LA-UR-11-06749

Approved for public release; distribution is unlimited.

Title: Smart Load and Generation Schedule for Power System Restoration

Author(s): Carleton Coffrin
Pascal van Hentenryck
Russell Bent

Intended for: Power Engineering Society General Meeting 2012

Smart Load and Generation Scheduling for Power System Restoration

Carleton Coffrin, *Student Member, IEEE*, Pascal Van Hentenryck, *Member, IEEE*, and Russell Bent

***Abstract*—This paper studies the applicability of the linearized DC model in optimizing power restoration after significant network disruptions. In such circumstances, no AC base-point solution exists and the objective is to maximize the served load. The paper demonstrates that the accuracy of the linearized DC model degrades with the size of the disaster and that it can significantly underestimate active and apparent power. To remedy these limitations, the paper proposes an Angle-Constrained DC Power Flow (ACDCPF) model that enforces constraints on the line phase angles and has the ability to shed load and generation across the network. Experimental results on N-3 contingencies in the IEEE30 network and power restoration instances from disaster recovery show that the ACDCPF model provides significantly more accurate approximations of active and apparent power. In the restoration context, the ACDCPF model is shown to be much more reliable and produces significant reduction in the size of the blackouts.**

***Index Terms*—power flow, dc power flow, power system analysis, power system restoration.**

NOMENCLATURE

$\lVert\widetilde{V}_i\rVert$	Voltage magnitude of bus i, volts
θ°_i	Phase angle of bus i, radians
θ°_{ij}	Phase angle for line i to j, i.e., $\theta^\circ_i - \theta^\circ_j$
$\widetilde{Z}$	Impedance
x	Reactance
$\widetilde{Y}_{bus}$	The nodal admittance matrix
$b^y(i,j)$	A susceptance from the $\widetilde{Y}_{bus}(i,j)$ matrix
$g^y(i,j)$	A conductance from the $\widetilde{Y}_{bus}(i,j)$ matrix
p_i	Active power at bus i, MW
q_j	Reactive power at bus i, MVar
p_{ij}	Active power on a line from i to j, MW
q_{ij}	Reactive power on a line from i to j, MVar
$c(i,j)$	Capacity on a line from i to j, MVA
$\mathcal{PN}$	A power network
N	A set of buses from a power network
L	A set of lines from a power network

I. INTRODUCTION

RESTORING a power system after a significant disruption (e.g., a cascading blackout or a seasonal hurricane) is an important task with consequences on both human and economic welfare. Power system components must be repaired and then re-energized without causing additional network instability. The restoration effort should be prioritized to minimize the size of the blackout by jointly optimizing repairs and power restoration. Our earlier work [1] approached this joint optimization using a sequence of optimization problems based on the linearized DC model.

Power restoration problems are daunting for a variety of reasons. First, since no *typical* operating base point is known, solving the resulting AC power flow problems is often challenging [2]. Second, good restoration plans jointly optimize the routing of repair crews and the scheduling of component energizing. The resulting optimization is a mixed integer nonlinear programming problem which is extremely challenging computationally. As a consequence, the power restoration algorithm proposed in [1] uses the linearized DC model in several steps, which makes it possible to model the problem in terms of mixed integer programs instead of mixed integer nonlinear programs.

The linearized DC model has been adopted as a general-purpose tool for a variety of power system optimizations in recent years (e.g., [2], [3]). However, its accuracy has been the topic of much discussion: Most papers (e.g., [2], [4]) take an optimistic outlook, while others (e.g., [5], [6]) are more pessimistic. This issue is of particular interest for power restoration which involves human and economic welfare. It is critical that the linearized DC solution be a reasonable approximation of a high-quality AC power flow solution to avoid causing additional network instability.

This paper studies the adequacy of using the linearized DC model for power restoration. It shows that, for power restoration, the linearized DC model may underestimate line loading significantly and produce solutions that are not feasible in an AC solver. Moreover, the experimental results suggest that large line phase angles are a good indicator of inaccurate active and apparent power estimations. The paper then proposes an Angle-Constrained DC Power Flow (ACDCPF) model that enforces constraints on the line phase angles and has the ability to shed load and generation across the network. The practicality of the approach is demonstrated on more than 11,000 damage contingencies in the IEEE30 network and validated on real-world power restoration problems arising in disaster management. The paper shows that the ACDCPF model produces solutions that are highly correlated with the AC power flow model and that these improvements in accuracy come with a reasonably small cost in load shedding. In particular, in the restoration context, the ACDCPF model is shown to be much more reliable and to produce significant reductions in the size of the blackouts compared to the linear DC model.

The rest of the paper is organized as follows, Section II gives a brief review on power system modeling. Section III motivates

C. Coffrin and P. Van Hentenryck are with the Department of Computer Science, Brown University, Providence RI 02912, USA.

R. Bent is with Los Alamos National Laboratory, Los Alamos NM 87545.

the paper by presenting a short review of the restoration order problem. Section IV investigates the accuracy of the linearized DC model under multiple line outages. Section V discusses the phase angle constraints and their impact. Section VI presents our ACDCPF model and its empirical evaluation. Section VII discusses a number of issues not addressed elsewhere in the paper. Section VIII concludes the paper and discusses future work.

II. Power System Modeling - A Brief Review

a) AC Power Flow: The ground truth in this paper is the single-phase AC power model, which is widely accepted as a high-quality approximation of the steady-state behavior of real-world power flows. The single-phase AC power flow model uses Kirchhoff's current law and Ohm's law in the complex plane to define the active and reactive power injected at each power bus n, i.e.,

$$p_n = \sum_{m \in N} ||\widetilde{V}_n||||\widetilde{V}_m||(g^y(n,m)\cos(\theta^\circ_{nm}) + b^y(n,m)\sin(\theta^\circ_{nm}))$$
$$q_n = \sum_{m \in N} ||\widetilde{V}_n||||\widetilde{V}_m||(g^y(n,m)\sin(\theta^\circ_{nm}) - b^y(n,m)\cos(\theta^\circ_{nm})).$$

These AC power flow equations can be solved by iterative solution techniques such as the Newton-Raphson method [7], [8], [9]. Convergence of these methods is not guaranteed and, when the system is heavily loaded, the solution space is riddled with infeasible low-voltage solutions which are not usable in practice [10]. In fact, finding a solution to the AC power flow when a base-point solution is unavailable is often "maddeningly difficult" [2].

b) The Linearized DC Model: The linearized DC model is derived from the AC power flow model through a series of approximations justified by operational considerations. In particular, it is assumed that (1) the susceptance is large relative to the impedance $|b(n,m)| \gg |g(n,m)|$; (2) the phase angle difference θ°_{nm} is small enough to ensure $\sin(\theta^\circ_{nm}) \approx \theta^\circ_{nm}$; and (3) the voltage magnitudes $||\widetilde{V}||$ are close to 1.0 and do not not vary significantly. Under these assumptions, the AC power flow equations reduce to

$$p_n = \sum_{m \in N}^{n \neq m} b^y(n,m)(\theta^\circ_n - \theta^\circ_m) \quad (1)$$

From a computational standpoint, the linearized DC model is much more appealing than the AC model: It forms a system of linear equations that admit reliable algorithms. These linear equations can also be embedded into optimization frameworks for decision support in power systems [1], [2], [3], [11], [12], [13]. However, it is important to verify whether the assumptions of the linearized DC model holds for each application domain.

c) Implementation Choices: The linearized DC model is so pervasive that authors often forget to mention important implementation details. Indeed, reference [5] recently demonstrated that small changes to the model formulation may have a significant impact on its accuracy. Moreover, there are conflicting suggestions about how the $\widetilde{Y}_{bus}$ matrix is derived (e.g., using $1/x$ or $-\Im(\widetilde{Z}^{-1})$) [14], [15], [16], [17]). Our goal is to make the AC and DC power models as similar as possible and our implementation reflects this choice.

In particular, we use the same susceptance value $b^y(n,m)$ in the AC and DC models and adopt the $\widetilde{Y}_{bus}$ calculation described and implemented in MATPOWER [18].

By necessity, the AC solvers use a slack bus to ensure the flow balance on the network when the total power consumption is not known a priori (due to line losses for instance). As a consequence, we also use a slack bus for the various DC models considered in this paper so that the AC and DC models can be accurately compared. It should be emphasized that the ACDCPF model proposed later in the paper does not need a slack bus and the only reason to use a slack bus in the ACDCPF model is to allow for meaningful comparisons between the DC and AC models. This issue is discussed in more length in Section VII.

III. Power Restoration and Linearized DC Models

The research in this paper is motivated by the joint repair and restoration of the transmission network after significant damages from a natural disaster. The goal is to schedule the repair and to re-energize the electrical components in order to minimize the size of the blackout. This joint repair/restoration problem is extremely challenging computationally and is solved through a sequence of optimization models [1]. Several of the models in the sequence use a linearized DC model and it is legitimate to question whether this is adequate in the presence of significant disruptions of the transmission network.

For the purpose of this paper, it is sufficient to consider only one of the optimization problems proposed in [1]: The Restoration Order Problem (ROP). Conceptually, the ROP can be formulated as follows: A collection of D power system components have been disrupted and must be re-energized one at a time. The goal is to find a restoration order $d_1, d_2, d_3, \ldots$ for all the components $d_i \in D$ in order to maximize the served load over time. In [1], the ROP is modeled as a generalization of Optimal Transmission Switching (OTS) (e.g., [11], [19]). More precisely, the ROP is built from a collection of OTS models, one for each restoration step, which are connected together to optimize the restoration order globally.

Ideally, the ROP problem should be solved using an AC power flow model. However, simply finding an AC power flow solution for such disruptions can be quite challenging, since no base-point solution is available. Furthermore, ROP and OTS problems require discrete decision variables to indicate which components are energized, producing highly complex mixed integer nonlinear programming models that are beyond the capabilities of existing optimization technology. As a result, OTSs and ROPs are modeled in terms of linearized DC models. Note that there is a significant difference between the OTS and the ROP models. In OTS models, lines are switched from a current working base-point, while the ROP has no base-point because the disruption is arbitrary and extensive. As a result, it is not clear how accurate the resulting DC model is and whether it can be used in practice for restoration problems. This paper answers both questions.

Model 1 Linearized DC Power Flow (LDC).

Inputs:

$\mathcal{PN} = \langle N, L \rangle$ the power network
b^y susceptance from a $\widetilde{Y}_{bus}$ matrix
s slack bus index

Variables:

$\theta^\circ_i \in (-\infty, \infty)$ - phase angle on bus i (radians)

Subject to:

$$\theta^\circ_s = 0 \quad \text{(M1.1)}$$

$$p_n = \sum_{m \in N}^{n \neq m} b^y(n,m)(\theta^\circ_n - \theta^\circ_m) \quad \forall n \in N \;\; n \neq s \quad \text{(M1.2)}$$

To compare the DC and AC models for ROPs, we can ignore the optimization process and focus on the solution returned, i.e., a restoration ordering. The quality of an ordering can be evaluated by a series of power flow calculations. Each step implements a change in the network topology since an additional component comes online and the power flow calculation gives the increase in served loads (or, equivalently, the reduction in the blackout). Therefore, to understand the accuracy of the linearized DC model in the restoration context, it is sufficient to study the accuracy of the linearized DC model in isolation when it is subject to significant topological changes. Model 1 presents the linearized DC model (LDC) implementing our modeling assumptions. The model takes as inputs a power network $\mathcal{PN}$, susceptance values b^y, and the index s of the slack bus. The goal is to find the phase angles of all buses. Constraint M1.1 fixes the phase angle of the slack bus to 0 and constraint M1.2 implements the power flow model as defined in Equation 1. It is important to note that the power balance constraint is not posted for the slack bus. This allows the slack bus to pick up any unserved loads and balance the power in the network as is done in a traditional AC power flow model.

IV. DC Power Flow with Multiple Line Outages

To understand the accuracy of the linearized DC model under significant disruptions, we begin with a comprehensive empirical study of the IEEE30 system. We consider 11,521 damage contingencies, some with as many as three line outages (about 7% of the total network). Despite its ubiquity for optimization with N-1 contingency constraints, the accuracy of the linearized DC model has only been evaluated for specific application domains (e.g., [2], [6]). To our knowledge, this is the first direct study of the DC model accuracy for N-1, N-2, and N-3 contingencies.[1]

To compare AC and DC models, we measure the same information (e.g., active power, phase angles, ...) in both models and plot their correlation. Specifically, for some measurement data (e.g., the active power of a line), the x-axis gives the data value in the AC model and the y-axis gives the data value in the DC model. As a result, the closer the plot is to the line $x = y$, the more the AC and DC models agree. We will focus primarily on apparent power since it is of particular interest to applications with line capacities. Obviously, in the DC model, apparent power is approximated by active power. The AC model is initialized with the voltages set to 1.0 and the phase angles to zero. For N-3 contingencies, it fails to converge in about 1% of the cases, as described later in the paper.

Figure 1 presents the correlation of apparent power for all N-1, N-2, and N-3 outages on the IEEE30 benchmark, giving us a total of 11,521 damaged networks. Each data point in the plots represents the apparent power of a line and, for brevity, the results are grouped by the number of outages and superimposed onto the same correlation graph. The plots also use red triangles for data points obtained from networks that feature *large* line phase angles (i.e., $|\theta^\circ_{nm}| > \pi/12$). This makes it possible to understand the link between large line phase angles and discrepancies in apparent power.

Figure 1 highlights a number of interesting phenomena. First, observe that the overall accuracy of the model degrades significantly as the number of damaged components increases, which is concerning for power restoration applications. Second, the linearized DC model underestimates apparent power systematically and the more significant errors are almost always associated with large line phase angles. Finally, the plots indicate a general trend for the apparent power to lean to the right for large line loads. This is due to line losses which are not captured in the DC model. This limitation can be addressed in the linearized DC model as discussed in [2], [5], [21], [20] and these solution techniques are completely orthogonal to the proposals in this paper.

Figure 2 dives deeper into these results and investigates the worst-case damage scenarios (i.e., N-3 contingencies) in more detail by presenting results for active power (left), bus phase angles (center), and reactive power (right). Once again, the color red represents large line phase angles. The left plot depicts the correlation of active power. It indicates that the linearized DC model underestimates active power on large line loads, which are also characterized by large line angles. The center plot depicts the correlation of the bus phase angles. It shows that the linearized DC model systematically underestimates the bus phase angles and the errors increase with large line phase angles. The right histogram depicts the number of lines whose reactive power fall within a certain range in the AC power flow. The color of each bar reflects the percentage of data points which are marked red in the other plots. The histogram reveals that N-3 contingencies produce a significant amount of reactive power on many lines, almost all of which exhibit large line phase angles. These results thus indicate that large line angles are correlated with both under-approximations of active power and large reactive power, and hence produce significant errors in estimating apparent power.

To increase our intuition, it is also worthwhile to consider one particular bus from the IEEE30 system. Bus 1 in the IEEE30 is connected to buses 2 and 3 with impedances of $0.0192 + i0.0575$ and $0.0452 + i0.1652$ respectively (and thus susceptance values of -15.65 and -5.632 respectively). We investigate the N-1 contingencies around this bus to show how large angle differences and reactive flows are connected. Table I presents the results for three scenarios: Normal operations,

[1] References [4], [5], and a companion paper to this one [20] are excellent studies of the linearized DC model accuracy but they only consider normal operating conditions.

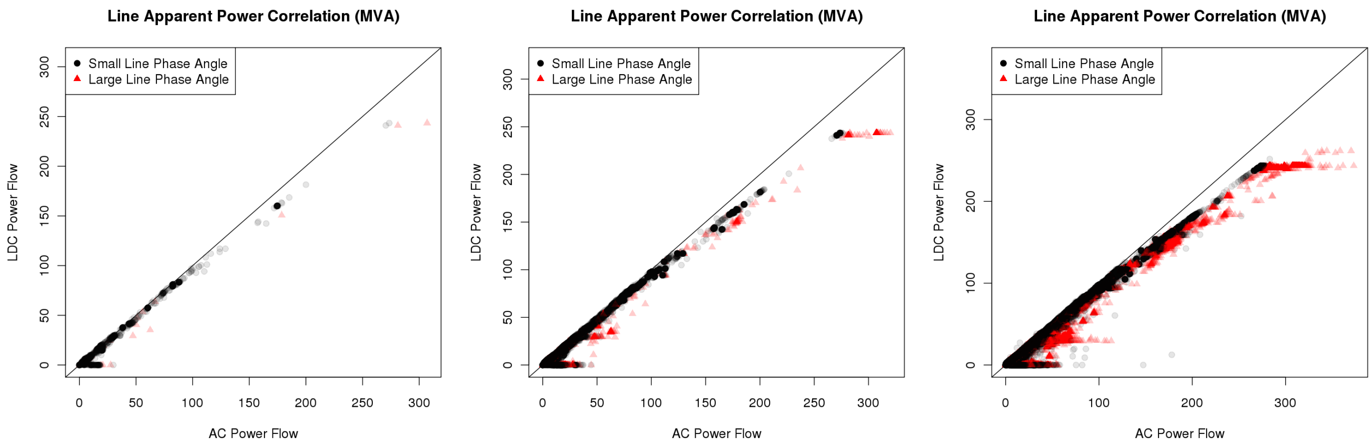

Fig. 1. Accuracy of Apparent Power in N-1(left), N-2(center), N-3(right) Contingencies using the Linearized DC Model.

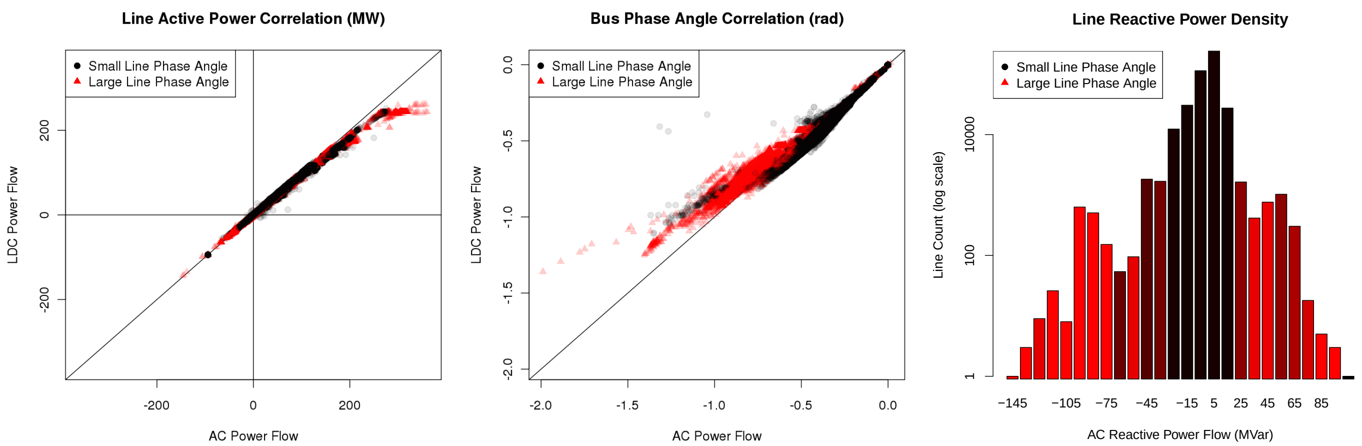

Fig. 2. Accuracy Details of the N-3 Damage Contingencies using the Linearized DC Model.

TABLE I
DAMAGE TO LINES CONNECTING BUS 1 IN THE IEEE30 SYSTEM.

Line	AC Model			LDC Model	
	$\theta^\circ_n - \theta^\circ_m$	p_{nm}	q_{nm}	$\theta^\circ_n - \theta^\circ_m$	p_{nm}
Normal Operation					
$1-2$	0.09338	173.2	-21.09	0.1023	160.1
$1-3$	0.1315	87.74	4.566	0.1479	83.28
Line 1-3 Damaged					
$1-2$	0.1478	270.4	-40.96	0.1556	243.4
$1-3$	–	–	–	–	–
Line 1-2 Damaged					
$1-2$	–	–	–	–	–
$1-3$	0.4862	304.0	43.00	0.4322	243.4

line 1-3 is damaged, and line 1-2 is damaged. Note that, in normal operations, two thirds of the active power is flowing on the line 1-2, so the contingency on that line is likely to be more interesting. The results indicate that, under normal operating conditions and when line 1-3 is damaged, the active power flows are very similar in both models and the phase angles are small. However, when line 1-2 is damaged, there is a large discrepancy in apparent power between the two models although the line angles are rather similar: The active flow in DC power model is 20% lower than the AC value, the angle on line 1-3 is large and approaches 0.5 radians, and a large reactive flow exists.

In summary, the results show that the linearized DC model becomes increasingly less accurate under significant network disruptions. Many of these disruptions create large line phase angles, which lead to under-estimations of active power and significant reactive power.

V. Constraints on Phase Angle Differences

The linearized DC model (i.e., Model 1) is a system of linear equations that can be solved very efficiently, particularly for sparse matrices which is the case for power systems. However, since the phase angles are not restricted in the model, it potentially violates a fundamental assumption of the derivation, i.e., $\sin(\theta^\circ_n - \theta^\circ_m) \approx \theta^\circ_n - \theta^\circ_m$. Moreover, the AC model guarantees that $-1 \leq \sin(\theta^\circ_n - \theta^\circ_m) \leq 1$ while the approximation of the sine term in the linearized DC model is unconstrained. This is problematic because the linearized DC model can produce solutions which are infeasible for the AC power model. In normal operating conditions, this is not likely. But this is certainly possible when large disruptions occur. Obviously, it is possible to state the constraints $-1 \leq \theta^\circ_n - \theta^\circ_m \leq 1$ and use linear programming to obtain a feasible solution but this does not guarantee an accurate approximation (the error is about 20% at the extremes). Even for resonable phase angles the linearized DC model will under-estimate them, as shown in Figure 3. High model accuracy requires much stronger constraints.

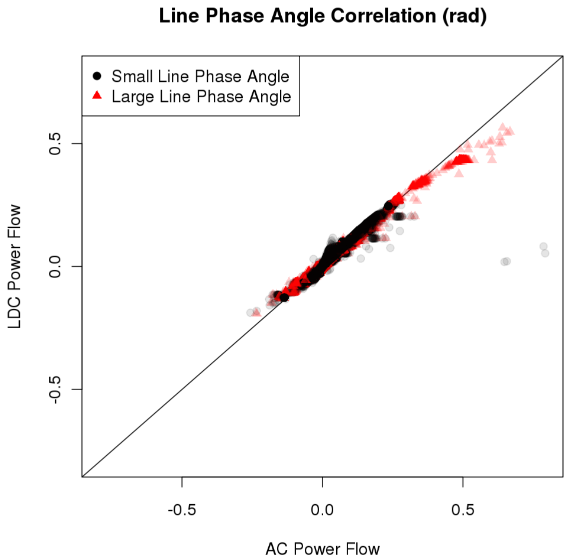

Fig. 3. Accuracy of Line Phase Angles in N-3 Damage Contingencies using the Linearized DC Model.

VI. Angle-Constrained DC Power Flow

This section proposes an Angle-Constrained DC Power Flow (ACDCPF) model that addresses the limitations discussed in the previous section and is particularly appropriate for power restoration. It is based on three key ideas:

1) Impose constraints on the line phase angles to avoid the power underestimations of the linearized DC model;
2) Use load and generation shedding to ensure accuracy of the model;
3) Use an objective function to maximize the served load.

The ACDCPF model is a linear program depicted in Model 2. The model receives as inputs the power network $\mathcal{PN}$, the susceptance values b^y, and the slack bus index s, and the maximum generation G_i and a desired load L_i for each bus i. These last inputs are implicit in the linearized DC model since p_i is always equal to $G_i - L_i$. Its decision variables are the traditional bus phase angles θ_i°, as well as new decision variables g_i and l_i that represent the amount of generation and load at each bus i. The objective function (M2.1) maximizes the served load and hence the ACDCPF model only sheds load to ensure feasibility. Constraint (M2.2) models Kirchhoff's current law and ensures flow conservation for every bus. Constraint (M2.3) enforces the phase angle constraints to remedy the accuracy issues of the linearized DC model. Constraint (M2.4) fixes the angle of the slack bus. It is not necessary for the ACDCPF model in practice and is only introduced here to make meaningful comparisons between the ACDCPF and AC models. Note also that constraint (M2.3) can be posted for every bus, because generator dispatching and load shedding is used to balance load and generation. The ACDCPF model is close to Optimal Power Flow (OPF) models that support flexible generation but typically not load shedding [22], [23], [24]. Our experimental results indicate that the difference in phase angles should be no more than $\pm\pi/12$ (15 degrees) to ensure high accuracy. This tight constraint introduces no more than 1.1% error in active flow on each line

Model 2 Angle-Constrained DC Power Flow (ACDCPF).

Inputs:
$\mathcal{PN} = \langle N, L\rangle$ the power network
b^y susceptance from a $\widetilde{Y}_{bus}$ matrix
s slack bus index
G_i maximum generation at bus i
L_i desired load at bus i

Variables:
$\theta_i^\circ \in (-\infty, \infty)$ - phase angle on bus i (radians)
$g_i \in (0, G_i)$ - generation at bus i
$l_i \in (0, L_i)$ - load at bus i

Maximize:
$$\sum_{n \in N} l_n \tag{M2.1}$$

Subject to:
$$g_n - l_n = \sum_{m \in N}^{n \neq m} b^y(n,m)(\theta_n^\circ - \theta_m^\circ) \quad \forall n \in N \tag{M2.2}$$
$$-\pi/12 \leq \theta_n^\circ - \theta_m^\circ \leq \pi/12 \quad \forall \langle n,m\rangle \in L \tag{M2.3}$$
$$\theta_s^\circ = 0 \tag{M2.4}$$

due to the sine approximation, but clearly the exact choice for this constraint is context-dependent. Observe that line phase angle constraints are very different from bus phase angle constraints (e.g. $-\pi/6 \leq \theta^\circ \leq \pi/6$) that have been employed previously in [1], [13], [11], [19].

It is important to emphasize that, in power restoration, the ACDCPF model is not actually performing load shedding: Rather it decides how much load can be served after a component has been repaired without exacerbating the instability of the network.[2]

A. Case Study: The IEEE30 Network

Section IV showed that the accuracy of the linearized DC model may degrade with significant line damages and that large line phase angles are indicative of such degradation. This section repeats these experiments with the ACDCPF model. Since the ACDCPF model may shed load and generation, it is important to formulate the AC solver appropriately for comparison purposes. In particular, we use the active power values obtained from the ACFCPF model and we obtain the reactive power values by scaling the active power using the power factor, thus shedding active and reactive power proportionally. This shedding approach provides resonable reactive voltage support for a wide range of networks.

Figure 4 revisits the correlation of apparent power under different damage conditions. The figure depicts the N-1 (left), N-2 (center), and N-3 (right) damage contingencies on the IEEE30 system and the results indicate that the ACDCPF model is performing very well. The main source of errors is now due to line losses on heavily loaded lines. As we said earlier, corrections for line losses are well studied and can be integrated in the ACDCPF model (e.g., [2], [5], [20], [21]). There is only one outlier in 10,638 solved N-3 contingencies, which is rather promising. Figure 5 dives in to the worst case

[2]Load shedding may not be acceptable in other settings where the load must be fully served. However, the ACDCPF may see other interesting uses with the advent of demand response.

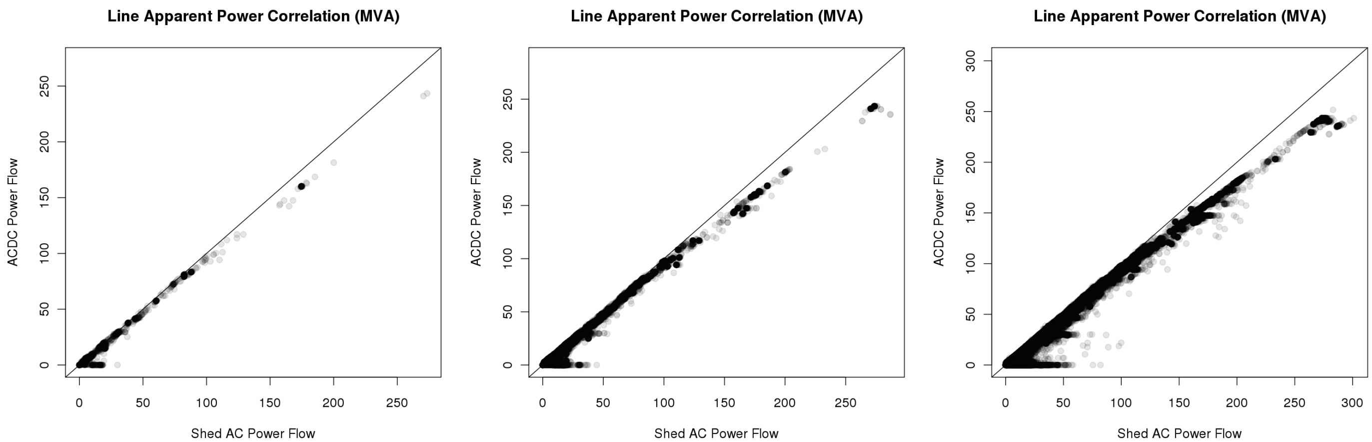

Fig. 4. Accuracy of Apparent Power in N-1(left), N-2(center), N-3(right) Contingencies using the ACDCPF Model.

Fig. 5. Accuracy Details of the N-3 Damage Contingencies using the ACDCPF Model.

TABLE II
AC SOLVABILITY OF IEEE30 DAMAGE CONTINGENCIES

		Linear. DC	ACDCPF Model	
Damage	Contingencies	Solved	Solved	μ(Active Shed)
N-1	41	41	41	0.86%
N-2	820	819	820	2.10%
N-3	10,660	10,602	10,638	3.73%

N-3 damage scenarios and presents results for active power (left), bus phase angles (center), and reactive power (right). The figure now shows a strong correlation for active power and a significant reduction in reactive power. The bus phase angles are also much better correlated.

These results indicate that line phase angle constraints largely remedy the inaccuracies of the linearized DC model and stabilize reactive power flows. This is at the cost of load shedding and it is important to quantify this loss. Table II summarizes the average load shedding for each contingency class. The results indicate that the loss in active power is reasonable: It is about 2% for N-2 contingencies and about 4% for N-3 contingencies on average. The table also reports how many times the ACDCPF solution can be transformed into a feasible AC solution and contrasts that result with the linearized DC model. The results show that the ACDFPF model leads to an AC solution 99.8% of the time and solves 30 more contingencies than the linearized DC model.

B. Case Study: Power Restoration

We now investigate the ACDCPF model for power restoration. In the context of restoration, the damage is so extensive that the full load cannot be served and load shedding always takes place. Figure 6 investigates an illustrative hurricane disaster scenario based on the United States power infrastructure. Additional scenarios are omitted for space considerations. In all of the graphs, two restoration plans are compared, an unconstrained (DCPF RP) and an angle-constrained plan (ACDCPF RP). The x-axis indicates how the power flow solution changes as new restoration actions are executed.

The DC restoration timeline (left) shows the restoration plans proposed by the DCPF and ACDCPF models. The ACDCPF RP restores about 10% less power because it sheds load to satisfy the angle constraint.

The AC restoration timeline (center) is much more illuminating: It depicts the restoration plans obtained by the AC model when initialized with the DC solutions. First, observe that a number of data points are missing: This is caused by the inability of the AC solver to converge, which is not surprising given that there is no hand-tuning of the AC model and solution convergence under unfamiliar conditions is considered ”maddeningly difficult” [2]. The dotted lines compensate for the missing data, by adopting the previous power flow solution when available. The AC restoration timeline thus demonstrates that the DCPF model does not easily admit an AC solution

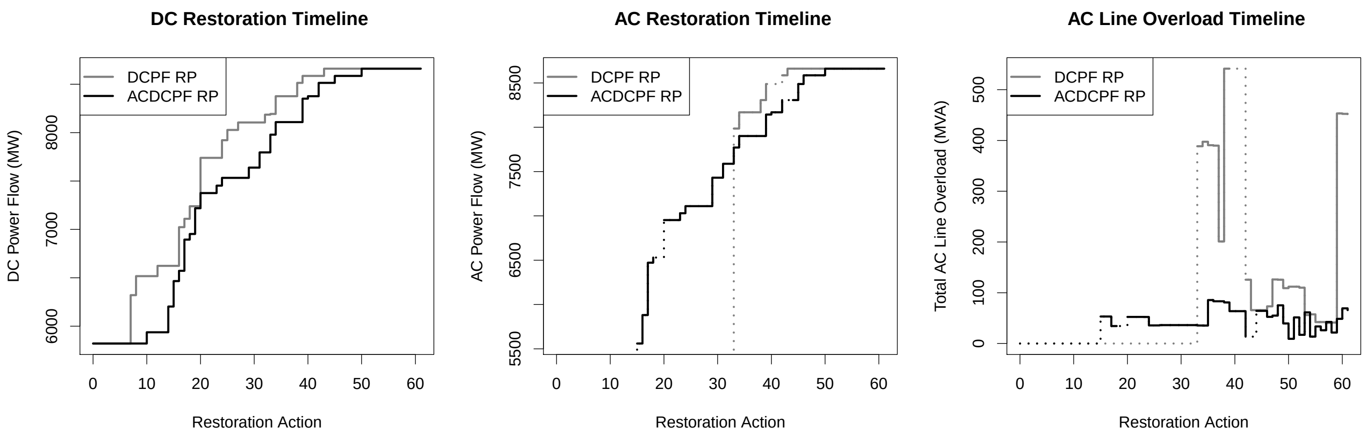

Fig. 6. The Restoration Order Problem With and Without Angle Constraints and the Impacts on AC Power Flow.

under significant disruptions, while the ACDCPF model consistenly produces solutions. In addition, and more importantly, *the AC restoration plan for the ACDCPF model is very close to its DC counterpart, while the DCPF model produces a substantially worse AC restoration plan.*

The right figure shows the total apparent power that exceeds all of the line capacities at various stages of restoration. Indeed, the IEEE30 case study indicates that large phase angles coincide with significant under-approximation of apparent power and that phase angle constraints mitigate these effects. The right figure confirms these results for the ROP by investigating the amount of power that exceeds line capacities. Less than half of the data points for the DCPF model are available, but those data points can have staggering line overloads, which can reach up to 500 MVA. On the contrary, the ACDCPF model significantly reduces line overloads and ensures they are consistently under 100 MVA.

Together, these results indicate that the ACDCPF model makes significant improvements to power restoration. Although its DC restoration plan restores 10% less power over time, its AC restoration plan restores significantly more power and keeps line overloads under control, which is not the case for the DCPF model.

VII. Discussion

A. Impact of the Slack Bus

It is important to emphasize that the ACDCPF does not need a slack bus: *The ACDCPF model is capable of shedding load and generation appropriately to balance the network without resorting to a slack bus or dedicated heuristics*. The slack bus was only introduced in this paper to allow for a natural comparison between the DC and AC models. The slack bus formulation is in fact undesirable as it accumulates line losses along the path to the slack bus, increases the load at the slack bus, and potentially causes larger phase angles and inaccuracies in apparent power (as shown in [5]). So the inaccuracies presented in this paper are likely over-estimated because of the slack bus.

B. Line Losses

It should also be clear that line capacities can easily be added to the ACDCPF model which is a linear program. In fact, our ACDCPF model for power restoration uses capacity constraints on the lines. Line capacities have an interesting connection to line phase angle constraints. In an AC model, a line capacity $c(n,m)$ imposes the constraint

$$\sqrt{p_{nm}^2 + q_{nm}^2} \leq c(n,m)$$

which, in linearized DC models, simplifies to

$$-c(n,m) \leq p_{nm} \leq c(n,m)$$

Expanding the definition of p_{nm} (Equation 1) and dividing by the susceptance $b(n,m)$ gives

$$-\frac{c(n,m)}{b(n,m)} \leq \theta_n^\circ - \theta_m^\circ \leq \frac{c(n,m)}{b(n,m)}.$$

Therefore, line capacity constraints in the ACDCPF model can be viewed as line phase angle constraints. In practice, line capacity constraints may be more or less restrictive than line phase angle constraints and there is no harm in posting both constraints in a linear program. Interestingly, for the standard IEEE benchmarks and those provided with MATPOWER, capacity constraints are significantly less constraining than a $\pm\pi/12$ phase angle constraint. The relationship between line angle constraints and line capacity constraints is fortunate because many existing power system optimization tools support line capacity constraints. A simple preprocessing step can transform line phase angle constraints into equivalent line capacity constraints to ensure model accuracy in existing power system optimization tools.

VIII. Conclusion

This paper studied the applicability of the linearized DC model in the context of disaster recovery and power restoration. It presented experimental results showing that the linearized DC model becomes increasingly less accurate under significant network disruptions. Indeed, the linearized DC model may significantly under-estimate active power in these networks, which also exhibit substantial reactive power. Moreover, the experimental results indicated that these inaccuracies are strongly correlated with large line phase angles.

The paper then presented an ACDCPF model which enhances the linearized DC model with strong line phase angle constraints and the ability to shed load and generation. Experimental results indicate that the ACDCPF model remedies many of the limitations of the linearized DC model, provides highly accurate approximation of active power, and results in solutions with significantly less reactive power. Moreover, the load shedding performed by the ACDCPF model is small and was about 4% on average for N-3 contingencies on the IEEE30 system and 10% for power restoration problems on the United States infrastructure with hurricane scenarios similar to those used at the National Hurricane Center. The ACDCPF model is particularly appropriate for power restoration applications where load shedding is a necessity and the goal is to maximize the load served over time. The experimental results show that the ACDCPF model is much more reliable and produces significant reduction in the blackout area compared to the DCPF model.

Our future research will expand the ACDCPF model, its analysis, and its applications in a variety of directions. From a functionality standpoint, the ACDCPF model should ideally be enhanced to approximate line losses, which are the major source of error at this stage, as well as reactive power in order to obtain more accurate power calculations. From an analysis standpoint, it would be interesting to compare the ACDCPF model with an AC solver using slack bus distribution. Such an AC solver models the real power systems more accurately and provides a better basis for comparison, since our ACDCPF model can shed load and generation at all buses. Finally, from an application standpoint, it would be interesting to study the applicability of our ACDCPF model to other types of applications exhibiting significant topological changes, such as long-term planning and vulnerability analysis.

REFERENCES

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

[2] T. Overbye, X. Cheng, and Y. Sun, "A comparison of the ac and dc power flow models for lmp calculations," in *System Sciences, 2004. Proceedings of the 37th Annual Hawaii International Conference on*, jan. 2004, p. 9 pp.

[3] J. Salmeron, K. Wood, and R. Baldick, "Worst-case interdiction analysis of large-scale electric power grids," *Power Systems, IEEE Transactions on*, vol. 24, no. 1, pp. 96 –104, feb 2009.

[4] K. Purchala, L. Meeus, D. Van Dommelen, and R. Belmans, "Usefulness of DC power flow for active power flow analysis," *Power Engineering Society General Meeting, 2005. IEEE*, pp. 454–459, 2005.

[5] B. Stott, J. Jardim, and O. Alsac, "DC Power Flow Revisited," *Power Systems, IEEE Transactions on*, vol. 24, no. 3, pp. 1290–1300, 2009.

[6] O. Ceylan, A. Ozdemir, and H. Dag, "Branch outage solution using particle swarm optimization," in *Power Engineering Conference, 2008. AUPEC '08. Australasian Universities*, dec. 2008, pp. 1 –5.

[7] R. Baldick, *Applied Optimization: Formulation and Algorithms for Engineering Systems*. Cambridge University Press, 2009.

[8] M. Crow, *Computational Methods for Electric Power Systems, Second Edition (Electric Power Engineering Series)*. CRC Press, 2009.

[9] J. Grainger and W. S. Jr., *Power System Analysis*. McGraw-Hill Science/Engineering/Math, 1994.

[10] Y. Tamura, H. Mori, and S. Iwamoto, "Relationship between voltage instability and multiple load flow solutions in electric power systems," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-102, no. 5, pp. 1115 –1125, may 1983.

[11] E. Fisher, R. O'Neill, and M. Ferris, "Optimal transmission switching," *Power Systems, IEEE Transactions on*, vol. 23, no. 3, pp. 1346 –1355, aug. 2008.

[12] D. Bienstock and S. Mattia, "Using mixed-integer programming to solve power grid blackout problems," *Discrete Optimization*, vol. 4, no. 1, pp. 115 – 141, 2007.

[13] C. Coffrin, P. Van Hentenryck, and R. Bent, "Strategic stockpiling of power system supplies for disaster recovery." *Proceedings of the 2011 IEEE Power & Energy Society General Meetings (PES 2011), Detroit, MI*, july 2011.

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

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

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

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

[18] C. Murillo-Sánchez and D. Gan, "Matpower," *Ithaca*, 1997.

[19] K. Hedman, R. O'Neill, E. Fisher, and S. Oren, "Optimal transmission switching with contingency analysis," *Power Systems, IEEE Transactions on*, vol. 24, no. 3, pp. 1577 –1586, aug. 2009.

[20] C. Coffrin, P. Van Hentenryck, and R. Bent, "Approximating Line Losses and Apparent Power in AC Power Flow Linearizations," *Submitted to, Proceedings of the 2012 IEEE Power & Energy Society General Meetings (PES 2012), San Diego, CA*, july 2012.

[21] C. A. N. A Borghetti, M Paolone, "A Mixed Integer Linear Programming Approach to the Optimal Configuration of Electrical Distribution Networks with Embedded Generators," *Proceedings of the 17th Power Systems Computation Conference (PSCC'11), Stockholm, Sweden*, 2011.

[22] D. Kirschen and H. Van Meeteren, "Mw/voltage control in a linear programming based optimal power flow," *Power Systems, IEEE Transactions on*, vol. 3, no. 2, pp. 481 –489, may 1988.

[23] O. Alsac, J. Bright, M. Prais, and B. Stott, "Further developments in lp-based optimal power flow," *Power Systems, IEEE Transactions on*, vol. 5, no. 3, pp. 697 –711, aug 1990.

[24] H. Dommel and W. Tinney, "Optimal power flow solutions," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-87, no. 10, pp. 1866 –1876, oct. 1968.

Carleton Coffrin received a B.Sc. in Computer Science and a B.F.A. in Theatrical Design from the University of Connecticut, Storrs, CT and a M.S. from Brown University, Providence, RI. He is currently a Ph.D. candidate in the department of computer science at Brown University.

Pascal Van Hentenryck is a Professor of Computer Science at Brown University. He is a fellow of the Association for the Advancement of Artificial Intelligence, the recipient of the 2002 ICS INFORMS award, the 2006 ACP Award, honorary degrees from the University of Louvain and the University of Nantes, and the Philip J. Bray award for teaching excellence. He is the author of five MIT Press books. Most of this research in optimization software systems has been commercialized and is widely used in academia and industry.

Russell Bent received his PhD in Computer Science from Brown University in 2005 and is currently a staff scientist at LANL in the Energy and Infrastructure Analysis Group. His publications include deterministic optimization, optimization under uncertainty, infrastructure modeling and simulation, constraint programming, algorithms, and simulation. Russell has published one book and over thirty articles in peer-reviewed journals and conferences in artificial intelligence and operations research. A full list of his publications can be found at *http://public.lanl.gov/rbent/*.