

Planning with MIP for Supply Restoration in Power Distribution Systems

Sylvie Thiébaux¹, Carleton Coffrin², Hassan Hijazi¹, John Slaney¹

^{1,2}Optimisation Research Group, NICTA

¹Australian National University

²University of Melbourne

Australia

first.last@nicta.com.au

Abstract

The next generation of power systems faces significant challenges, both in coping with increased loading of an aging infrastructure and incorporating renewable energy sources. Meeting these challenges requires a fundamental change in the operation of power systems by replacing human-in-the-loop operations with autonomous systems. This is especially acute in distribution systems, where renewable integration often occurs. This paper investigates the automation of power supply restoration (PSR), that is, the process of optimally reconfiguring a faulty distribution grid to resupply customers. The key contributions of the paper are (1) a flexible mixed-integer programming framework for solving PSR, (2) a model decomposition to obtain high-quality solutions within the required time constraints, and (3) an experimental validation of the potential benefits of the proposed PSR operations.

1 Motivation

Optimisation technology is widely used in modern power systems [Momoh, 2001] and has resulted in dramatic savings [Ott, 2010] (on the order of billions of dollars annually). But the increasing role of demand response, the integration of renewable sources of energy, and the desire for more automation in fault detection and recovery pose new challenges for the planning and control of electrical power systems [Miller, 2011]. Power grids now need to operate in more stochastic environments and under varying operating conditions, while still ensuring system reliability and security.

We investigate the automation of Power Supply Restoration (PSR), a fundamental task in the operation of distribution systems. PSR consists in generating a sequence of switching operations to reconfigure a faulty network in such a way as to isolate the faults and resupply as many customers as possible as quickly as possible. PSR is subject to a range of constraints and secondary optimisation criteria, and has aggressive computational runtime requirements—minutes at most.

Currently, due to the current low level of automation in distribution systems, PSR is most often performed by human operators, sometimes aided by rule-based algorithms capable of

issuing switching recommendations in very simple fault situations. The goal of the present work is to leverage techniques from artificial intelligence and operations research to replace current rule-based systems for PSR with robust, flexible optimisation technology, capable of delivering higher-quality solutions with a greater degree of autonomy.

Existing approaches in the literature typically produce sub-optimal solutions, either by relying on incomplete optimisation methods, or by severely limiting the configurations considered. Moreover, they often ignore features which are important for automation, including multiple faults, electrical power flows, capacity constraints, or sequencing of the actions in the restoration plans. See Section 6 for more details of the related literature and references.

This paper presents a mixed-integer programming framework for PSR which does not suffer from any of these limitations. We propose to automate PSR by decomposing it into two optimisation problems. First, the problem of finding an optimal final network *configuration* is solved, ignoring intermediate plan steps. Second, a *sequencing* problem is solved to determine how best to transition the network into this optimal configuration. We observe that solving the two problems jointly to optimality does not meet the runtime requirements of PSR, even for modest network sizes and number of faults. Whilst our decomposition *may* produce suboptimal plans, our experiments show that in practice, plan utility under the two-step approach is nearly indistinguishable from the optimal.

The paper is organised as follows. Section 2 describes PSR in more detail. Section 3 deals with the problem of finding a final configuration. Section 4 deals with finding sequences of actions, either globally or under the two step approach. Section 5 presents experimental results and Section 6 concludes with detail about related and future work.

2 Power Supply Restoration

Network Structure: A medium voltage power distribution system (see Figure 1) can be viewed as a network of buses that are connected by lines equipped with switches whose position is open or closed. Conceptually, buses are the network's nodes and lines are the network's edges; an edge is disabled when the line switch is open. Power flows into the network via circuit-breakers and stops at open switches. A network element is *fed* by a (closed) circuit-breaker if there is a path between them going only through closed switches. Consumers




Figure 1: Semi-rural Network from [Thiébaux and Cordier, 2001] in its nominal radial configuration. The large squares are circuit-breakers, the small squares are switches and the thin rectangles are buses. Colors differentiate feeders.

are connected to buses and are supplied with power when their bus is fed.

Distribution networks have a meshed topology, which is often configured radially: the switches are set so that the path taken by the power from each circuit-breaker forms a distinct tree called a feeder, each element being fed by at most one circuit-breaker. However, the advent of distributed generation is gradually turning distribution systems into active meshed networks for which the radiality assumption does not hold. In this paper we consider both radial and meshed topologies, and use the word “feeder” to refer to any part of the network (radial or not) fed by the same set of circuit-breakers.

Faults: Especially in bad weather conditions, power distribution systems are frequently subject to faults (e.g. short-circuits) which cannot be eliminated by protection systems. Multiple faults are not rare (e.g. due to lightning). Such permanent faults lead the circuit-breakers feeding faulty network elements to open to protect the network from overloads. This leaves the *entire* area fed by the tripped circuit-breaker(s) without power. For instance in Figure 1, if a fault occurs on B13, CB1 will open and all consumers on buses B10 to B14 will be without power.

Restoration Plans: When faults occur, the faulty network elements must be located—a problem we assume solved—and the network reconfigured to isolate them and restore power to as many customers as possible. In this paper, we consider *automated* PSR, where reconfiguration is performed by operating *remote controlled* switches. Some switches are not manoeuvrable remotely, e.g. because they are temporarily dysfunctional or require manual operation (manual switches are not shown in Figure 1). Restoration plans are sequences of remote-controlled switching operations. Parallel plans are possible in theory, but control room operators prefer sequential plans as they can more easily detect problems during their execution.

By opening switches we may isolate suspected network elements, and by closing switches redirect power from healthy parts of the network to the areas that need supply. In the B13 fault example, we can open S12 and S13 to isolate the fault, re-close CB1 to supply buses upstream of the fault (B10, B11, B12) and close S53 to resupply the area downstream (B14) via the orange feeder (CB5). S12 must open before CB1 is closed, or CB1 will feed the fault again and reopen. Similarly, opening S13 must precede closing S53.

Automated PSR must be completed within a time bound (1–5 minutes); longer service interruptions lead to heavy fines from power regulators. Following automated PSR, repair crews are dispatched to fix the faulty equipment and operate manual switches. We do not consider repair and crew dispatching in this paper. See e.g [Coffrin *et al.*, 2012] for techniques relevant to this type of problem.

Constraints: A valid restoration plan satisfies the following constraints at each step. Firstly, no fault must be fed. Secondly, network configurations must comply with a range of power constraints including maintaining voltage, current, and power flow within certain ranges. Most importantly for present purposes, circuit-breakers and lines have capacities, which the power flow must not exceed. These may prevent redirecting power via certain paths and resupplying all customers, e.g. CB5 may not have enough capacity to resupply B14. A third relevant constraint is the radiality of the network if loops and double feeding of buses are to be avoided.

Finally, current PSR practices often restrict the set of plans considered to so-called “level-1” plans which only operate switches located on faulty feeders, including the tie switches connecting them to directly adjacent feeders. Our plan example for the fault on B13 is level-1 since CB1, S12, S13 and tie-switch S53 are on the faulty feeder. Level-1 plans can be unduly restrictive as they do not allow load on healthy feeders to be shed or transferred to help resupply faulty feeders.

In this paper, we consider level- k plans for arbitrary k . Let the fault distance of a switch (or a circuit-breaker) be the minimum number of open switches (excluding itself) that must be traversed to reach a faulty element from the switch in the state immediately following the fault. A level- k plan only operates switches with fault-distance strictly less than k . For instance if before closing S53 we had the purple feeder take part of the orange feeder’s load by opening S51 and closing S41, we would have a level-2 plan.

Objectives: A good plan will optimise certain parameters under those constraints. The primary objective is to resupply as many customers as possible, as fast as possible, giving priority to critical customers such as hospitals. Secondary objectives include minimising the number of switching operations (or more generally their cost on equipment), keeping the final network configuration close to the nominal one, and balancing load to avoid congestion during demand peaks.

Evidently, the “as fast as possible” part of the primary objective impacts on the ordering of actions in the plan. However, almost all existing literature ignores the choice of ordering and focuses on the problem of finding an optimal final configuration. As we will show, finding an optimal plan, or even optimally ordering the operations required to reach a




Figure 2: Load supplied by level-1 plans (left) and level-2 plans (right) as a function of plan step, under different operation orderings (naïve vs optimal) and network topology (radial vs meshed). “Naïve” order has all openings first and then all closings. Faults are on B14, B51, and B33 of Figure 1. In the plan descriptions, S means “close S ” and $\bar{S}$ means “open S ”. Observe that the level-1 plans (left) are shorter but resupply less load than the level-2 plans (right), that meshed topologies reduce level-2 plan length and avoid the service interruptions (curve drops) observed with radial topologies, and that naive operation ordering yields poor utility (the area under the curve is smaller).

given final configuration, are in practice harder than finding the final configuration itself.

All objectives above are affected by the topology and plan level allowed. An optimal level-1 plan is typically (but not necessarily) shorter than an optimal level- k plan for $k > 1$, but may resupply fewer customers. A meshed topology will reduce the plan length since fewer openings are required, and will often avoid temporary service interruptions when off-loading healthy feeders in plans of level > 1 . Figure 2 illustrates the effects on plan utility of operations ordering, radial vs meshed topologies, and level-1 vs level-2 plans.

3 Finding Optimal Configurations

This section describes our MIP model for finding an optimal final configuration, without sequencing actions. To simplify notation, we assume that at most one line connects a given pair of buses, that faults occur only on buses, and that circuit-breakers are associated with a set of “generation” buses at which power may be injected. In the following, we write $\mathcal{B} = 1 \dots n$ for the set of buses, $\mathcal{F} \subseteq \mathcal{B}$ for the set of faulty buses, $\mathcal{G} \subseteq \mathcal{B}$ the set of generation buses, $\mathcal{L} \subset \{(i, j) \in \mathcal{B} \times \mathcal{B} | i < j\}$ for the set of lines, and $\mathcal{S} \subseteq \mathcal{L}$ for the set of non-manoeuvrable (hence “static”) switches.¹

Power networks are unlike communication networks in that the flow of energy is governed by physical laws. The steady-state AC power flow equations² are widely accepted as an accurate model of power flow. However, these form a system of non-convex non-linear equations that can be difficult

¹We could omit static switches from the model. Instead use them to conveniently model constraints on the switches that can be operated, e.g., the restriction to level- k plans.

²Where p_{ij} , q_{ij} , $\bar{g}_{ij}$ and $\bar{b}_{ij}$ are the real power flow, reactive power flow, conductance and susceptance for line ij and $|V_i|$ and θ_i are the voltage magnitude and phase angle at bus i , the AC power flow equations are:

$$\begin{aligned} p_{ij} &= |V_i|^2 \bar{g}_{ij} - |V_i| |V_j| (\bar{g}_{ij} \cos(\theta_i - \theta_j) + \bar{b}_{ij} \sin(\theta_i - \theta_j)) \\ q_{ij} &= -|V_i|^2 \bar{b}_{ij} - |V_i| |V_j| (\bar{g}_{ij} \sin(\theta_i - \theta_j) - \bar{b}_{ij} \cos(\theta_i - \theta_j)) \end{aligned}$$

to solve and optimise. To enable fast and reliable algorithms, we (and many others [Momoh, 2001]) adopt the DC power flow approximation [Powell, 2004].³ Mathematically, this approximate flow is similar to a DC circuit: the power flow p_{ij} on a line is proportional to the line susceptance $\bar{b}_{ij}$ and the difference $(\theta_i - \theta_j)$ between the phase angles of the buses it connects, that is $p_{ij} = -\bar{b}_{ij}(\theta_i - \theta_j)$.

Basic Model: Model 1 is our basic model for finding an optimal configuration. For clarity, we describe it using logical constraints; their linearisations are obtained via the standard transformations. Constants are indicated by an overline bar.

The set $\mathcal{F}$ of faulty lines, the load $\bar{l}_i$ on each bus before the incident, as well as the pre-incident positions of the switches ($\bar{y}_{ij}$) and circuit-breakers ($\bar{x}_i$) are assumed to be given. Also given are the capacities of lines ($\bar{p}_{ij}$) and circuit breakers ($\bar{g}_i$), line susceptances ($\bar{b}_{ij}$), and maximum phase angle differences ($\bar{\theta}_{ij}$) between adjacent buses which are assumed to be small (e.g. $\pi/12$). The variables refer to the network state in the final configuration we seek. They encode the final positions of the switches y_{ij} and circuit-breakers x_i , the generation g_i at generation buses, the line flows p_{ij} , and the bus phase angles θ_i . The boolean f_i indicates whether a bus is fed, in which case all its pre-incident load $\bar{l}_i$ is served. For each line a *single* flow variable p_{ij} , $i < j$ measures the flow in the $i \rightarrow j$ direction and can therefore be negative.

Constraint (M1.3) states that faulty buses must not be fed. (M1.4) states that non-manoeuvrable switches keep their initial positions. (M1.5) encodes Kirchhoff’s current law, which enforces flow conservation at the buses. Constraints (M1.6-7) define the line flow as per the DC power

³The DC power flows are a linear approximation derived from the AC power flows through a series of approximations justified by operational considerations. In particular, the DC power flows do not capture reactive power and approximate real power by assuming 1. that conductance is small in comparison to susceptance ($\bar{g}_{ij} \simeq 0$), 2. that voltage is close to 1.0 (p.u.) ($|V_i| = |V_j| \simeq 1$) and 3. that phase angle differences are small (so $\cos(\theta_i - \theta_j) \simeq 1$ and $\sin(\theta_i - \theta_j) \simeq \theta_i - \theta_j$). Under these assumptions, the AC power flows reduce to $p_{ij} = -\bar{b}_{ij}(\theta_i - \theta_j)$.

Model 1 Optimal Configuration

Inputs:

$i \in \mathcal{B}$	$\bar{l}_i$	- load on bus i
$i \in \mathcal{B}$	$\bar{g}_i$	- generation capacity at bus i ($= 0$ if $i \notin \mathcal{G}$)
$i \in \mathcal{G}$	$\bar{x}_i$	- pre-incident position of circuit-breaker i
$(i, j) \in \mathcal{L}$	$\bar{y}_{ij}$	- pre-incident position of switch ij
$(i, j) \in \mathcal{L}$	$\bar{b}_{ij}$	- susceptance of line ij
$(i, j) \in \mathcal{L}$	$\bar{\theta}_{ij}$	- max. phase angle difference for line ij
$(i, j) \in \mathcal{L}$	$\bar{p}_{ij}$	- capacity of line ij

Variables:

$i \in \mathcal{B}$	$g_i \in [0, \bar{g}_i]$	- generation at bus i
$i \in \mathcal{B}$	$\theta_i \in (-\infty, \infty)$	- phase angle at bus i
$i \in \mathcal{B}$	$f_i \in \{0, 1\}$	- is bus i fed?
$i \in \mathcal{B}$	$x_i \in \{0, 1\}$	- position of circuit-breaker i
$(i, j) \in \mathcal{L}$	$y_{ij} \in \{0, 1\}$	- position of switch ij
$(i, j) \in \mathcal{L}$	$p_{ij} \in [-\bar{p}_{ij}, \bar{p}_{ij}]$	- flow in direction $i \rightarrow j$

Maximise:

$$1. \sum_{i \in \mathcal{B}} \bar{w}_i^l \bar{l}_i f_i \quad (\text{M1.1})$$

$$2. - \sum_{i \in \mathcal{G}} \bar{w}_i^m |x_i - \bar{x}_i| - \sum_{(i,j) \in \mathcal{L}} \bar{w}_{ij}^m |y_{ij} - \bar{y}_{ij}| \quad (\text{M1.2})$$

Subject to:

$$i \in \mathcal{F} \quad \neg f_i \quad (\text{M1.3})$$

$$i \in \mathcal{S} \quad y_{ij} = \bar{y}_{ij} \quad (\text{M1.4})$$

$$i \in \mathcal{B} \quad g_i + \sum_{j:(j,i) \in \mathcal{L}} p_{ji} = \bar{l}_i f_i + \sum_{j:(i,j) \in \mathcal{L}} p_{ij} \quad (\text{M1.5})$$

$$(i, j) \in \mathcal{L} \quad y_{ij} \rightarrow (p_{ij} = -\bar{b}_{ij}(\theta_i - \theta_j)) \quad (\text{M1.6})$$

$$(i, j) \in \mathcal{L} \quad \neg y_{ij} \rightarrow (p_{ij} = 0) \quad (\text{M1.7})$$

$$(i, j) \in \mathcal{L} \quad y_{ij} \rightarrow (-\theta_{ij} \leq \theta_i - \theta_j \leq \bar{\theta}_{ij}) \quad (\text{M1.8})$$

$$(i, j) \in \mathcal{L} \quad (f_i \neq f_j) \rightarrow \neg y_{ij} \quad (\text{M1.9})$$

$$i \in \mathcal{G} \quad \neg x_i \leftrightarrow (g_i = 0) \quad (\text{M1.10})$$

$$\theta_0 = 0 \quad (\text{M1.11})$$

flow model if the line switch is closed, and as zero if it is open. (M1.8) enforces the maximal phase angle difference between adjacent buses. Capacity constraints are enforced in the domain declarations of the flow and generation variables p_{ij} and g_i . (M1.9) states that two adjacent buses one of which is fed and one of which is not, must be separated by an open switch. This forces appropriate switches open to isolate the faulty buses and any other areas that cannot be resupplied. Similarly, (M1.10) closes circuit-breakers if they generate power and opens them otherwise. Finally, (M1.11) breaks translation symmetry among the θ_i variables.

The primary objective (M1.1) is for the final configuration to maximise the sum of the fed buses loads. Alternatively, we could replace the load by the number of customers. The weights $\bar{w}_i^l$ typically reflect the presence of critical customers. The secondary objective (M1.2) is to minimise the Hamming distance between the pre-incident and final configurations. Here the weights $\bar{w}_i^m$ and $\bar{w}_{ij}^m$ can be used to capture the cost of operating certain devices.

Model Extensions: Model 1 is suitable for meshed networks with arbitrary plan levels. However, having flexibility to support many variants of the model is critical in practice as different distribution grid operators have unique preferences and constraints on how their system is operated. We now discuss four model extensions of practical interest.

First, the model easily supports alternative objectives, such

as load balancing, which is achieved by minimising a variable α under the constraints $|g_i| \leq \bar{g}_i \alpha$ and $|p_{ij}| \leq \bar{p}_{ij} \alpha$. Second, enforcing only Level- k plans is achieved by adding all switches whose fault distance is k or more to the set $\mathcal{S}$ of non-maneuvrable switches. Fault distance can be computed in polynomial time using dynamic programming. Third, in networks where line-losses have a significant impact, a linear approximation of line-losses can be added to the power flow equations, as in [Coffrin *et al.*, 2011].

Finally, if a radial (tree) topology is required, we introduce a boolean variable z_{ij} indicating whether a strictly positive flow is allowed in direction $i \rightarrow j$. Unlike p_{ij} , there are two instances of these flow indicators per line, one in each direction and they are linked to the p_{ij} with the constraints $\neg z_{ij} \rightarrow (p_{ij} \leq 0)$, $\neg z_{ji} \rightarrow (p_{ij} \geq 0)$. The tree structure is enforced by the constraint,

$$\sum_{j:(j,i) \in \mathcal{L}} z_{ji} + \sum_{j:(i,j) \in \mathcal{L}} z_{ij} \leq 1 \quad i \in \mathcal{B} \setminus \mathcal{G}$$

which ensures that the flow entering each bus comes from at most one other bus. For generator buses, $i \in \mathcal{G}$, x_i must be added to the left-hand side. Finally, $z_{ij} + z_{ji} = y_{ij}$ ensures that flows through an open switch are not allowed and that flows through a closed switch are unidirectional.

4 Finding Switching Sequences

We now turn to the full PSR problem of finding a sequence of switching operations resupplying as much load as possible, as fast as possible, whilst complying with the constraints stated in the previous section at any step of the plan. Regardless of the network topology (radial/meshed), finding an optimal configuration is NP-complete, even when restricted to level-1 plans and when all non-faulty buses can be resupplied.⁴ Finding an optimal sequence is therefore NP-hard, but it is unknown whether it belongs to NP because there is no obvious polynomial bound on the length of the optimal plan. In practice, finding sequences appears substantially more difficult than finding configurations.

Optimal Switching Sequence: Model 2 is our MIP model for the problem of finding an operation sequence optimal over an horizon of T time steps. The primary objective (M2.1) is to maximise the area under the curve of the load supplied as a function of the time t for $1 \leq t \leq T + 1$. The number (or cost) of switchings, and then the Hamming distance between the pre-incident and final configuration are used to break ties.

The variables and constraints include copies of those of Model 1, one copy per time step. In addition, $2 \times T$ boolean variables per device (i.e. switch or circuit-breaker) indicate whether opening (o_{ij}^t and o_i^t) or closing (c_{ij}^t and c_i^t) the device occurs at time t . Additional constraints (M2.4-5) state that the initial positions of the circuit breakers at time 1 are those right *after* the incident (since the incident might have caused some circuit-breaker openings), while switches keep their pre-incident positions. (M2.6-7) update the positions of the devices at each time step, taking into account possible closing/opening. (M2.8) allows at most one operation per

⁴This can be proven by reduction from BIN-PACKING.

time step. Since the utility at the end of a non-empty optimal plan is always strictly greater than at the start, the objective suffices to force the plan to be contiguous and packed at the start of the timeline.

Model 2 Optimal Switching Sequence

Inputs:

same inputs as in Model 1 and add:

$i \in \mathcal{G}$ $\bar{x}'_i$ - post-incident position of circuit-breaker i
 T - plan horizon

Variables:

$T+1$ copies v^t of the variables of Model 1 indexed by t and add:

$t < T, i \in \mathcal{G}$ $o_i^t \in \{0, 1\}$ - breaker i opens at time t ?
 $t < T, i \in \mathcal{G}$ $c_i^t \in \{0, 1\}$ - breaker i closes at time t ?
 $t < T, (i, j) \in \mathcal{L}$ $o_{ij}^t \in \{0, 1\}$ - switch ij opens at time t ?
 $t < T, (i, j) \in \mathcal{L}$ $c_{ij}^t \in \{0, 1\}$ - switch ij closes at time t ?

Maximise:

$$1. \sum_{t \leq T+1} \sum_{i \in \mathcal{B}} \bar{w}_i^t l_i f_i^t \quad (\text{M2.1})$$

$$2. - \sum_{t \leq T} \left(\sum_{i \in \mathcal{G}} \bar{w}_i^t (c_i^t + o_i^t) + \sum_{(i,j) \in \mathcal{L}} \bar{w}_{ij}^t (c_{ij}^t + o_{ij}^t) \right) \quad (\text{M2.2})$$

$$3. - \sum_{i \in \mathcal{G}} |x_i^{T+1} - \bar{x}_i| - \sum_{(i,j) \in \mathcal{L}} |y_{ij}^{T+1} - \bar{y}_{ij}| \quad (\text{M2.3})$$

Subject to:

$T+1$ copies of the constraints of Model 1, variables indexed by t

$$i \in \mathcal{G} \quad x_i^1 = \bar{x}'_i \quad (\text{M2.4})$$

$$(i, j) \in \mathcal{L} \quad y_{ij}^1 = \bar{y}_{ij} \quad (\text{M2.5})$$

$$t < T, i \in \mathcal{G} \quad x_i^{t+1} = x_i^t + c_i^t - o_i^t \quad (\text{M2.6})$$

$$t < T, (i, j) \in \mathcal{L} \quad y_{ij}^{t+1} = y_{ij}^t + c_{ij}^t - o_{ij}^t \quad (\text{M2.7})$$

$$t < T \quad \sum_{i \in \mathcal{G}} c_i^t + o_i^t + \sum_{(i,j) \in \mathcal{L}} c_{ij}^t + o_{ij}^t \leq 1 \quad (\text{M2.8})$$

Two-Step Decomposition: As will be apparent from our experiments, optimal sequences cannot be found in real-time (minutes) for interesting problem sizes. We therefore investigate a more scalable two-step approach: compute the optimal configuration as in Section 3 and then order the operations in the *shortest* plan required to achieve it. Clearly, the only operations in the shortest plan are to toggle the devices whose initial and final positions differ. This can always be done without violating the constraints, just by performing all of the openings before the closings. This naïve ordering gives plans of poor utility, so instead we optimally order the actions in the shortest plan. It is unknown whether this problem is in P, but it is in NP since the plan length is bounded by the number of devices hence plan validity and utility can be checked in polynomial time. In any case, finding an optimal ordering can be accommodated as a slight variation of Model 2. Let $\hat{x}_i$ and $\hat{y}_{ij}$ be the device positions in the final configuration. We set the plan horizon to:

$$T = |\{i \in \mathcal{G} \mid \hat{x}_i \neq \bar{x}'_i\} \cup \{(i, j) \in \mathcal{L} \mid \hat{y}_{ij} \neq \bar{y}_{ij}\}|;$$

we remove objectives (M2.2-3) and add constraints:

$$\sum_{t \in T} c_{ij}^t = (\hat{y}_{ij} \wedge \neg \bar{y}_{ij}), \quad \sum_{t \in T} o_{ij}^t = (\neg \hat{y}_{ij} \wedge \bar{y}_{ij}), \\ \sum_{t \in T} c_i^t = (\hat{x}_i \wedge \neg \bar{x}'_i), \quad \sum_{t \in T} o_i^t = (\neg \hat{x}_i \wedge \bar{x}'_i).$$

5 Experimental Results

Our experiments aim at illustrating a number of important features of the problems and approaches we considered: (1)

that despite its much greater generality, our MIP model for computing the optimal final configuration is competitive with state of the art heuristic search approaches; (2) that the two-step approach is much more practical than computing the optimal switching sequence; (3) that naïve sequencing leads to significant degradation of plan utility whilst optimised sequencing leads to plans whose utility is hardly distinguishable from that of the optimal sequence; and (4) that meshed topologies and arbitrary plans are advantageous in terms of utility over radial topologies and level-1 plans. All experiments were run on an Intel Core i7 2.60GHz processor with 8GB of memory. Problems were solved with Gurobi 5.0.1 with its default settings [Gurobi Optimization Inc., 2012].

Small Networks: Our first set of experiments is targeted at (2) and (3). To be able to compute the optimal switching sequence, we considered small networks: the semi-rural network in figure 1 but with the manual switches and corresponding buses included, which has 7 breakers, 45 lines and 45 buses; we also consider the part left of b13 of the smaller suburban network in [Botea *et al.*, 2012, Fig 4.], with 10 breakers, 37 buses and 38 lines.

The results are presented as a function of the percentage of faulty buses, and each point in the graphs is the average over 50 randomly generated fault sets with a given percentage. We experimented with up to 50% faulty buses, even though it is unlikely that automated PSR would be used in disasters of this magnitude. All problems were attempted with all combinations of network topology (meshed/radial), plan level (level-1, unrestricted), and approach (optimal sequence, naïve two-steps and optimised two-steps), leading to 6000 problem instances per network. To compute the optimal switching sequence, plan horizon was set to 15 for semi-rural and 25 for suburban. Following generation, the utilities of all plans produced by all approaches were evaluated with these horizons. We imposed a run-time limit per problem of 2 min for semi-rural and 5 min for suburban.

The left column of Figure 3 shows some of the results for semi-rural with unrestricted plans. Results with level-1 plans and with the suburban network are very similar. The top graph shows the average run-times of the 3 approaches as a function of the percentage of faulty buses. There is an order of magnitude difference between the naïve ordering approach (< 0.1 sec), the optimised two steps approach ($\simeq 1$ sec) and the optimal ($\simeq 100$ sec). Run-time (and plan length) increase with the number of faults until a critical point is reached where the problem becomes so overconstrained that plan length starts decreasing. Radial plans are longer (especially in the unrestricted plan case), and harder to generate.

The bottom graph shows the average of the ratio of the utility⁵ obtained with the optimised two-steps approach (resp. the naïve ordering) to that of the optimal. It is clear from the graph that the the average utility of the optimised two steps approach is extremely close to that of the optimal, whilst the performance ratio of the naïve ordering degrades to 0.7.

⁵For the graphs to be more meaningful, the load supplied right after the incident corresponds to a baseline utility of 0.




Figure 3: Run-Time and Utility Comparisons

Large Networks: Experiments on larger networks target (1) and (2), and also (3) albeit only with the two-step approach, since optimal sequence generation is not feasible. We used the larger suburban network in [Botea *et al.*, 2012] which has 81 breakers, 210 lines and 207 buses, with a more realistic number of up to 10 faults.

The run-time of the naive two-step approach is dominated by that of solving the optimal final configuration MIP model. It is therefore worthwhile to compare its efficiency with that of the heuristic search approach in [Botea *et al.*, 2012] which only considers the problem of finding a final configuration re-supplying buses downstream of faults assumed to be already isolated, assuming radiality, level-1 plans, and that all non-faulty buses can be resupplied. Botea *et al.* [2012] report run-times of the order of 1 sec for their large network. As can be seen from the naive two-steps run-time curves in the top graph of the middle column of Figure 3, our more general approach can find optimal final configurations without any of the restrictive assumptions in ≈ 0.5 sec.

As expected, the optimised two-step approach is exponentially slower, but remains practical for any realistic number of faults. Moreover, the bottom graph in the middle column shows that the ratio of the naive to optimised two-step plan utility degrades further for larger networks. Our experiments with intermediate sized networks confirm this trend.

Network Topology and Plan Level: Now targeting (4), the right-hand column of Figure 3 illustrates the tradeoffs between the radial and meshed topology and the level 1 and unrestricted plans. The top graph shows the ratio of the utility obtained with a radial topology to that obtained with a meshed topology for the 4 possible combinations of two-step approaches and level-1/unrestricted plans. In all settings the meshed topology leads to utility gains. The gain becomes significant when unrestricted plans are generated with the naive

ordering. This is because radiality increases plan length and creates service interruptions on healthy buses in all but level-1 plans, and that the effects of these two factors are amplified by the poor quality of the ordering.

The bottom figure shows the ratio of utility of level-1 plans to that of unrestricted plans. We observe a similar effect where the utility of unrestricted plans is significantly worse than that of level-1 plans, when used in combination with radial topology and naive ordering. For all other settings, unconstrained plans yield a utility gain of up to 5%. Note that a 5% gain can translate into fines being averted.

6 Conclusion, Related and Future Work

In this paper we have proposed an efficient and general MIP framework for automating PSR. It extends previous work by incorporating the DC power flow model, allowing general configurations, and sequencing operations to obtain high-quality solutions within the required scale and time constraints. This lifts several restrictions in previous research and is a significant step towards automating the next generation power distribution systems. The framework was validated on several power distribution system benchmarks. The results indicate that our approach can find good switching sequences for networks with several hundred buses in real time, and that PSR systems would benefit by allowing level- k plans and meshed network topologies both enabled by this work.

Related Work: Literature on PSR is abundant and includes methods based on special purpose procedures and rule-based systems [Liu *et al.*, 1988; González *et al.*, 2011], meta-heuristics [Fukuyama, 1996; Toune *et al.*, 2002; Carvalho *et al.*, 2006], state space search [Morelato and Monticelli, 1989; Botea *et al.*, 2012], knowledge compilation [Hadzic *et al.*, 2007], AI planning [Bertoli *et al.*, 2002; Hoffmann *et al.*,

2006], and mathematical programming [Ciric and Popovic, 2000]. However, none of these are as comprehensive or general as the framework proposed here.

In particular, the vast majority of methods are limited to finding a final *configuration* of the network.⁶ They do not consider the critical issue of *sequencing* switching operations. An exception considering sequencing is [Carvalho *et al.*, 2006], which advocates greedily pairing openings and closings required to achieve a suboptimal radial configuration and ordering the pairs “using dynamic programming”, without explaining how. No indication of the quality of the resulting plans is provided.

Even then, existing approaches to finding final configurations typically rely on incomplete methods that produce suboptimal solutions, or on problem simplifications such as ignoring power flows and capacity constraints, restricting the search to level-1 plans, assuming a single fault, or assuming that all non-faulty lines can be resupplied. None adopt the DC power flow model, which prevents them from considering meshed network topologies. For instance, the AI planning work in [Hoffmann *et al.*, 2006] completely ignores power flows and capacity constraints, resulting in a polynomial-time problem [Helmert, 2006]. The heuristic search approach by [Botea *et al.*, 2012] is restricted to radial topologies, level-1 plans, and assumes that all non-faulty buses can be resupplied. Our results show that our optimal MIP model can find level-*k* plans for arbitrary network topologies faster than heuristic search can find level-1 plans for radial topologies.

MIP has been used to find a final configuration by e.g. [Nagata *et al.*, 1995; Ciric and Popovic, 2000]. However, in addition to the limits mentioned above, these formulations assume that the fault is already isolated. They decide which lines are fed in the final configuration, but not the optimal set of switching operations required to achieve such feeding.

Another line of work enables interactive reconfiguration by compiling power flow and other constraints into a BDD which can quickly indicate to an operator that a switching operation would violate constraints [Hadzic *et al.*, 2007]. Unlike the problem considered here, switching operations are not automatically chosen. At present, this work is limited to radial configurations and discretised flow values based on a max-flow model (as opposed to DC power flows). It would be interesting to extend it to more accurate power flow models that can incorporate line losses.

Researchers have also considered the related problem of repairing power transmission systems following a major disaster, in such a way as to maximise the area under the curve of the restored network *capacity* as a function of time [Van Hentenryck *et al.*, 2011; Coffrin *et al.*, 2012]. The focus is on the selection of the next item to repair in such a way as to maximise the maximum power flow the network can absorb, under constraints stemming from the repair crew routing problem. No consideration is given to fault isolation and to the low-level switching operations required to achieve the maximum power flow after each repair.

⁶In many cases, the problem considered is further limited to finding a configuration resupplying areas downstream of faults assumed to be already isolated.

Future Work: In the future, we would like to validate our results against AC power flows and experiment with more complex power flow models, including more elaborate linear models which incorporate reactive power and voltage [Coffrin and Van Hentenryck, 2012], quadratic convex models similar to [Taylor and Hover, 2012], or even directly with non-linear AC models. Note that the class of mathematical programming model (MIP, MIQP, MINLP) obtained with our approach depends only on the power flow model considered and not on other aspects of the PSR problem. We would also like to incorporate other features of power systems in our model, including transformers, shunts, capacitors and switches that cannot be operated under load.

We also plan to resolve the open questions about the complexity of finding the optimal switching sequence and the optimised ordering. Especially if the former turns out to be PSPACE-complete, it would be interesting to address it by considering automated planning technology to handle the problem at a topological level, and mathematical programming to check capacity and other numerical constraints and provide bounds to the planner. An instantiation of the planning modulo theory framework in [Gregory *et al.*, 2012] might be suitable for this purpose (see [Piacentini *et al.*, 2013] for a first use of PMT in power systems optimisation).

Finally, the assumption that fault locations are perfectly known is an important limitation of our current framework. As research in this direction indicates, uncertainty in the fault location considerably complicates PSR [Thiébaux *et al.*, 1996; Bertoli *et al.*, 2002; Bonet and Thiébaux, 2003; González *et al.*, 2011]. Existing works on the joint fault location/restoration problem either ignore power flows altogether, rely on ad-hoc procedures, or do not produce plans that actively try to gain information about fault location. Addressing the joint problem in its full generality is a major challenge which we look forward to investigating next.

Acknowledgements

This work is supported by NICTA’s Optimisation Research Group, as part of its *Future Energy Systems* project. We thank the project members and Dana Nau for useful discussions. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

References

- [Bertoli *et al.*, 2002] P. Bertoli, A. Cimatti, J.K. Slaney, and S. Thiébaux. Solving power supply restoration problems with planning via symbolic model checking. In *Proc. 15th European Conference on Artificial Intelligence (ECAI)*, pages 576–580, 2002.
- [Bonet and Thiébaux, 2003] B. Bonet and S. Thiébaux. GPT meets PSR. In *Proc. 13th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 102–112, 2003.
- [Botea *et al.*, 2012] A. Botea, J. Rintanen, and D. Banerjee. Optimal reconfiguration for supply restoration with

- informed A* search. *IEEE Transactions on Smart Grid*, 3(2):583–593, 2012.
- [Carvalho *et al.*, 2006] P.M.S. Carvalho, L.A.F.M. Ferreira, and L.M.F. Barruncho. Optimization approach to dynamic restoration of distribution systems. *Electrical Power & Energy Systems*, 29:222–229, 2006.
- [Ciric and Popovic, 2000] R.M. Ciric and D.S. Popovic. Multi-objective distribution network restoration using heuristic approach and mix integer programming model. *Electrical Power & Energy Systems*, 22:497–505, 2000.
- [Coffrin and Van Hentenryck, 2012] C. Coffrin and P. Van Hentenryck. A linear-programming approximation of AC power flows. *CoRR*, abs/1206.3614, 2012.
- [Coffrin *et al.*, 2011] C. Coffrin, P. Van Hentenryck, and R. Bent. Approximating line losses and apparent power in AC power flow linearizations. In *Proc. Power & Energy Society General Meeting (PES)*, 2011.
- [Coffrin *et al.*, 2012] C. Coffrin, P. Van Hentenryck, and R. Bent. Last-mile restoration for multiple interdependent infrastructures. In *Proc. 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 455–463, 2012.
- [Fukuyama, 1996] Y. Fukuyama. A parallel genetic algorithm for service restoration in electric power distribution systems. In *Proc. IEEE International Conference on Fuzzy Systems*, pages 275–282, 1996.
- [González *et al.*, 2011] A. González, F.M. Echavarren, L. Rouco, T. Gómez, and J. Cabetas. Fault location and service restoration method for large-scale distribution networks. In *Proc. Power & Energy Society General Meeting (PES)*, 2011.
- [Gregory *et al.*, 2012] P. Gregory, D. Long, M. Fox, and J.C. Beck. Planning modulo theories: Extending the planning paradigm. In *Proc. 22nd International Conference on Automated Planning and Scheduling (ICAPS)*, 2012.
- [Gurobi Optimization Inc., 2012] Gurobi Optimization Inc. Gurobi optimizer reference manual, 2012.
- [Hadzic *et al.*, 2007] T. Hadzic, A. Wasowski, and H. Reif Andersen. Techniques for efficient interactive configuration of distribution networks. In *Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 100–105, 2007.
- [Helmert, 2006] M. Helmert. New complexity results for classical planning benchmarks. In *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 52–62, 2006.
- [Hoffmann *et al.*, 2006] J. Hoffmann, S. Edelkamp, S. Thiébaux, R. Englert, F. dos S. Liporace, and S. Trüg. Engineering benchmarks for planning: the domains used in the deterministic part of IPC-4. *J. Artif. Intell. Res. (JAIR)*, 26:453–541, 2006.
- [Liu *et al.*, 1988] C.C. Liu, S.J. Lee, and S.S. Vankata. An expert system operational aid for restoration and loss reduction of distribution systems. *IEEE Transactions on Power Systems*, 3:619–626, 1988.
- [Miller, 2011] J. Miller. Power system optimization smart grid, demand dispatch, and microgrids. <http://www.netl.doe.gov/smartgrid/referenceshelf>, 2011.
- [Momoh, 2001] J.A. Momoh. *Electric Power System Applications of Optimization (Power Engineering (Willis))*. CRC Press, 2001.
- [Morelato and Monticelli, 1989] A.L. Morelato and A.J. Monticelli. Heuristic search approach to distribution system restoration. *IEEE Transactions on Power Delivery*, 4, 1989.
- [Nagata *et al.*, 1995] T. Nagata, H. Sasaki, and R. Yokoyama. Power system restoration by joint usage of expert system and mathematical programming approach. *IEEE Transactions on Power Systems*, 10(3):1473–1479, 1995.
- [Ott, 2010] A. Ott. Unit commitment in the PJM day-ahead and real-time markets. <http://www.ferc.gov/eventcalendar/Files/20100601131610-Ott,%20PJM.pdf>, 2010.
- [Piacentini *et al.*, 2013] C. Piacentini, V. Alimisis, M. Fox, and D. Long. Combining a temporal planner with an external solver for the power balancing problem in an electricity network. In *Proc. 23rd International Conference on Automated Planning and Scheduling (ICAPS)*, 2013.
- [Powell, 2004] L. Powell. *Power System Load Flow Analysis*. McGraw Hill professional, 2004.
- [Taylor and Hover, 2012] J.A. Taylor and F.S. Hover. Convex models of distribution system reconfiguration. *IEEE Transactions on Power Systems*, 27(3):1407–1413, 2012.
- [Thiébaux and Cordier, 2001] S. Thiébaux and M.-O. Cordier. Supply restoration in power distribution systems — a benchmark for planning under uncertainty. In *Proc. 6th European Conference on Planning (ECP)*, pages 85–95, 2001.
- [Thiébaux *et al.*, 1996] S. Thiébaux, M.-O. Cordier, O. Jehl, and J.-P. Krivine. Supply restoration in power distribution systems: A case study in integrating model-based diagnosis and repair planning. In *Proc. 12th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 525–532, 1996.
- [Toune *et al.*, 2002] S. Toune, H. Fudo, T. Genji, Y. Fukuyama, and Y. Nakanishi. Comparative study of modern heuristic algorithms to service restoration in distribution systems. *IEEE Transactions on Power Systems*, 17:173–181, 2002.
- [Van Hentenryck *et al.*, 2011] P. Van Hentenryck, C. Coffrin, and R. Bent. Vehicle routing for the last mile of power system restoration. In *Proc. 17th Power Systems Computation Conference (PSCC)*, 2011.