

PowerModelsRestoration.jl: An open-source framework for exploring power network restoration algorithms

Noah Rhodes^a, David M. Fobes^b, Carleton Coffrin^{b,*}, Line Roald^a

^a University of Wisconsin-Madison, Madison, WI, USA

^b Los Alamos National Laboratory, Los Alamos, NM, USA

ARTICLE INFO

Keywords:

Power system restoration
N-k
Nonlinear optimization
Convex optimization
AC power flow
Julia language
Open-source

ABSTRACT

With the escalating frequency of extreme grid disturbances, such as natural disasters, comes an increasing need for efficient recovery plans. Algorithms for optimal power restoration play an important role in developing such plans, but also give rise to challenging mixed-integer nonlinear optimization problems, where tractable solution methods are not yet available. To assist in research on such solution methods, this work proposes PowerModelsRestoration, a flexible, open-source software framework for rapidly designing and testing power restoration algorithms. PowerModelsRestoration constructs a mathematical modeling layer for formalizing core restoration tasks that can be combined to develop complex workflows and high performance heuristics. The efficacy of the proposed framework is demonstrated by proof-of-concept studies on three established cases from the literature, focusing on single-phase positive sequence network models. The results demonstrate that PowerModelsRestoration reproduces the established literature, and for the first time provide an analysis of restoration with nonlinear power flow models, which have not been previously considered.

1. Introduction

As the threat of exogenous grid disturbances, *e.g.*, natural disasters or sophisticated targeted attacks, continues to intensify, so does the importance of expedient power network restoration. There is thus an increasing need for decision support tools that can assist network operators in identifying optimal restoration plans, and ultimately provide autonomous self-healing capabilities to the grid. Unfortunately, high-fidelity modeling of power system restoration has proven notoriously difficult due to the challenge of finding an AC-feasible operating point in networks where hundreds to thousands of components have been damaged [1,2]. However, recent advances in convex relaxations of the AC power flow equations have shown promise for single-time-point N-k analysis [3].

While traditional work on power system restoration has focused on system reenergization following a blackout where the vast majority of components are undamaged [4–6], the work in this paper considers the longer-term problem of restoring power supply following extreme physical impacts, such as hurricanes or earth quakes, where a large number of components must be repaired before they can be re-energized. Specifically, it focuses on component restoration ordering, *i.e.*, selecting the sequence in which components should be prioritized for restoration, in order to minimize energy not served over time. This

prioritization task results in a sequential network design problem, which, in practice, can be remarkably challenging to solve [7], due in part to the combination of discrete and continuous optimization variables as well as nonlinear, non-convex constraints. Consequently, significant network approximations and/or heuristic methods are often required to make the problem tractable. Optimizing the restoration order can provide a significant reduction in energy not served relative to what is achieved through applying simple rules or heuristics [8]. However, previous work [2,9] have demonstrated that popular approximations, *e.g.*, DC power flow, fail to capture important aspects of the problem, leading to suboptimal restoration ordering, infeasible intermediate solutions, and higher-than-necessary energy not served.

Designing and validating the effectiveness of different approaches to this problem requires significant research, and is an essential step towards the aspiration of a resilient, self-healing grid. To support this design challenge, the core contribution of this work is a novel software framework, PowerModelsRestoration, which enables rapid exploration of power network restoration algorithms. By building on the PowerModels framework [10], PowerModelsRestoration is able to consider a broad range of power flow formulations, spanning the full AC equations [11], convex relaxations [12,13], and active-power-only approximations [14]. This restoration framework includes exact restoration algorithms, modeled as mixed-integer nonlinear programs

* Corresponding author.

E-mail addresses: nrhodes@wisc.edu (N. Rhodes), dfobes@lanl.gov (D.M. Fobes), cjc@lanl.gov (C. Coffrin), roald@wisc.edu (L. Roald).

<https://doi.org/10.1016/j.epsr.2020.106736>

Received 4 October 2019; Received in revised form 20 April 2020; Accepted 2 August 2020

Available online 26 August 2020

0378-7796/ © 2020 Published by Elsevier B.V.

(MINLPs), and in the future will contain heuristic restoration algorithms, such as largest capability first. Some of the notable features of PowerModelsRestoration include: (1) support for AC-based restoration, (2) restoration plan quality guarantees provided by convex relaxations, (3) incorporation of storage devices in power system restoration, and (4) tools for simulating restoration plans with the AC power flow equations. With these features, PowerModelsRestoration aspires to be a valuable tool for rapidly exploring the wide variety of possible restoration algorithms, and providing a baseline implementation of established restoration algorithms for resilience analysis.

Utilizing the proposed software framework, this work develops case studies that highlight the benefits and drawbacks of established power flow formulations. On the one hand, the results demonstrate the flexibility and value of the software framework, while on the other hand provide new insights into restoration nonlinear formulations, not previously considered to our knowledge. An unexpected contribution of the paper is the insight that convex relaxations of the power flow equations may provide a valuable tool for balancing accuracy and performance when developing power restoration plans.

This work begins with a brief overview of the power system restoration context that motivates this work in Section 2, followed by a review of mathematical programming and the PowerModels framework in Section 3. The mathematical modeling layer of PowerModelsRestoration is presented in Section 3.3. Section 4 illustrates how the modeling layer can be combined into more complex analysis workflows. Section 5 develops restoration studies to validate the framework and Section 6 concludes the paper.

2. Power restoration background and modelling

This section provides an overview of the main modeling concepts that currently form the basis of PowerModelsRestoration.

2.1. Restoration of a physically damaged grid

As noted above, traditionally in power system restoration work, it is assumed that the majority of components are still physically intact. For transmission grid restoration, the main objective of traditional methods is therefore to create a sequence in which components can be re-energized, considering, e.g., the black-start capabilities of power plants and the dynamics of cold-load pick-up and power plant re-synchronization [15]. Distribution system restoration has focused on a problem called Power Supply Restoration (PSR) [16–18], which is an operational task that considers how to reconfigure a distribution system's topology for the purposes of fault isolation and power resupply to a maximum possible amount of loads.

In contrast, we consider a different time scale, where the initial restoration and re-synchronization of available generators has already taken place, but a large number of system components are still damaged, and require physical repairs. This problem has been the topic of transmission system restoration research for many years [4] and forms the basis for the software implementation discussed in this paper. Related work for distribution grids includes [19], which utilizes a connectivity-based model of the power network.

2.2. Modelling assumptions

In our current formulation, we assume: (1) a single-phase equivalent network representation; (2) full awareness of the system's state, including energized and damaged components; (3) remote controllability of all components, including, e.g., generator outputs and breaker set-points; and (4) the ability to partially shed loads at buses. In the context of a transmission system, where instruments such as measurement

units, remote control devices, and circuit breakers are typically available throughout the system, these assumptions are reasonable. Additionally, transmission loads usually represent aggregations of many smaller loads, which can be switched via substation re-configuration procedures, enabling gradual load shedding.

For distribution grids, these assumptions represent a simplification, which are often considered non-realistic. In distribution grids the number of measurement devices is typically limited and information about damaged items requires manual surveying. There are typically fewer switches and breakers, so branches cannot be re-energized before several of them have been repaired, and loads are often discrete. Finally, distribution systems are usually unbalanced (especially so during restoration), so a three-phase representation would be more appropriate. However, given the limited amount of previous work on distribution grids, which did not directly account for power flow or voltage magnitude constraints (but has considered limited number of branches [18]), we believe that modeling framework demonstrated here can also provide interesting insights into distribution restoration. Further details on how we plan to extend this framework to more realistic distribution grid models is provided in the future work section.

2.3. Problem formulations

Based on the above assumptions, we present two problems essential to both transmission and distribution grid restoration.

Minimum Restoration Set Problem (MRSP) In the interest of fault resiliency and routine maintenance, power systems are engineered with many redundant components. Consequently, only a subset of the damaged components may need to be restored to meet the needs of all of the loads in the system. The goal of the Minimum Restoration Set Problem (MRSP) is to identify the smallest set of components that must be repaired to support the entire system load. From an optimization perspective, the MRSP is a network design task, similar to Transmission Network Expansion Planning [20], but tailored to the power restoration context.

This problem includes discrete variables for each damaged component indicating whether the component will be repaired or remain out-of-service. To model whether a choice of repairs is feasible, the problem includes a set of equality constraints representing the power flow throughout the grid, as well as constraints on generation capacity, transmission line ratings, and voltage magnitudes for all in-service components. The problem focuses only on a single time point (e.g., near-term peak demand). A more detailed discussion on the MRSP problem formulation can be found in [8].

Restoration Ordering Problem (ROP) The Restoration Ordering Problem (ROP) represents the most fundamental power system restoration optimization task. This problem most closely reflects the goal of network operators during power restoration, i.e., to minimize the amount of load shed, given a limited amount of resources. The objective of the ROP is to identify the sequence in which the components should be restored, while minimizing unserved load over time, typically referred to as *energy not served (ENS)*. From an optimization perspective, the ROP is a generalization of the MRSP, a so-called sequential network design task [8,21].

The ROP problem is divided into finite time steps, and in each time step a limited number of components can be repaired (representing resource constraints on, e.g., the number of repair teams). After a component has been repaired, it can either immediately be put into service, or wait for other components to be restored. In our model, we only include variables to indicate whether a component is out-of-service (damaged, or repaired but not yet energized) or in-service (not damaged, or repaired and energized). After a component has been energized, it must remain in service for the remainder of the restoration.

Similar to the MRSP, at each step the ROP includes a model of both the power flow and the generation and transmission constraints. Our current implementation also assumes a constant load profile, representing the daily peak load (assuming restoration takes days or weeks to be completed).

2.4. Computational challenges

The MRSP and ROP tasks include both continuous and discrete decision variables. Such problems present significant computational challenges, particularly when combined with detailed representations of the power flow, such as the nonlinear, non-convex AC power flow model. Particularly the ROP problem, which includes temporal constraints linking the restoration sequence, is very challenging to solve in practice.

Current optimization solvers are inadequate for reliably solving the MINLP presented by AC power system restoration on a large scale [3]. A common approach is to optimize the restoration problems with approximate linear power flow models (e.g., DC Power Flow), which can utilize outstanding commercial optimization tools, such as Gurobi [22] and CPLEX [23]. However, this approach requires the conversion of restoration plans into AC-feasible plans, proven to be significantly challenging [2,24]. Fortunately, recent works have proposed new promising approximations [25] and relaxations [3] to address these shortcomings. PowerModelsRestoration brings these approximations and relaxations together within a common framework, convenient for exploring multiple formulations in terms of accuracy and computational efficiency.

3. PowerModelsRestoration

Although many approximations or relaxations of the power flow equations [3] have been proposed, it is unclear which are suitable for power system restoration [2,9]. To address this, it is necessary to (1) develop several implementations of power system restoration models with differing levels of modeling accuracy for performance analysis and model validation, and (2) devise workflows that decompose complex power restoration tasks into multiple subproblems that can be solved effectively at scale [7,8]. A central goal of this work is to support researchers in addressing these challenges.

3.1. PowerModels

The PowerModels framework [10], which was designed to streamline model development and assist in exploring the accuracy and computational efficiency of different power flow formulations, is a core inspiration for this work. PowerModels is implemented in the mathematical programming framework JuMP [26], which has emerged as a commercial-grade tool for optimization modeling in Julia [27], a high-performance programming language for numerical computing. PowerModels [10] was created as a modeling layer specializing in mathematical programs for power systems focusing on enabling easy switching between a variety of power flow formulations, which has been particularly useful for research on efficient formulations of AC Optimal Power Flow (OPF) [13].

PowerModels allows the development of formulation-agnostic power system optimization models using power-network-aware abstractions. These abstractions are transformed into concrete mathematical models upon specification of a formulation. An extensive list of power flow formulations is provided with the PowerModels documentation, but below we illustrate a simple example of how these abstractions work.

Example of PowerModels Abstractions A power network optimization task, e.g., OPF, involves several expressions that include the complex voltage product, $V_i V_j^*$. Depending on the formulation specified by the user, PowerModels replaces these generic expressions with a specific

real number implementation. For example, specifying the *ACPowerModel* formulation (the AC power flow formulation in polar coordinates) results in the following mapping:

$$V_i V_j^* \Rightarrow |V_i| |V_j| \cos(\theta_i - \theta_j) + i |V_i| |V_j| \sin(\theta_i - \theta_j) \quad (1)$$

If the *DCPowerModel* formulation (the traditional DC power flow approximation) is specified, the mapping is,

$$V_i V_j^* \Rightarrow 1.0 + i(\theta_i - \theta_j) \quad (2)$$

while the *SOCWRPowerModel* relaxation from [12] results in,

$$V_i V_j^* \Rightarrow W_{ij}^R + i W_{ij}^I \quad (3a)$$

$$(W_{ij}^R)^2 + (W_{ij}^I)^2 \leq W_{ii}^R W_{jj}^R \quad (3b)$$

For more general expressions, this mapping may be complex and can add many auxiliary variables and constraints.

3.2. Extension to PowerModelsRestoration

This work introduces PowerModelsRestoration, a framework for studying power system restoration optimization. As an extension of PowerModels, PowerModelsRestoration can support a variety of power flow formulations, including AC Power Flow (*ACPowerModel*), DC Power Flow (*DCPowerModel*), and the SOC convex relaxation (*SOCWRPowerModel*) [12]. Furthermore, its flexible design similar to PowerModels enables and encourages user-driven extension. For example, while here we present a single-phase equivalent reformulation of the network, development to extend the framework to support multi-phase distribution network restoration via *PowerModelsDistribution* is underway.

3.3. Restoration optimization specifications

In the following, we provide working examples of the two canonical restoration problems introduced above, the MRSP and ROP. Detailed derivations of the mathematical models for each problem are available in [8,9], and a more complete description of all constraints and variables can be found in software documentation for PowerModels¹ and PowerModelsRestoration². The goal of our presentation here is to demonstrate how the PowerModelsRestoration modeling layer is leveraged to implement these mathematical programs for power system restoration with a variety of power flow formulations. We also include functional code examples, which serve as concise illustrations of how the problems are constructed in PowerModelsRestoration v0.5. Note that while PowerModelsRestoration includes variables and constraints for energy storage devices, we have omitted these components in the examples below for brevity and clarity.

Minimum Restoration Set Problem (MRSP) We first discuss the implementation of the MRSP problem. Code Block 1 shows the specification of the MRSP introduced informally in Section 2 as implemented in PowerModelsRestoration. First, optimization variables are added to the model, including all variables typically present in an OPF problem, which are continuous, as well as a discrete 0/1 variable representing the status of a component, indicated by the suffix `_indicator`. Next, constraints linking the optimization variables are added; cross referencing Code Block 1 with the OPF formulation in PowerModels [10] highlights the model similarities, e.g., the re-use of the function `constraint_power_balance`, which represents the nodal power balance constraints. The key differences between the MRSP and standard OPF formulations are the `_damage` constraints, which consider whether a component is active based on the status of its indicator variable. Additionally, the constraints of the form

¹ <https://lanl-ansi.github.io/PowerModels.jl/stable>

² <https://lanl-ansi.github.io/PowerModelsRestoration.jl/stable>

`constraint_component_damage` implement the component dependency requirement that a bus must be active if anything connecting to it is active. Finally, an objective function is added to minimize the number of restored components.

Code Block 1 Abstract Minimum Restoration Set Model

```

1 function build_mrsp(pm::AbstractPowerModel)
2 # the mathematical program's optimization variables
3 variable_bus_damage_indicator(pm)
4 variable_bus_voltage_damage(pm)
5
6 variable_branch_damage_indicator(pm)
7 variable_branch_power(pm)
8
9 variable_gen_damage_indicator(pm)
10 variable_gen_power_damage(pm)
11
12 # component by component power system constraints
13 constraint_model_voltage_on_off(pm)
14
15 for i in ids(pm, :ref_buses)
16 constraint_theta_ref(pm, i)
17 end
18
19 for i in ids(pm, :bus)
20 constraint_bus_damage_soft(pm, i)
21 constraint_power_balance(pm, i)
22 end
23
24 for i in ids(pm, :gen)
25 constraint_gen_damage(pm, i)
26 end
27
28 for i in ids(pm, :branch)
29 constraint_branch_damage(pm, i)
30 constraint_ohms_yt_from_damage(pm, i)
31 constraint_ohms_yt_to_damage(pm, i)
32
33 constraint_voltage_angle_difference_damage(pm, i)
34
35 constraint_thermal_limit_from_damage(pm, i)
36 constraint_thermal_limit_to_damage(pm, i)
37 end
38
39 # a minimum restoration count objective
40 objective_min_restoration(pm)
41 end

```

Restoration Ordering Problem (ROP) We next discuss the implementation of the multi-period ROP. Code Block 2 shows a specification of the ROP introduced in Section 2. As previously mentioned, the ROP is a multi-time point generalization of the MRSP. The sequential nature of the ROP leverages PowerModels' *multinetwork* feature, which enables multiple networks with linking constraints to be optimized simultaneously. In this specification, N copies of the network are specified in the data model, each representing a single time-point network restoration state, which are linked together with inter-temporal constraints. It is assumed that that $n = 1$ represents the initial state, and all damaged components can be restored by $n = N$. The number of restoration periods $N - 1$ is specified by the user, and the per period limit on the number of restored items (constant for all periods) is automatically chosen such that all items can be restored by the final period.

The first for-loop of Code Block 2 constructs the independent restoration models for each time period, reusing much of the modeling seen in Code Block 1. The notable differences are (1) new variables for power demands (loads) and bus shunts, which are required for partial power restoration in each period, and (2) a cardinality constraint, which limits the number of restoration actions in each period.

Perhaps the most crucial constraints in the ROP specification are the inter-temporal constraints appearing in the second outer for-loop of Code Block 2. The `energized` constraints ensure that once a component is activated it remains active throughout the remainder of the restoration, while the `increasing` constraints ensure that once a load is served, it remains so for the remainder of the restoration. The `restore_all_items` constraint ensures that all initially damaged components are repaired and energized by the restoration completion. Finally, the objective function aims to maximize the load delivered throughout the restoration process.

Code Block 2 Abstract Restoration Ordering Model

```

1 function build_rop(pm::AbstractPowerModel)
2 # add multiple copies of the restoration network
3 for (n, network) in nws(pm)
4 variable_bus_damage_indicator(pm, nw=n)
5 variable_bus_voltage_damage(pm, nw=n)
6
7 variable_branch_damage_indicator(pm, nw=n)
8 variable_branch_power(pm, nw=n)
9
10 variable_gen_damage_indicator(pm, nw=n)
11 variable_gen_power_damage(pm, nw=n)
12
13 variable_load_power_factor(pm, nw=n, relax=true)
14 variable_shunt_admittance_factor(pm, nw=n,
15 relax=true)
16
17 constraint_restoration_cardinality_ub(pm, nw=n)
18
19 constraint_model_voltage_damage(pm, nw=n)
20
21 for i in ids(pm, :ref_buses, nw=n)
22 constraint_theta_ref(pm, i, nw=n)
23 end
24
25 for i in ids(pm, :bus, nw=n)
26 constraint_bus_damage_soft(pm, i, nw=n)
27 constraint_power_balance_shed(pm, i, nw=n)
28 end
29
30 for i in ids(pm, :gen, nw=n)
31 constraint_gen_damage(pm, i, nw=n)
32 end
33
34 for i in ids(pm, :load, nw=n)
35 constraint_load_damage(pm, i, nw=n)
36 end
37
38 for i in ids(pm, :shunt, nw=n)
39 constraint_shunt_damage(pm, i, nw=n)
40 end
41
42 for i in ids(pm, :branch, nw=n)
43 constraint_branch_damage(pm, i, nw=n)
44 constraint_ohms_yt_from_damage(pm, i, nw=n)
45 constraint_ohms_yt_to_damage(pm, i, nw=n)
46
47 constraint_voltage_angle_difference_damage(pm,
48 i, nw=n)
49
50 constraint_thermal_limit_from_damage(pm, i,
51 nw=n)
52 constraint_thermal_limit_to_damage(pm, i, nw=n)
53 end
54 end
55
56 # add inter-temporal constraints across the networks
57 network_ids = sort(collect(nw_ids(pm)))
58 n_1 = network_ids[1]
59 for n_2 in network_ids[2:end]
60 for i in ids(pm, :gen, nw=n_2)
61 constraint_gen_energized(pm, i, n_1, n_2)
62 end
63 for i in ids(pm, :bus, nw=n_2)
64 constraint_bus_energized(pm, i, n_1, n_2)
65 end
66 for i in ids(pm, :branch, nw=n_2)
67 constraint_branch_energized(pm, i, n_1, n_2)
68 end
69 for i in ids(pm, :load, nw=n_2)
70 constraint_load_increasing(pm, i, n_1, n_2)
71 end
72 n_1 = n_2
73 end
74
75 n_final = last(network_ids)
76 constraint_restore_all_items(pm, n_final)
77
78 objective_max_load_delivered(pm)
79 end

```

4. Using PowerModelsRestoration

In this section we demonstrate how the mathematical programs for power system restoration highlighted in the previous section can be leveraged in simple Julia scripts to perform more complex power restoration analysis.

Choosing a solver PowerModelsRestoration constructs a mathematical program (i.e., a JuMP model), which can be solved using a general purpose optimization solver. Many solvers are available in JuMP, including Ipopt for continuous nonlinear programs (NLP) [28], Cbc for mixed-integer linear programs (MIP) [29], and Juniper for mixed-integer nonlinear programs (MINLP) [30]. Better performance may be gained by using commercial solvers, e.g., Gurobi or KNITRO.

Restoration Optimization and Simulation A core workflow for power systems restoration is to compute a restoration plan using a relaxed or approximated power flow formulation and validate the resulting plan with a higher fidelity power flow representation [2,9]. Different from the formulation in [2,9], which was based on running an AC power flow with the provided restoration sequence and generation setpoints, we implement a *multi-period AC OPF* that only uses the ordering of the components to fix the binary variables values, and allow the redispatch to adjust the value of all other variables, such as choosing new generation setpoints. Code Block 3 provides an example of this workflow implemented in PowerModelsRestoration. First, the DC power flow approximation is used to compute a restoration plan, and the resulting plan is then simulated via the inter-temporal AC maximal load delivery optimization (i.e., ACPPowerModel).

Code Block 3 AC Redispatch of a DC Restoration Plan

```
using PowerModels, PowerModelsRestoration
using Ipopt, Cbc

# load the network data with component damage
case = PowerModels.parse_file("case5_strg_damaged.m")
restore_case = replicate_restoration_network(case, count=3)

# optimize a restoration plan using the DC model
restore_result = run_rop(restore_case, DCPowerModel,
    Cbc.Optimizer)

# update the data with the restoration plan
clean_status!(restore_result["solution"])
update_status!(restore_case, restore_result["solution"])

# simulate the restoration plan with the AC model
result = run_restoration_redispatch(restore_case,
    ACPPowerModel, Ipopt.Optimizer)
```

Code Block 4 Building a Three-Step AC Restoration Plan

```
using PowerModels, PowerModelsRestoration, JuMP
using Ipopt, Juniper
ipopt = optimizer_with_attributes(Ipopt.Optimizer,
    "print_level"=>0)
solver = optimizer_with_attributes(Juniper.Optimizer,
    "nl_solver"=>ipopt)

# load the network data with component damage
case = PowerModels.parse_file("case5_strg_damaged.m")

# solve the minimum restoration set problem
result = run_mrsp(case, ACPPowerModel, solver)

# update the case data with the minimum restoration set
clean_status!(result["solution"])
update_status!(case, result["solution"])

# compute a three-step restoration plan
restore_case = replicate_restoration_network(case, count=3)
restore_result = run_rop(restore_case, ACPPowerModel, solver)
```

Accelerating Restoration Optimization Another core workflow for improving the solve time of a large ROP instance is to first solve the MRSP and then solve the ROP including only the subset of damaged components identified by the MRSP solution [8,9]. Code Block 4 presents an example of this workflow. Communication between the MRSP and ROP models occurs in the network data specification, where the component *status* field is leveraged to indicate components that should be ignored in the ROP. This example uses the AC power flow variant and specifies 3 time periods for the ROP. Other power flow models can be used by changing the power flow specification to, e.g., SOCWRPowerModel.




Fig. 1. A network diagram of a five-bus example illustrating core restoration decision problems. The generators on bus 1 have a small capacity, representing distributed energy resources. Note that bus 5 has energy storage.

5. Case study

The core features of PowerModelsRestoration are demonstrated by the following case study, using a small example for illustrative purposes, as well as two typical test cases representing both a transmission grid and a single-phase representation of a distribution grid. Beyond demonstrating the software, these simulations provide new insights into the impact of the different formulations and workflows. Throughout this section the following solvers were used: Ipopt v3.12 [31] for the NLP version of AC formulation, Gurobi v8.1 for the MIP and MISOCP models given by the DC and SOC formulations, and Juniper v0.6 [32] using Ipopt and Gurobi [22] as sub-solvers for the MINLP version of AC model.

5.1. Comparison of different power flow formulations

The first case study investigates the impact of power flow formulations on computational efficiency and quality of the solutions. For this investigation, we use a small 5-bus network [33], illustrated in Fig. 1. Using a small system allows for benchmarking simpler formulations against more detailed, but also more computationally intensive formulations. For the purposes of this study, all components in the network were damaged (except buses) and the workflow illustrated by code block 3 was utilized. We first solve the ROP problem using a DC, AC, and SOC power flow formulation, respectively. Each solution provides us with an optimal restoration sequence (from the optimal value of the decision variables), as well as an *estimated ENS* (given by the objective function value). For each solution, we then solve a multi-step OPF using the full AC formulation to redispatch generators to obtain the *true ENS*, except for the solution obtained with the AC ROP. For the AC ROP, the optimization already accounts for the full AC model and thus the estimated ENS would be the same as the true ENS.

We first compare the performance of different power flow formulations in terms of (i) the solve speed of the ROP problem, (ii) the estimated ENS, as obtained directly from the ROP, and (iii) the true ENS as computed by the AC redispatch (shown in Table 1). We observe that the DC ROP is by far the fastest, requiring ≈ 0.1 s to reach an optimal solution, and the SOC ROP is also relatively quick, (2.11s). By comparison, the AC ROP requires more than 15 minutes even for the 5-bus case, and can only solve to local optimality which results in a solution with a comparatively higher ENS.

In terms of solution quality, we first observe that using the AC power flow formulation provides inferior results, despite a higher fidelity model. This is due to higher problem complexity; the problem is non-convex even after the binary variables are fixed, and only a locally

Table 1
Comparison of power flow formulations (5-bus) .

Power Flow Formulation	Estimated ENS [MWh]	True ENS [MWh]	Solve Time [s]
DC	2900	2902.5	0.12
SOC	2902.4	2902.5	2.11
AC	3953.9	-	914.72




Fig. 2. (a) True Energy Not Served (ENS) for the 5-bus case with the AC, SOC and DC models. ENS is obtained from restoration redispatch using the full AC formulation. (b) True ENS for the ROP and MRSP + ROP problems using the SOC formulation in the fully damaged 5-bus case. (c) Estimated and true ENS per restoration period obtained from the DC ROP solution and corresponding AC redispatch.

optimal solution can be found. Furthermore, we observe that the DC ROP returns the lowest estimated ENS. However, by comparing the estimated and true ENS, we note that the DC ROP problem is overly optimistic, and has a higher ENS (i.e. serves less load) than predicted. A similar observation can be made for the SOC problem, although the gap between the estimated and true ENS is notably smaller. This indicates that the SOC ROP provides more accurate results than the DC ROP solution, with relatively small sacrifices in solve time. However, in this particular case, both the SOC and DC find a similar restoration sequence and are able to supply the same amount of load in the simulation. To provide further details about the solutions, the ENS for each time step of the AC solutions is presented in Fig. 2a. This plot highlights that the AC solution has a higher ENS in some of the time steps, and that the DC and SOC restoration sequences have very similar ENS in each period.

5.2. MRSP as a preprocessing step

The second case study considers the performance gains of applying MRSP as a preprocessing step, as well as the impact on solution quality. For this study, we use the same fully damaged 5-bus test case as above and the SOC power flow formulation. We compare two versions of the ROP problem, the original (as in the above section), and one where we apply MRSP as a preprocessing step to reduce the number of considered elements (i.e., the MRSP + ROP problem). This workflow is similar to the example shown in code block 4, and the results are shown in Table 2.

The MRSP + ROP problem returns a solution in 0.103 seconds (including the solve time for both the MRSP and the ROP problems), which is $\sim 20 \times$ faster than directly solving the ROP problem. This speed-up can be achieved because the MRSP problem reduces the number of components considered for repair from 11 to 6, meaning fewer restoration periods and fewer binary variables per time step in the ROP problem. We expect this speed enhancement to be more significant for larger systems containing more damaged components.

In terms of solution quality, the MRSP + ROP produces a result with an ENS of 3081.4 MWh, notably higher than the restoration sequence from the full ROP (2902.4 MWh). To illustrate why this happens, we have plotted the ENS per time step for both the ROP and MRSP + ROP in Fig. 2b. It can be observed that by repairing only the minimum restoration set, the MRSP + ROP problem reaches full load delivery (i.e.,

Table 2
Performance of SOC Power Flow Formulation with the MRSP pre-processor on the 5-bus case .

Workflow	ENS	Time [s]
ROP	2902.4	2.11
MRSP + ROP	3081.4	0.103

ENS = 0) before the ROP solution. However, while the ROP solution repairs a larger number of components than is strictly necessary before reaching full load delivery, it serves more load while the repairs are ongoing.

5.3. Larger restoration problems with DC power flow

The third case study investigates the efficiency of the DC power flow formulation in larger networks. For these investigations, we use the IEEE 118-bus transmission case [34] and a single-phase equivalent (SPE) representation of the IEEE 123-bus distribution feeder case³, obtained from [35]. In the 118-bus case, we create a damage scenario by simulating significant localized damage to a third of the network, representing, e.g., an event such as a hurricane or an earthquake that causes severe damage to only one area of the transmission network. The damage scenario is generated randomly with 35% to branches and generators in Area 1 of the network, which includes buses 1–23, 25–32, 113–115, and 117. In the 123-bus distribution feeder, we simulate the impact of a windstorm by randomly damaging 25% of the branches in the network. For these two cases, we apply the DC ROP with subsequent AC redispatch. We assume an ability to repair multiple elements per period, such that we have 10 available repair periods for the 118-bus case, and 15 periods for the 123-bus case. At the end of the simulation, we record the solution time, estimated ENS and true ENS (summarized in Table 3).

Even for these larger test cases, the DC ROP problem solves within a few seconds, demonstrating the tractability of DC power flow based restoration algorithms. However, analysis of the solution quality reveals some important differences between the two test cases. For the IEEE 118-bus transmission case, there is a large discrepancy between the estimated and true ENS, indicating that the DC ROP drastically overestimates the amount of load that can be delivered. In contrast, for the IEEE 123-bus distribution feeder case, the estimated and true ENS are nearly equivalent.

To understand these differences, we first plot the estimated and true ENS for the 118-bus case in Fig. 2c. We observe that the true ENS is higher than the predicted ENS for all time periods. This can be explained by considering the structure of the network. The 118-bus case represents a well connected network with 35 synchronous condensers. In this network, reactive power plays a significant role in ensuring feasibility, which is not recognized by the DC power flow formulation. Instead, the DC power flow formulation delays the repair of the synchronous condensers because they are unnecessary for a feasible DC solution.

We do not plot the results for the 123-bus case, because there is almost no difference between the estimated and true ENS in any of the

³ The single-phase representation contains only 56 buses, but we will still refer to the case as the 123-bus case.

Table 3
Performance of DC Power Flow Formulation for the 118- and 123-bus cases.

Test Case	System Type	Estimated ENS [MWh]	True ENS [MWh]	Solve Time [s]
118 bus	Transmission	226.2	323.6	6.81
123 bus	Distribution SPE	1395.5	1396.4	0.955

time periods. Again, this can be explained by the structure of the case; the distribution system in the IEEE 123-bus case is a radial network where the substation is the only source of both active and reactive power. Due to the radial network structure and other network parameters, the restoration problem focuses on connectivity over the details of power flow, which is captured equally well by the DC and AC models. The estimated and true ENS therefore match, since the ROP prioritizes feeder branches to reconnect to the substation first. However, it remains to be demonstrated how close the obtained solution is to the true optimal solution in this case. While the single phase equivalent representation and the assumptions regarding the restoration process are less not entirely realistic for a distribution network, our results demonstrate how the network topology plays a role in whether or not a power flow formulation provides accurate results.

6. Conclusion

This work highlights the need for a new generation of power restoration algorithms to support grid resiliency in the context of large-scale disturbances, and illustrates the formidable computational challenges faced by such algorithms. To that end, PowerModelsRestoration is proposed as a flexible framework for the rapid exploration of power restoration algorithms. A validation of the framework was conducted on seminal network cases, including a 5-bus test case, the IEEE 118-bus transmission network and the 123-bus distribution feeder network. This validation replicated key results from related works [8,9], but show that the DC model is able to provide feasible (though potentially sub-optimal) restoration sequences when generators are redispatched in a multi-period AC OPF framework. The paper provides novel insights into the challenges of non-convex restoration formulations and the successes of convex relaxations, which had not been considered previously. Overall, PowerModelsRestoration represents a strong foundation for continued research on power restoration algorithms and resiliency analysis.

One important future direction of this work is to improve the support for restoration in distribution grids. While we have notably included support for distributed energy resources in PowerModelsRestoration, distribution restoration requires a number of additional considerations, including sectionalizing, limited availability of switching devices and loads, which can no longer be assumed to be continuous variables, as well as the consideration of three-phase, unbalanced grids, which could be achieved through integration with the PowerModelsDistribution package.

Another important avenue for future work is the integration of metaheuristics, such as hierarchical or adaptive decomposition schemes [7], for scaling the restoration ordering problem to much larger problems sizes, e.g., those with thousands of buses. Furthermore, our current formulation only considers the ordering of power component repair, while ignoring other important aspects such as optimal repair crew dispatching [8], and stability margin verification during the restoration process. As such, the work presented here represents a sub-problem to the overall challenge of power grid recovery. However, addressing the challenges of this foundational sub-problem will enable us to efficiently solve the more comprehensive and complex problems in the future.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by funding from the U.S. Department of Energys (DOE) Office of Electricity and Office of Energy Efficiency and Renewable Energy as part of the CleanStart-DERMS project of the Grid Modernization Laboratory Consortium and by the MACSER project, funded by U.S. Department of Energys (DOE) Office of Science.

References

- [1] T. Overbye, X. Cheng, Y. Sun, A comparison of the AC and DC power flow models for LMP calculations, *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*, (2004), p. 9pp.
- [2] C. Coffrin, P.V. Hentenryck, R. Bent, Accurate load and generation scheduling for linearized dc models with contingencies, 2012 IEEE Power and Energy Society General Meeting, (2012), pp. 1–8, <https://doi.org/10.1109/PESGM.2012.6345338>.
- [3] C. Coffrin, R. Bent, B. Tasseff, K. Sundar, S. Backhaus, Relaxations of ac maximal load delivery for severe contingency analysis, *IEEE Trans. Power Syst.* 34 (2) (2019) 1450–1458, <https://doi.org/10.1109/TPWRS.2018.2876507>.
- [4] M.M. Adibi, L.H. Fink, Power system restoration planning, *IEEE Trans. Power Syst.* 9 (1) (1994) 22–28, <https://doi.org/10.1109/59.317561>.
- [5] M. Adibi, P. Clelland, L. Fink, H. Happ, R. Kafka, J. Raine, D. Scheurer, F. Trefny, Power system restoration - a task force report, *IEEE Trans. Power Syst.* 2 (2) (1987) 271–277, <https://doi.org/10.1109/TPWRS.1987.4335118>.
- [6] D.S. Kirschen, T.L. Volkman, Guiding a power system restoration with an expert system, *IEEE Trans. Power Syst.* 6 (2) (1991) 558–566, <https://doi.org/10.1109/59.76698>.
- [7] C. Coffrin, P. Van Hentenryck, R. Bent, Last-mile restoration for multiple inter-dependent infrastructures, *AAAI*, 12 (2012), pp. 455–463.
- [8] P. Van Hentenryck, C. Coffrin, R. Bent, et al., Vehicle routing for the last mile of power system restoration, *Proceedings of the 17th Power Systems Computation Conference (PSCC11)*, Stockholm, Sweden, Citeseer, 2011.
- [9] P. Van Hentenryck, C. Coffrin, Transmission system repair and restoration, *Math. Program.* 151 (1) (2015) 347–373, <https://doi.org/10.1007/s10107-015-0887-0>.
- [10] C. Coffrin, R. Bent, K. Sundar, Y. Ng, M. Lubin, Powermodels.jl: an open-source framework for exploring power flow formulations, 2018 Power Systems Computation Conference (PSCC), (2018), pp. 1–8, <https://doi.org/10.23919/PSCC.2018.8442948>.
- [11] M.B. Cain, R.P. O' Neill, A. Castillo, *History of Optimal Power flow and Formulations*, Technical Report, Federal Energy Regulatory Commission, 2012.
- [12] R.A. Jabr, Radial distribution load flow using conic programming, *IEEE Trans. Power Syst.* 21 (3) (2006) 1458–1459, <https://doi.org/10.1109/TPWRS.2006.879234>.
- [13] C. Coffrin, H.L. Hijazi, P. Van Hentenryck, The qc relaxation: a theoretical and computational study on optimal power flow, *IEEE Trans. Power Syst.* 31 (4) (2016) 3008–3018, <https://doi.org/10.1109/TPWRS.2015.2463111>.
- [14] B. Stott, J. Jardim, O. Alsac, Dc power flow revisited, *IEEE Trans. Power Syst.* 24 (3) (2009) 1290–1300, <https://doi.org/10.1109/TPWRS.2009.2021235>.
- [15] L. Sun, Z. Lin, Y. Xu, F. Wen, C. Zhang, Y. Xue, Optimal skeleton-network restoration considering generator start-up sequence and load pickup, *IEEE Trans. Smart Grid* 10 (3) (2019) 3174–3185, <https://doi.org/10.1109/TSG.2018.2820012>.
- [16] S. Thiébaux, C. Coffrin, H. Hijazi, J. Slaney, Planning with mip for supply restoration in power distribution systems, *Twenty-Third International Joint Conference on Artificial Intelligence*, (2013).
- [17] H.L. Hijazi, S. Thiébaux, Optimal ac distribution systems reconfiguration, 2014 Power Systems Computation Conference, (2014), pp. 1–7, <https://doi.org/10.1109/PSCC.2014.7038420>.
- [18] J. Li, X. Ma, C. Liu, K.P. Schneider, Distribution system restoration with microgrids using spanning tree search, *IEEE Trans. Power Syst.* 29 (6) (2014) 3021–3029, <https://doi.org/10.1109/TPWRS.2014.2312424>.
- [19] Y. Tan, F. Qiu, A.K. Das, D.S. Kirschen, P. Arabshahi, J. Wang, Scheduling post-disaster repairs in electricity distribution networks, *IEEE Trans. Power Syst.* 34 (4) (2019) 2611–2621, <https://doi.org/10.1109/TPWRS.2019.2898966>.
- [20] R. Romero, A. Monticelli, A. Garcia, S. Haffner, Test systems and mathematical models for transmission network expansion planning, *IEE Proc. - Gener. Trans. Distrib.* 149 (1) (2002) 27–36, <https://doi.org/10.1049/ip-gtd:20020026>.
- [21] T. Kalinowski, D. Matsypura, M.W. Savelsbergh, Incremental network design with maximum flows, *Eur. J. Oper. Res.* 242 (1) (2015) 51–62, <https://doi.org/10.1016/j.ejor.2014.10.003>.
- [22] Gurobi Optimization, Inc., Gurobi optimizer reference manual, 2014, (Published online at <http://www.gurobi.com>).
- [23] I. IBM, IBM ILOG CPLEX Optimization Studio, 2014, (<http://www-01.ibm.com/software/commerce/optimization/cplex-optimizer/>).

- [24] C. Coffrin, P.V. Hentenryck, Transmission system restoration with co-optimization of repairs, load pickups, and generation dispatch, *International Journal of Electrical Power & Energy Systems* 72 (2015) 144–154, <https://doi.org/10.1016/j.ijepes.2015.02.027>. The Special Issue for 18th Power Systems Computation Conference.
- [25] C. Coffrin, P. Van Hentenryck, A linear-programming approximation of ac power flows, *INFORMS J. Comput.* 26 (4) (2014) 718–734, <https://doi.org/10.1287/ijoc.2014.0594>.
- [26] I. Dunning, J. Huchette, M. Lubin, Jump: a modeling language for mathematical optimization, *SIAM Rev.* 59 (2) (2017) 295–320, <https://doi.org/10.1137/15M1020575>.
- [27] J. Bezanson, A. Edelman, S. Karpinski, V. Shah, Julia: a fresh approach to numerical computing, *SIAM Rev.* 59 (1) (2017) 65–98, <https://doi.org/10.1137/141000671>.
- [28] A. Wächter, L.T. Biegler, On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming, *Math. Program.* 106 (1) (2006) 25–57.
- [29] J. Forrest, T. Ralphs, S. Vigerske, LouHafer, B. Kristjansson, jpfasano, EdwinStraver, M. Lubin, H.G. Santos, rlougee, M. Saltzman, coin-or/cbc: Version 2.9.9, 2018DOI: 10.5281/zenodo.1317566.
- [30] O. Kröger, C. Coffrin, H. Hijazi, H. Nagarajan, Juniper: An open-source nonlinear branch-and-bound solver in julia, in: W.-J. van Hoes (Ed.), *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, Springer International Publishing, Cham, 2018, pp. 377–386.
- [31] A. Wächter, L.T. Biegler, On the implementation of primal-dual interior point filter line search algorithm for large-scale nonlinear programming, *Math. Prog.* 106 (1) (2006) 25–57, <https://doi.org/10.1007/s10107-004-0559-y>.
- [32] O. Kröger, C. Coffrin, H. Hijazi, H. Nagarajan, Juniper: An Open-source Nonlinear Branch-and-bound Solver in Julia, *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, Springer International Publishing, 2018, pp. 377–386.
- [33] F. Li, R. Bo, Small test systems for power system economic studies, *IEEE PES General Meeting*, (2010), pp. 1–4, <https://doi.org/10.1109/PES.2010.5589973>.
- [34] S. Babaeinejadsarookolae, A. Birchfield, R.D. Christie, C. Coffrin, C. DeMarco, R. Diao, M. Ferris, S. Fliscounakis, S. Greene, R. Huang, C. Jozs, R. Korab, B. Lesieutre, J. Maeght, D.K. Molzahn, T.J. Overbye, P. Panciatici, B. Park, J. Snodgrass, R. Zimmerman, The power grid library for benchmarking ac optimal power flow algorithms, 2019.
- [35] S. Bolognani, S. Zampieri, On the existence and linear approximation of the power flow solution in power distribution networks, *IEEE Trans. Power Syst.* 31 (1) (2016) 163–172.