, we overcome this challenge by using a mixed-integer program as a neighborhood oracle within the local search step of VNS.

Scenario-based Variable Neighborhood Decomposition Search (SBVNDS) Given that we have a powerful exact method in Algorithm 1 as well as a VNS in Algorithm 3, the natural algorithm hybridizes these approaches to get Algorithm, SBVNDS. The algorithm proceeds exactly the same as Algorithm 1, except that the exact solver for $Solve(P(S'))$ is replaced by VNS in line 2.

Empirical Results

The algorithms were implemented using the CPLEX C++ API with Concert technology as a 32 threaded application on Intel XEON 2.29 GHz processors. Since these are planning problems, in principle, practitioners could utilize days of CPU time to produce a plan. However, in order to produce a wide range of results, we limited the algorithms to 48 hours of CPU time. Our problems are based on a modified version of the IEEE 34 bus systems (Kersting 1991) (see Fig. 4) that are representative of medium sized distribution systems.

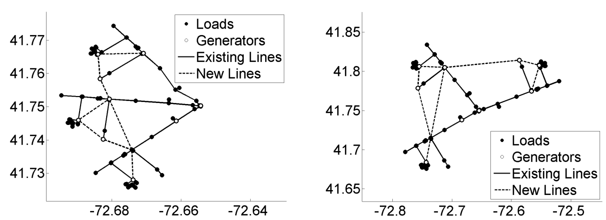

(a) Urban (b) Rural

Figure 4: We generated two variations of the IEEE 34 bus problem. Each problem contains three copies of the IEEE 34 system to mimic situations where there are three normally independent distribution circuits that could support each other during extreme events. These problems include 100 scenarios, 109 nodes, 118 possible generators, 204 loads, and 148 edges, resulting in problems with $> 90k$ binary variables. The difference between rural (a) and urban (b) is the distances between nodes (expansion costs and line impedances). The cost of single and three phase underground lines is between $\$40k$ and $\$1500k$ per mile (Governor and General Assembly of Virginia 2005) and we adopt the cost of $\$100k$ per mile and $\$500k$ per mile, respectively. The cost of single and three phase switches is estimated to be $\$10k$ and $\$15k$, respectively (Bialek 2013). Finally, the installed cost of natural gas-fired CHP in a microgrid is estimated to be $\$1500k$ per MW (EIA 2010). Full details of the problems are available at `http://public.lanl.gov/rbent/`.

Scenarios for this paper are based on damage caused by ice storms, whose intensity tends to be homogeneous on the scale of distribution systems (Sa 2002). Intensities are modeled as damage rates per mile on power poles and are transformed into the probability a power line segment of one mile length is damaged (a pole has failed). Empirically, we find that 100 randomly created scenarios is sufficient to capture the salient features of the distribution. Each scenario contains two sets of line failures, one for hardened lines ($\mathcal{D}'_s$) and a second for lines that are not hardened ($\mathcal{D}_s$).

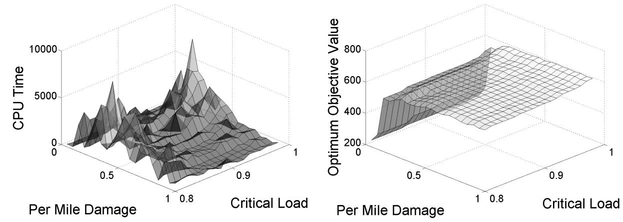

(a) CPU time (b) Objective value

Figure 5: Sensitivity of the CPU time and objective value to changes in λ on the Urban problem for SBD when hardened lines are not damageable. Due to short distances, the solution favors hardening many lines. The required hardening is relatively insensitive to the amount of damage and λ. However, there are spikes in problem difficulty at transitions in λ that require additional load service.

Table 1 provides results when hardened lines are not damaged or are damaged at rates of $\frac{1}{100}$ or $\frac{1}{10}$ of the unhardened rate. There are a number of important observations in these tables. First, CPLEX by itself is computationally uncompetitive. Only when the hardened lines are not damaged does CPLEX complete within the time limit. These problems are "easier" because hardened lines are robust and relatively inexpensive, enabling CPLEX to eliminate many solutions. The objective value for Greedy is always worse than optimal. The exact method SBD is much more computationally efficient than CPLEX and is able to solve many more problems to optimality indicating that CPLEX is unable to recognize the scenario structure in the problems. However, SBD is sensitive to which scenarios are included (function l), and if poor choices are made, it begins to resemble CPLEX. However, the meta-heuristic SBVNDS is able to overcome these limitations. It is much faster than SBD, and almost always achieves the optimal solution. This indicates that heuristic methods based on combining powerful techniques like VNS with strong exact algorithms are very good on this type of 2-stage mixed integer programming problems.

Critical load constraint Figures 5 and 6 show some results for rural and urban problems when the required fraction of critical load served is varied. In general, peaks in CPU time correspond to discrete jumps in the amount of load served as λ increases.

Chance constraints Fig. 7 shows results when the resiliency criteria are relaxed to the chance constraints in

Urban, Hardened lines are not damageable (a)

	CPLEX		Greedy	SBD		SBVNDS	
	CPU	OBJ	OBJ	CPU	OBJ	CPU	OBJ
10%	19984.7	322.9	1044.5	465.8	322.9	289.9	353.7
25%	166352	635.4	1643.5	8028.3	635.4	811.4	635.4
50%	TO	X	2021.2	2840.7	647.7	791.3	647.7
75%	TO	X	1874.2	991.1	652.1	692.5	652.1
100%	TO	X	1934.4	712.7	654.1	662.5	654.1

Rural, Hardened lines are not damageable (b)

	CPLEX		Greedy	SBD		SBVNDS	
	CPU	OBJ	OBJ	CPU	OBJ	CPU	OBJ
10%	33083.5	2337.0	3274.8	1837.9	2337.0	503.3	2337.0
25%	32170.8	2390.3	3427.6	571.0	2390.3	457.8	2390.3
50%	20840.3	2397.6	3449.9	471.2	2397.6	421.2	2397.6
75%	15556.1	2400.4	3452.7	337.5	2400.4	299.8	2400.4
100%	17225.9	2400.6	2780.6	385.8	2400.6	346.9	2400.6

Urban, Hardened lines are damaged at a $\frac{1}{100}$ rate (c)

	CPLEX		Greedy	SBD		SBVNDS	
	CPU	OBJ	OBJ	CPU	OBJ	CPU	OBJ
10%	159166	445.8	1061.7	2232.9	445.8	2721.3	476.5
25%	TO	X	1441.9	14299.2	662.9	2994.7	701.5
50%	TO	X	1571.2	2848.7	646.0	1917.7	760.2
75%	TO	X	1787.3	16040.6	687.6	1481.4	687.6
100%	TO	X	2744.8	24270.3	1320.5	2157.5	1330.5

Rural, Hardened lines are damaged at a $\frac{1}{100}$ rate (d)

	CPLEX		Greedy	SBD		SBVNDS	
	CPU	OBJ	OBJ	CPU	OBJ	CPU	OBJ
10%	77947.9	2363.0	3375.4	759.0	2363.0	576.9	2363.0
25%	TO	X	8238.6	TO	X	919.4	6744.3
50%	TO	X	12336.0	TO	9288.9	4361.8	7121.0
75%	TO	X	23099.5	TO	X	23142.6	11500.0
100%	TO	X	16600.7	TO	X	5879.5	9797.3

Urban, Hardened lines are damaged at a $\frac{1}{10}$ rate (e)

	CPLEX		Greedy	SBD		SBVNDS	
	CPU	OBJ	OBJ	CPU	OBJ	CPU	OBJ
10%	TO	X	859.1	5265.1	460.8	2505.7	594.1
25%	TO	X	1742.2	12530.3	961.2	2843.2	961.2
50%	TO	X	3133.8	34822.7	1417.2	3363.5	1555.2
75%	TO	X	3472.0	TO	X	7486.5	1894.2
100%	TO	X	10479.1	TO	X	32289.8	7959.4

Rural, Hardened lines are damaged at a $\frac{1}{10}$ rate (f)

	CPLEX		Greedy	SBD		SBVNDS	
	CPU	OBJ	OBJ	CPU	OBJ	CPU	OBJ
10%	TO	X	7503.3	141718.0	4325.9	7756.8	4424.8
25%	TO	X	18021.3	TO	X	21993.5	7371.9
50%	TO	X	28865.0	TO	12017.7	74729.0	12031.2
75%	TO	X	31887.0	TO	13522.2	107165.0	13500.8
100%	TO	X	32901.9	TO	16794.4	114354.0	16778.2

Table 1: These tables compare the performance of the algorithms when hardened lines cannot be damaged (a, b), are damaged at $\frac{1}{100}$ the rate of unhardened lines (c, d), and damaged at $\frac{1}{10}$ the rate of unhardened lines (e, f). The columns denoted by CPU and OBJ refer to CPU time and objective value, respectively. We omit the CPU time of Greedy as it is always less than 60 CPU seconds. The rows refer to the probability a 1 mile segment of a line is damaged.

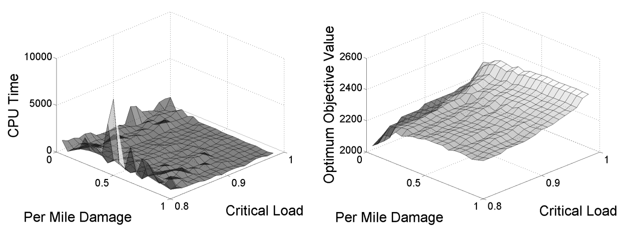

(a) CPU time (b) Objective value

Figure 6: Sensitivity of the CPU time and objective value to changes in λ for SBD on the Rural problem when hardened lines are not damageable. Because of long distances, the solution favors adding generation and is sensitive to the amount of damage and λ.

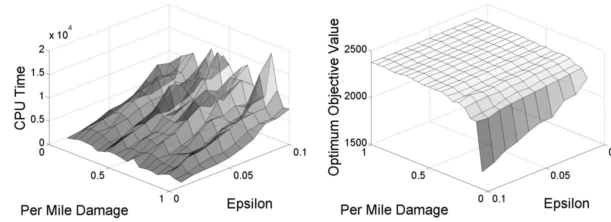

(a) CPU time (b) Objective value

Figure 7: These figures show how the CPU time and solution quality changes when chance constraints (ϵ) is modified for the Rural network, when hardened lines are not damageable. These plots are generated by SBD.

Eq. 29 and ϵ is varied. Interestingly, CPU time is not impacted too greatly by damage rates. Also, the solution is relatively insensitive to the choice of ϵ as damage rates increase, indicating that an "easier" problem with small ϵ could be used to approximate a solution to the harder problems.

Conclusions

We formulated, proposed and tested new algorithms to solve the ORDGDP. Our primary contribution is an algorithm that combines the benefits of an exact method based on scenario decomposition with variable neighborhood search. This algorithm is shown to scale well to problems that are difficult for exact methods, without sacrificing solution quality. Future directions include: 1) Using a more accurate model of the 3-phase AC power flow equations to better exclude infeasible solutions. Options include the DistFlow approximation in (Baran and Wu 1989) and no-good cuts. 2) Scaling to entire city-sized distribution networks. We considered a feeder system connected to a single substation in this paper. However, distribution grids in a city can span multiple substations. In general, we expect city-sized networks can be partitioned into subproblems to reduce complexity and is a topic of future work. 3) Including a variation of the restoration problem posed by (Coffrin, Hentenryck, and Bent 2012).

Acknowledgments

The work was funded by the DOE-OE Smart Grid R&D Program in the Office of Electricity in the US Department of Energy, the DTRA Basic Research Program, and the Center for Nonlinear Studies. It was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396.

References

Baran, M., and Wu, F. 1989. Optimal Capacitor Placement on Radial Distribution Systems. *IEEE Transactions on Power Delivery* 4(1):725–734.

Bent, R.; Berscheid, A.; and Toole, G. L. 2010. Transmission network expansion planning with simulation optimization. In *AAAI-10*. AAAI Press.

Bialek, T. 2013. Private Communication with Tom Bialek of San Diego Gas & Electric.

Central Hudson Gas and Electric. 2014. `http://www.centralhudson.com/valueforourvalley/`. Accessed: Sept. 2014.

Chen, R. L.-Y., and Phillips, C. A. 2013. k-edge failure resilient network design. *Electronic Notes in Discrete Mathematics* 41(0):375 – 382.

Chen, R. L.-Y.; Cohn, A.; Fan, N.; and Pinar, A. 2014. Contingency-Risk Informed Power System Design. *IEEE Transactions on Power Systems* 29(5):2087–2096.

Coffrin, C.; Hentenryck, P. V.; and Bent, R. 2012. Last-mile restoration for multiple interdependent infrastructures. In *AAAI-12*. AAAI Press.

Delgadillo, A.; Arroyo, J.; and Alguacil, N. 2010. Analysis of Electric Grid Interdiction with Line Switching. *IEEE Transactions on Power Systems* 25(2):633–641.

EIA. 2010. http://www.meede.org/wp-content/uploads/Commercial-and-Industrial-CHP-Technology-Cost-and-Performance-Data-Analysis-for-EIA_June-2010.pdf.

Executive Office of the President. 2013. Economic benefits of increasing electric grid resilience to weather outages. Technical report, Executive Office of the President.

Garcia, P.; Pereira, J.; Carneiro, S.; da Costa, V.; and Martins, N. 2000. Three-phase power flow calculations using the current injection method. *IEEE Transactions on Power Systems* 15(2):508–514.

Garg, M., and Smith, J. C. 2008. Models and algorithms for the design of survivable multicommodity flow networks with general failure scenarios. *Omega* 36(6):1057–1071.

Garg, V. K.; Jayram, T. S.; and Narayanaswamy, B. 2013. Online optimization with dynamic temporal uncertainty: Incorporating short term predictions for renewable integration in intelligent energy systems. In *AAAI-13*. AAAI Press.

Golari, M.; Fan, N.; and Wang, J. 2014. Two-stage stochastic optimal islanding operations under severe multiple contingencies in power grids. *Electric Power Systems Research* 114(0):68 – 77.

Governor and General Assembly of Virginia. 2005. Placement of utility distribution lines underground. Technical report.

Hentenryck, P. V.; Gillani, N.; and Coffrin, C. 2012. Joint assessment and restoration of power systems. In *ECAI*, 792–797.

Jabr, R. A. 2013. Robust Transmission Network Expansion Planning With Uncertain Renewable Generation and Loads. *IEEE Transactions on Power Systems* 28(4):4558–4567.

Jain, S.; Narayanaswamy, B.; and Narahari, Y. 2014. A multi-armed bandit incentive mechanism for crowdsourcing demand response in smart grids. In *AAAI-14*, 721–727. AAAI Press.

Johnson, D. S.; Lenstra, J. K.; and Kan, A. H. G. R. 1978. The complexity of the network design problem. *Networks* 8(4):279–285.

Kersting, W. 1991. Radial distribution test feeders. *IEEE Transactions on Power Systems* 6(3):975–985.

Khushalani, S.; Solanki, J. M.; and Schulz, N. N. 2007. Optimized restoration of unbalanced distribution systems. *IEEE Transactions on Power Systems* 22(2):624–630.

Lazic, J.; Hanafi, S.; Mladenovic, N.; and Urosevic, D. 2010. Variable neighbourhood decomposition search for 0-1 mixed integer programs. *Comp. & OR* 37(6):1055–1067.

Li, J.; Ma, X.-Y.; Liu, C.-C.; and Schneider, K. P. 2014. Distribution System Restoration With Microgrids Using Spanning Tree Search. *IEEE Transactions on Power Systems* PP(99):1–9.

Mansfield, M., and Linzey, W. 2013. Hurricane sandy multi-state outage & restoration report. Technical Report 9308, National Association of State Eenergy Officials.

Munoz, F. D.; Hobbs, B. F.; Ho, J. L.; and Kasina, S. 2014. An Engineering-Economic Approach to Transmission Planning Under Market and Regulatory Uncertainties: WECC Case Study. *IEEE Transactions on Power Systems* 29(1):307–317.

Nace, D.; Piro, M.; Tomaszewski, A.; and Zotkiewicz, M. 2013. Complexity of a classical flow restoration problem. *Networks* 62(2):149–160.

Raidl, G. R.; Baumhauer, T.; and Hu, B. 2014. Speeding up logic-based benders' decomposition by a metaheuristic for a bi-level capacitated vehicle routing problem. In *HM 2014*, 183–197.

Reddy, P. P., and Veloso, M. M. 2011. Strategy learning for autonomous agents in smart grid markets. In *IJCAI*, 1446–1451.

Reddy, P. P., and Veloso, M. M. 2012. Factored models for multiscale decision-making in smart grid customers. In *AAAI-12*. AAAI Press.

Reddy, P. P., and Veloso, M. M. 2013. Negotiated learning for smart grid agents: Entity selection based on dynamic partially observable features. In *AAAI-13*. AAAI Press.

Sa, Y. 2002. *Reliability Analysis of Electric Distribution Lines*. Ph.D. Dissertation, McGill University, Montreal, Canada.

Salmeron, J.; Wood, K.; and Baldick, R. 2009. Worst-case interdiction analysis of large-scale electric power grids. *IEEE Transactions on Power Systems* 24(1):96–104.

Santoso, T.; Ahmed, S.; Goetschalckx, M.; and Shapiro, A. 2003. A stochastic programming approach for supply chain network design under uncertainty. Stochastic Programming E-Print Series. Institut fr Mathematik.

Shann, M., and Seuken, S. 2013. An active learning approach to home heating in the smart grid. In *IJCAI*.

Thibaux, S.; Coffrin, C.; Hijazi, H.; and Slaney, J. K. 2013. Planning with mip for supply restoration in power distribution systems. In *IJCAI*.

Tomaszewski, A.; Pióro, M.; and Żotkiewicz, M. 2010. On the complexity of resilient network design. *Networks* 55(2):108–118.

US Department of Energy. 2013. U.S. Energy sector vulnerabilities to climate change and extreme weather. Technical Report DOE/PI-0013, US Department of Energy.

Vanderbeck, F., and Wolsey, L. A. 2010. Reformulation and decomposition of integer programs. In *50 Years of Integer Programming*. 431–502.