

MODELING AND OPTIMIZATION TECHNIQUES FOR CRITICAL
INFRASTRUCTURE RESILIENCE

A Dissertation

Submitted to the Faculty

of

Purdue University

by

Michael L. Bynum

In Partial Fulfillment of the

Requirements for the Degree

of

Doctor of Philosophy

December 2018

Purdue University

West Lafayette, Indiana

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF DISSERTATION APPROVAL

Dr. Carl D. Laird

Davidson School of Chemical Engineering

Dr. Gintaras V. Reklaitis

Davidson School of Chemical Engineering

Dr. Zoltan Nagy

Davidson School of Chemical Engineering

Dr. Andrew L. Liu

School of Industrial Engineering

Approved by:

Dr. Sangtae Kim

Head of the School Graduate Program

To my loving and supportive wife and daughter.

ACKNOWLEDGMENTS

I would like to express my sincere gratitude for all of the support provided by my advisor, Professor Carl D. Laird, throughout my PhD studies. This dissertation would not be possible without him. Most importantly, I would like to thank him for the numerous opportunities he has provided, each of which has led me to a career that I believe will be truly fulfilling. I will always admire Carl for his quick intuition, vast knowledge, and passion for teaching and guiding students.

I would like to thank my wife, Andrea, for all of her support and encouragement and my daughter, Cora, for the inspiration she has given me. I would like to thank my parents for being there every step of the way. Finally, the advice, ideas, and knowledge shared from both my group members at Purdue and my collaborators at Sandia have been invaluable. In particular, I would like to thank Yankai, Arpan, Jianfeng, Santiago, Todd, Kate, JP, and Anya.

This research was funded in part by the U.S. Environmental Protection Agency (EPA) through its Office of Research and Development (DW8992450201) under an Interagency Agreement with the Department of Energy's Sandia National Laboratories. It has been subjected to the Agency's review and has been approved for publication. Note that approval does not signify that the contents necessarily reflect the views of the Agency. Mention of trade names products, or services does not convey official EPA approval, endorsement, or recommendation. This research was also funded in part by Sandia National Laboratories' Laboratory Directed Research and Development program. This research was also funded in part by the U.S. Department of Energy's Office of Electricity, under the Advanced Grid Modeling (AGM) program. This research was also funded in part by the U.S. Department of Homeland Security.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energys National Nuclear Security Administration under contract DENA0003525.

Disclaimer: This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

©2018 IEEE Part of this dissertation is reprinted with permission from “Tightening McCormick Relaxations Toward Global Solution of the ACOPF Problem” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., IEEE Transactions on Power Systems, 2018.

In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of Purdue University’s products or services.

Part of this dissertation is reprinted from “13th International Symposium on Process Systems Engineering (PSE 2018)”, Volume 44, Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., “Strengthened SOCP Relaxations for ACOPF with McCormick Envelopes and Bounds Tightening”, pages 1555–1560, 2018, with permission from Elsevier.

Part of this dissertation is reprinted from “Evaluating Demand Response Opportunities for Power Systems Resilience Using MILP and MINLP Formulations” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., to appear in AIChE Journal, 2018.

Part of this dissertation is reprinted with permission from “A Software Framework for Assessing the Resilience of Drinking Water Systems to Disasters with an Example Earthquake Case Study” by Klise, K.A., Bynum, M., Moriarty, D., and Murray, R., Environmental Modelling & Software, 2017.

TABLE OF CONTENTS

	Page
LIST OF TABLES	x
LIST OF FIGURES	xii
ABBREVIATIONS	xv
ABSTRACT	xvi
1 INTRODUCTION	1
1.1 The Importance of Critical Infrastructure Resilience	1
1.2 Existing Methods and Software Tools for Evaluating and Improving Critical Infrastructure Resilience	3
1.2.1 Power Systems	3
1.2.2 Water Distribution Systems	5
1.3 Research Problem Statement and Scope of Work	9
1.4 Thesis Outline	10
 I Models and Algorithms for Improving Power Systems Resilience	 14
2 MODELING POWER SYSTEMS	15
2.1 Alternating Current Power Flow Equations	15
2.2 The DC Approximation	18
3 STOCHASTIC MIXED INTEGER LINEAR PROGRAMS FOR IMPROV- ING POWER SYSTEMS RESILIENCE TO EXTREME WEATHER	21
3.1 A Resilient Dispatch Formulation	22
3.2 Scenario Construction	25
3.3 Extensions for Long Term Investments	26
3.3.1 Transmission Line Hardening	26
3.3.2 Increasing Transmission Capacity	27
3.4 Computational Results	28
3.5 Out-of-Sample Cross Validation	31
3.6 The Role of Energy Consumers	33
3.6.1 Problem Formulation	36
3.6.2 Numerical Results	38
3.7 Summary	42

	Page
4 THE EFFECT OF THE DC APPROXIMATION	45
4.1 A Stochastic Mixed Integer Nonlinear Programming Formulation	45
4.2 MISOCP Relaxation	48
4.3 Solution Approach	50
4.4 A Comparison of AC and DC Solutions	51
4.5 Summary	59
5 RELAXATIONS AND REFINEMENT TECHNIQUES FOR AC CONSTRAINED POWER FLOW PROBLEMS	61
5.1 A Review of Convex Relaxations for the ACOPF Problem	62
5.1.1 Second-Order Cone Programming Relaxation	64
5.1.2 McCormick Relaxations of the Rectangular Form	64
5.1.3 Quadratic Relaxations of the Polar Form	65
5.2 Numerical Results	66
5.3 The Impact of the Reference Bus	73
5.4 Summary	74
6 DECOMPOSING OPTIMIZATION-BASED BOUNDS TIGHTENING PROB- LEMS VIA GRAPH PARTITIONING	75
6.1 Decomposed Bounds Tightening Algorithm	79
6.2 Graph Partitioning	82
6.3 Application to Optimal Power Flow	84
6.4 Results	84
6.5 Summary	90

II Software Tools for Water Distribution System Resilience 91

7 WATER NETWORK TOOL FOR RESILIENCE	93
7.1 Modeling Framework	93
7.1.1 Software Availability	95
7.1.2 Earthquake Attenuation Models	95
7.1.3 Fragility Curves	97
7.1.4 Flexible Controls	98
7.1.5 Leak Model	98
7.1.6 Pressure-Driven Demand Hydraulics	99
7.1.7 Resilience Metrics	100
7.1.8 Monte Carlo Simulation	102
7.2 Earthquake Case Study	102
7.2.1 Earthquake Damage	104
7.2.2 Repair Strategies	109
7.2.3 Case Study Results	110

		Page
7.3	Summary	119
8	SUMMARY	121
8.1	Thesis Summary and Contributions	121
8.1.1	Stochastic Programming Formulations for Improving Power Systems Resilience to Extreme Weather	121
8.1.2	The Effect of the DC Power Flow Approximation	122
8.1.3	Relaxations and Refinement Techniques for AC Constrained Power Flow Problems	123
8.1.4	Decomposed Bounds Tightening	124
8.1.5	Water Network Tool for Resilience: A Python Package	124
8.2	Recommendations for Future Work	126
8.2.1	Stochastic Programming Formulations for Improving Power Systems Resilience to Extreme Weather	126
8.2.2	Solution Techniques for ACPF-based MINLP's	127
8.2.3	Water Distribution System Resilience	127
	REFERENCES	129

LIST OF TABLES

Table	Page
3.1 Summary of network sizes	28
3.2 Problem size statistics	29
4.1 Integer Solutions DC - pglib_opf_case14_ieee	55
4.2 Integer Solutions AC - pglib_opf_case14_ieee	55
4.3 Comparison of DR locations selected by DC and AC models for pglib_opf_case14_ieee across all values of Δ_b and N_Z	56
4.4 Average fraction of AC-feasible scenarios when implementing DC solution .	56
5.1 Optimality Gaps for relaxations of NESTA archive standard operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE	67
5.2 Optimality Gaps for relaxations of NESTA archive congested operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE	68
5.3 Optimality Gaps for relaxations of NESTA archive small angle difference operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE	69
5.4 Computational Performance for relaxations of NESTA archive standard operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE	70
5.5 Computational Performance for relaxations of NESTA archive congested operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE	71

Table	Page
5.6 Computational Performance for relaxations of NESTA archive small angle difference operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE	72
6.1 Number of branches removed in order to partition the root graph.	90
7.1 Minimum water service availability (WSA) and recovery time for each earthquake magnitude, location, and repair strategy. Values are computed using the median of the 50 realizations. A recovery time of 0 means that the median value never went below 90% water service availability.	117
7.2 Maximum population impacted and recovery time for each earthquake magnitude, location, and repair strategy. Values are computed using the median of the 50 realizations. A recovery time of 0 means that the median value never went above 10% population impacted.	118

LIST OF FIGURES

Figure	Page
3.1 The left and right panels illustrate results for THF and TCF, respectively, for case30. The y-axis indicates the expected load shed across all scenarios. The asterisk represents the <i>baseline</i> , which is the case where no preparation is done before the severe weather event. The points at zero hardened lines (or zero transmission lines with increased capacity) correspond to the RDF solution.	31
3.2 The left and right panels illustrate results for THF and TCF, respectively, for case2383wp. The y-axis indicates the expected load shed across all scenarios. The asterisk represents the <i>baseline</i> , which is the case where no preparation is done before the severe weather event. The points at zero hardened lines (or zero transmission lines with increased capacity) correspond to the RDF solution.	31
3.3 Incremental reduction in expected load shed for each additional hardened line (left) and each additional line with increased capacity (right) for case30.	32
3.4 Incremental reduction in expected load shed for each additional hardened line (left) and each additional line with increased capacity (right) for case2383wp.	32
3.5 Cross validation results for case30. The line labeled “True Scenarios” shows the expected load shed when the solution to the stochastic programming problem is evaluated on the same 100 scenarios used to solve the stochastic programming problem. The remaining lines show the expected load shed when the stochastic programming problem is solved with a different set of scenarios than those used to evaluate the expected load shed.	34
3.6 The value of DR on resilience as a function of the number of demand response contracts (N_z) and the allowable fraction of DR (Δ_b) with 100 scenarios.	41
4.1 Comparison of DC and AC results for pglib-opf_case14_ieee. Each figure shows the minimum number of demand response contracts as a function of Δ_b and N_z	53

Figure	Page
4.2 Comparison of DC and AC results for <code>pglib_opf_case30_as</code> . Each figure shows the minimum number of demand response contracts as a function of Δ_b and N_z	54
4.3 Fraction of scenarios feasible when integer solutions obtained with the DC model are evaluated with the AC model ($\Delta_b = 0.2$). The solid lines show the minimum number of hardened branches and the dashed lines show the fraction of feasible scenarios, both as a function of N_z	57
4.3 Fraction of scenarios feasible when integer solutions obtained with the DC model are evaluated with the AC model ($\Delta_b = 0.2$). The solid lines show the minimum number of hardened branches and the dashed lines show the fraction of feasible scenarios, both as a function of N_z	58
5.1 Impact of the reference bus on bounds tightening for <code>nesta_case24_ieee_rts_api</code> . The abscissa shows the iteration count, and the ordinate shows the average range between the upper and lower variable bounds.	74
6.1 Wallclock time for one iteration of BT	85
6.2 <code>pglib_opf_case30_ieee</code>	86
6.3 <code>pglib_opf_case57_ieee_sad</code>	86
6.4 <code>pglib_opf_case73_ieee_rts_api</code>	87
6.5 <code>pglib_opf_case118_ieee</code>	87
6.6 <code>pglib_opf_case162_ieee_dtc</code>	88
6.7 <code>pglib_opf_case240_pserc</code>	88
6.8 <code>pglib_opf_case300_ieee</code>	89
7.1 Water distribution network model showing A) pipe material and B) liquefaction potential.	103
7.2 Water distribution network model showing the fault line (dashed red line), node elevation, and epicenter of three earthquakes. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)	105
7.3 Fragility curves for A) tank, B) pump, and C) pipe damage.	107
7.4 A) PGA and B) location of pipe damage (in red) and tank damage (in blue) for a magnitude of 6.5 earthquake at the central epicenter location. The yellow star indicates the location of the central epicenter. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)	108

Figure	Page
7.5 (A) Leak rate and (B) tank level as a function of time for a single realization of a magnitude 6.5 earthquake at the central location using no conservation (RS1). In A), a particular leak is shown in red illustrating the possibility of a leak becoming more prominent as the system is restored. In B), damaged tanks are shown in red up until the time they are repaired, tanks are refilled as the system is restored. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)	111
7.6 Water service availability for a magnitude 6.5 earthquake at the central location using no conservation (RS1). Light gray is the 5th-95th percentile, dark gray is the 25th-75th percentile, and the black line is the median. Threshold for recovery time shown as a dotted line.	112
7.7 Firefighting capacity for a magnitude 6.5 earthquake at the central location using no conservation (RS1). Light gray is the 5th-95th percentile, dark gray is the 25th-75th percentile, and the black line is the median. Firefighting capacity is computed every 12 h (indicated by x's) at the location of a particular fire.	113
7.8 Population impacted (A) without conservation (RS1), (B) with conservation (RS2), and (C) with seismic-resistant pipes (RS3) for 50 realizations of a magnitude 6.5 earthquake at the central location. Light gray is the 5th-95th percentile, dark gray is the 25th-75th percentile, and the black line is the median. Threshold for recovery time shown as a dotted line.	114
7.9 Effect of location of the earthquake on maximum population impacted for a 6.5 magnitude earthquake at the central location and no conservation (RS1).	116

ABBREVIATIONS

MINLP	mixed-integer nonlinear program
MILP	mixed-integer linear program
NLP	nonlinear program
AC	alternating current
DC	linear approximation of alternating current power flows
ACPF	alternating current power flow
DCPF	linear approximation of alternating current power flows
SOC	second-order cone
SOCP	second-order cone program
DR	demand response
MISOCP	mixed-integer second order cone programming
ACOPF	alternating current optimal power flow
QC	quadratic convex
ISO	Independent System Operator
RTO	Regional Transmission Operator

ABSTRACT

Bynum, Michael L. PhD, Purdue University, December 2018. Modeling and Optimization Techniques for Critical Infrastructure Resilience. Major Professor: Carl D. Laird.

The resilience of critical infrastructure, such as water distribution systems and power systems, is critical for both the economy and public safety and health. However, methods and tools for evaluating and improving the resilience of these systems must be able to address the large network sizes, nonlinear physics, discrete decisions, and uncertainty. This dissertation focuses on the development of modeling and optimization techniques that address these difficulties, enabling the evaluation and improvement of power and water distribution system resilience.

In Part I, we present novel stochastic optimization models to improve power systems resilience to extreme weather events. We consider proactive redispatch, transmission line hardening, and transmission line capacity increases as alternatives for mitigating the effects of extreme weather. Our model is based on linearized or "DC" optimal power flow, similar to models in widespread use by independent system operators (ISOs) and regional transmission operators (RTOs). Our computational experiments indicate that each of these strategies can play a major role in power systems resilience.

We then extend the resilience formulations to investigate the role chemical process facilities, as industrial energy consumers, can play in improving electric grid resilience through demand response (DR). For process facilities to effectively negotiate demand response (DR) contracts and make investment decisions regarding flexibility, they need to quantify their additional value to the grid. We also reformulate the DR problems using the more accurate nonlinear alternating current power flow model to

investigate the effect of the linear DC approximation. Our numerical results demonstrate that the linearized model often underestimates the amount of DR needed, motivating scalable solution algorithms for Mixed-Integer Nonlinear Programming (MINLP) problems in power systems.

An important step in many MINLP algorithms is the global solution of a Nonlinear Programming (NLP) subproblem. For power systems applications, this involves global solution of NLP's containing the alternating current (AC) power flow model. This thesis presents several advances to aid in global optimization of AC power flow equations. We show that a strong upper bound on the objective of the alternating current optimal power flow (ACOPF) problem can significantly improve the effectiveness of optimization-based bounds tightening (OBBT) on a number of relaxations. Furthermore, we investigate the effect of the reference bus on OBBT. We find that, if reference bus constraints are included, relaxations of the rectangular form significantly strengthen existing relaxations and that the effectiveness of OBBT at a given iteration is directly related to the distance of the corresponding bus from the reference bus.

Ultimately, with OBBT alone, we are able to reduce the optimality gap to less than 0.1% on all but 5 NESTA test cases with up to 300 buses. However, the computational expense required for OBBT grows rapidly with the size of the network. We present a decomposition algorithm based on graph partitioning to drastically improve this performance. Our numerical results demonstrate that our decomposed bounds tightening (DBT) algorithm results in variable bounds nearly as tight as those obtained with traditional, full-space OBBT. Furthermore, the computational expense of the DBT algorithm scales far more favorably with problem size, resulting in drastically reduced wallclock times, especially for large networks.

In Part II, we describe the Water Network Tool for Resilience (WNTR), an new open source Python package designed to help water utilities investigate resilience of water distribution systems to hazards and evaluate resilience-enhancing actions. The WNTR modeling framework is presented and a case study is described that

uses WNTR to simulate the effects of an earthquake on a water distribution system. The case study illustrates that the severity of damage is not only a function of system integrity and earthquake magnitude, but also of the available resources and repair strategies used to return the system to normal operating conditions. While earthquakes are particularly concerning since buried water distribution pipelines are highly susceptible to damage, the software framework can be applied to other types of hazards, including power outages and contamination incidents.

1. INTRODUCTION¹

1.1 The Importance of Critical Infrastructure Resilience

The United States Presidential Policy Directive 21 (PPD-21) originally outlined the nation’s need for secure and resilient critical infrastructure and identified 16 critical infrastructure sectors, including the electric power sector and the water sector [The White House, 2013]. Subsequently, national focus on improving critical infrastructure resiliency has only increased. In particular, the resilience of electric power systems and water distribution systems is vital for both the economy and public health and safety. Physical components of these systems may be damaged by natural disasters (e.g., earthquakes and hurricanes) and by intentional acts (e.g., perpetrated by nation-state actors), ultimately impacting the ability to deliver power and/or clean water to consumers. The annual cost of power outages due to severe weather in the United States alone between 2003 and 2012 is estimated to be between 18 and 33 billion USD [Executive Office of the President, 2013].

Informally, resilience refers to the ability of a system to withstand and quickly recover from adverse events [Watson et al., 2014]. Here, we specifically focus on re-

¹Part of this chapter is reprinted from “Evaluating Demand Response Opportunities for Power Systems Resilience Using MILP and MINLP Formulations” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., to appear in *AICHE Journal*, 2018.

©2018 IEEE Part of this chapter is reprinted with permission from “Tightening McCormick Relaxations Toward Global Solution of the ACOPF Problem” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., *IEEE Transactions on Power Systems*, 2018. In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of Purdue University’s products or services.

Part of this chapter is reprinted from “13th International Symposium on Process Systems Engineering (PSE 2018)”, Volume 44, Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., “Strengthened SOCP Relaxations for ACOPF with McCormick Envelopes and Bounds Tightening”, pages 1555–1560, 2018, with permission from Elsevier.

Part of this chapter is reprinted with permission from “A Software Framework for Assessing the Resilience of Drinking Water Systems to Disasters with an Example Earthquake Case Study” by Klise, K.A., Bynum, M., Moriarty, D., and Murray, R., *Environmental Modelling & Software*, 2017.

silience to severe weather events, e.g., hurricanes and ice storms. A necessary first step toward improving resilience to such events is estimating the impact of these events on power and water distribution systems. Such estimation involves weather forecasting, predicting the effect of severe weather conditions on system components, and modeling the performance (operations) of the systems given sets of damaged components [Panteli and Mancarella, 2015, Ouyang and Duenas-Osorio, 2014, Guikema et al., 2010, Panteli et al., 2017a].

Panteli et al. [2017b] describe the concept of a “resilience trapezoid” toward evaluation and quantification of critical infrastructure resilience. The resilience trapezoid highlights the performance of a system during an adverse event (Phase I), immediately following an adverse event (Phase II), and during the restoration process (Phase III). Panteli et al. [2017b] subsequently introduced the $\Phi\Lambda E\Pi$ resilience metric system based upon the different phases of the resilience trapezoid. The $\Phi\Lambda E\Pi$ resilience metric system is composed of four components:

- Φ : The slope of system performance during Phase I
- Λ : The amount by which performance drops from Phase I to Phase II (i.e., the baseline performance minus the post-disturbance performance)
- E : The duration of Phase II
- Π : The slope of the system performance during Phase III

The resilience trapezoid and $\Phi\Lambda E\Pi$ metric system provides a conceptual framework for categorizing different approaches to improving resilience. In the remainder of this chapter, we review existing approaches for improving the resilience of power and water distribution systems within the context of the resilience trapezoid, review limitations in current methods and tools for evaluating and improving resilience, and outline the rest of this dissertation.

1.2 Existing Methods and Software Tools for Evaluating and Improving Critical Infrastructure Resilience

In this section, we review existing methods and tools for evaluating and improving the resilience of power and water distribution systems.

1.2.1 Power Systems

In the context of power systems operations, the primary considerations for Phase I of the resilience trapezoid (during an adverse event) are minimization of unserved demand due to component damage and the prevention of cascading outages (i.e., blackouts). Cascading outages are complex phenomena due to the complex interactions between system components, component failures, and protection schemes. Many different failure types may contribute to a blackout, including cascading overloads, transient instabilities, and voltage collapse [Baldick et al., 2008]. Typical mitigation strategies include remedial action schemes such as generation trip, brake insertion, fast valve/gen ramp, HVDC ramp, islanding, intentional load shed, excitation forcing, shunt capacitor/reactor switching, and series capacitor/reactor switching [Vaiman et al., 2013]. Chen et al. [2001] propose a decision-event tree to help operators respond rapidly to an event and prevent cascading outages. Wide area monitoring and backup protection systems have also been proposed for preventing blackouts [Tan et al., 2002, Zima and Andersson, 2004]. Song and Kezunovic [2007] present an early detection scheme that uses the vulnerability index, margin index, and power flow solutions to predict possible failures. Cascading outages may also be prevented by partitioning the network into islands by minimal cut sets [Wang and Vittal, 2004] and using graph partitioning methods in the context of mixed-integer linear programming (MILP) models [Fan et al., 2012].

Extensive research has been conducted to develop improved strategies for post-blackout restoration (Phase III). Wang et al. [2016b] divide restoration into three stages (preparation, system restoration, and load restoration) and provide a review

of research conducted to address each stage. Preparation primarily involves identification of the status of various system components [Adibi and Fink, 1994b]. System restoration deals with restarting generators and reconnecting and synchronizing the power grid. For example, Sun et al. [2011] propose a MILP model to determine an optimal generator restart sequence in order to maximize system generation capability. Load restoration (also known as distribution system restoration) deals with scheduling load pickup while maintaining system frequency. This process must be coordinated with generation capability [Wang et al., 2016b, Adibi and Fink, 1994b].

Effectively improving the performance of power systems in Phase II (i.e., increasing the post-disturbance system performance) can be challenging. Many hardening options are available to achieve this goal, including hardening transmission lines, elevating substations, the addition of distributed generation, the introduction of storage devices, and system network reconfiguration (e.g., switching) [Panteli et al., 2017b]. Computational approaches are generally required to rigorously select from among these various options. Zare et al. [2017] introduce a stochastic MILP optimization model for switch placement in power distribution systems in order to isolate faulted areas; network flows are not considered. In general, determining which option, or combination of options, to employ is computationally challenging, driving research toward more efficient approaches. Panteli et al. [2017b] propose hardening transmission lines according to each corridor’s Resilience Achievement Worth (RAW) index. Wang et al. [2018] perform Monte Carlo sampling to assess line outage impacts on generation unit scheduling. However, such heuristic approaches to improving resilience generally do not provide optimal solutions. They may provide significantly worse-than-optimal solutions in the presence of a diverse set of strategies and large system sizes.

Alternatively, rigorous optimization approaches based on mathematical programming can provide optimal solutions in these contexts, although the computational challenges can be significant. This difficulty is illustrated in the body of literature on bi-level and tri-level optimization models and corresponding solution algorithms for

critical infrastructure defense planning [Wu and Conejo, 2017, Salmeron et al., 2004, Alguacil et al., 2014]. These models consider critical infrastructure defense strategies in the context of an intelligent adversary, and focus on mitigating worst-case impacts to a power system. For example, Shao et al. [2017] and Wang et al. [2013] use similar strategies to mitigate the impact of component failures (contingencies) within the context of long-term (i.e., investment) and short-term (i.e., operations) planning, respectively. The former introduces a two-stage tri-level optimization model and imposes an upper limit on load shed. The latter introduces a two-stage robust optimization model to minimize a linear combination of operating cost and load shed, focusing again on the worst-case loss induced by component failure. Both of these models are solved using decomposition strategies in which the model is partitioned into a MILP master problem and bi-level subproblems. The bi-level subproblems (which are typically more computationally demanding) are used to find the worst-case contingencies, which are iteratively integrated into the master problem. Hardening strategies recommended by such models are designed to be effective against an omniscient adversary with full knowledge of the power system under consideration, e.g., a nation state actor.

1.2.2 Water Distribution Systems

Depending on their location and vulnerability, drinking water utilities are taking steps to enhance their resilience to earthquakes, floods, drought, hurricanes, winter storms, forest fires, tornadoes, contamination incidents, terrorist attacks, and other types of hazards [Critical Infrastructure Partnership Advisory Council (CIPAC), 2009, American National Standards Institute (ANSI), 2010, US Environmental Protection Agency (USEPA), 2014, 2015b, 2016]. Resilience is defined by the National Academies of Science as the ability to prepare and plan for, absorb, recover from, and successfully adapt to adverse events [National Academy of Sciences (NAS), 2012]. The ability to maintain drinking water service during and following such hazardous incidents is critical to ensure the well-being and continuity of daily life. Water system resilience

is important not only for individuals, but also for hospitals, schools, nursing homes, fire stations, restaurants, and for other industries that rely on water.

Natural disasters and other types of hazards have resulted in different types of water service disruptions: pipe breaks and leaks; power outages; failure of reservoirs, tanks, pumps, treatment plants, and other infrastructure; reduced water quality; loss of access to facilities and supplies; as well as financial, social, environmental and human health consequences [Critical Infrastructure Partnership Advisory Council (CIPAC), 2009, US Environmental Protection Agency (USEPA), 2015c, Eidinger and Davis, 2012]. Following large disruptive incidents like earthquakes, affected communities have experienced power outages and water service outages lasting from hours to weeks. For example, the 1994 magnitude 6.7 Northridge Earthquake, located outside Los Angeles, California, damaged seven reservoirs, over 60 transmission mains, and 1,000 distribution pipes. The quantity of water delivered was restored to pre-earthquake volumes after seven days, and the quality of water was restored and boil-water orders were lifted after 12 days. However, it took nine years to complete all repairs and restore full functionality of the water system [Davis, 2014]. Even though evidence shows that seismic-resistant pipes have a high survival rate following an earthquake [Eidinger and Davis, 2012], these upgrades have not been widely implemented because of the high cost for pipe replacement.

General guidance is available on water system resilience to disasters [Critical Infrastructure Partnership Advisory Council (CIPAC), 2009, American National Standards Institute (ANSI), 2010, US Environmental Protection Agency (USEPA), 2014, 2015b, 2016]; however, robust software tools to support utility-specific resilience assessment are not available. With such tools, water utilities could estimate potential damages to their system, understand the multitude of disruptions that could unfold over time, investigate redundancies, evaluate preparedness, and prioritize specific mitigation strategies, such as pipe replacement or adding redundancy to supply systems [American National Standards Institute (ANSI), 2010, US Environmental Protection Agency (USEPA), 2015c]. Systems modeling tools have the potential to meet this

need by combining disaster damage models with hydraulic and water quality models of water distribution systems. Additionally, systems modeling could incorporate changes in customer behavior during disasters (e.g., water usage), as well as utility response actions. This type of simulation approach could allow a water utility to design more effective mitigation activities before a disaster occurs.

Existing hydraulic and water quality modeling software tools were not designed to handle sudden failures resulting in inadequate pressure or rapid changes in system operation. Moreover, they were not designed to handle situations when sections of a water system become isolated, tanks or reservoirs drain, or the system operational rules cannot be met. For example, commonly used demand-driven (DD) hydraulic simulators, like EPANET [Rossman, 2000], assume customer demands are always met even if the pressure is insufficient to provide the demand. In reality, disasters can lead to situations where pressure in the system is unusually low and customer demand would not be met.

Several alternatives to DD simulators have been discussed in the literature. Wagner et al. [1988] presented pressure-driven demand (PDD) hydraulic equations for water distribution systems in which the demand supplied to a node is a function of the pressure at that node [Muranho et al., 2014]. During low-pressure conditions, customers receive a fraction of their expected demand. PDD simulators include WaterNetGen [Muranho et al., 2014, 2012], which is an open source software tool, and WaterGEMSTM [Wu et al., 2008], which is a commercial software tool. Quasi-PDD simulators (or semi-PDD) run DD simulations in an iterative manner and nodes are switched between constant-demand nodes, zero-demand nodes, and (sometimes) emitters depending on the domain in which the pressure falls [Trifunovic, 2012, Pathirana, 2010, Yoo et al., 2015]. When considering disaster scenarios, the difference between using DD and PDD simulation can be drastic [LauCELLi et al., 2012].

After a large-scale disruption, water demand in the network can change dramatically. Structural damage and emergency operational changes can result in isolated sections of the network or low pressure conditions that reduce the amount of water

delivered to customers. Policy changes, including do-not- drink orders, boil-water orders, or water conservation efforts, also decrease customer demand. When planning for an adequate drinking water supply during emergencies, water utilities need to account for the minimum acceptable water use per capita, the anticipated time scale of the disruption, the population impacted, and water quality standards that need to be upheld [US Environmental Protection Agency (USEPA), 2015a]. If acceptable water volume and quality cannot be delivered, potable water alternatives would have to be considered. Customer behavior can also change during emergencies, either temporarily increasing demand (e.g., filling up bathtubs to stockpile water) or decreasing demand (e.g., relying on bottled water because of a lack of confidence in the delivered water). These changes in customer demand can impact resilience; for example, conservation might increase water availability for firefighting.

Several tools have been developed to estimate the hydraulic performance of a water distribution system after an earthquake [Yoo et al., 2015, Mani et al., 2013, Shi and O'Rourke, 2008, Markov et al., 1994]. These tools use attenuation models and fragility curves developed by the American Lifeline Alliance to estimate network damage based on earthquake magnitude and location [American Lifelines Alliance (ALA), 2001a,b, 2005]. Attenuation models calculate peak ground acceleration (PGA) and peak ground velocity (PGV) as a function of earthquake magnitude, location and depth. Fragility curves determine the probability of damage as a function of ground movement. These models are generally built on empirical data from recent earthquakes, which includes information such as the characteristics of damaged pipes and the measured PGA. The Graphical Interactive Serviceability Analysis of Life-Lines subjected to Earthquakes (GISALLE) tool quantifies water service availability after an earthquake and uses the Loma Prieta Earthquake as the basis for a stochastic parametric study [Markov et al., 1994]. The Graphical Iterative Response Analysis for Flow Following Earthquakes (GIRAFFE) software builds upon GISALLE to include upgrades to the way low pressure conditions are handled when modeling pipe failure and includes a seismic wave model to estimate joint damage [Shi and O'Rourke,

2008]. Mani et al. [2013] include pipe leak models developed by Shi and O’Rourke [2008] and use the Tehran water distribution network as a case study. The Reliability EVALuation model for Seismic hazard for water supply NETwork (REVAS.NET) tool includes earthquake attenuation models and probabilistic scenarios with different repair strategies [Yoo et al., 2015]. Guidotti et al. [2016] extend previous work by including a general procedure for modeling resilience of critical network infrastructure. This work includes methods that account for dependencies between networks, such as the water distribution systems and the electric power network. The general procedure can be applied to a wide range of hazards and recovery actions. To date, these research efforts use quasi-PDD hydraulic simulations.

Several studies have simulated the damage to the Los Angeles Department of Water and Power caused by the 1994 Northridge Earthquake in California and the restoration activities that followed [Tabucchi, 2007, Romero et al., 2010, Shi and O’Rourke, 2008]. These tools combine GIRAFFE with the capability to demonstrate restoration actions, such as sending crews to investigate, isolate, and repair pipe breaks. Other modeling tools were developed to help the East Bay Municipal Water District in Northern California manage earthquake response in real-time, prioritize transmission line upgrades, and assess interdependencies with the electricity sector [Irias et al., 2011]. These tools combine real-time USGS ShakeMap data on ground movement after an earthquake [Wald et al., 2006] with a customized software tool to rapidly predict damage to specific water utility assets [Irias et al., 2011].

1.3 Research Problem Statement and Scope of Work

As we have seen, extensive research has been done on estimating the effects of natural disasters and severe weather events on physical components of both power systems and water distribution systems. Similarly, extensive work has been done on modeling, evaluating, and, in some cases, optimizing recovery strategies. However, although it is widely recognized that planning and investment before an event can significantly mitigate the effects of the event, existing techniques for effectively directing

resources toward this end are limited. Existing methods either are not computationally efficient enough for large scale implementation, do not incorporate network effects (e.g., power flow through a transmission network), are heuristic in nature (leading to sub-optimal solutions), or do not adequately account for uncertainty. Furthermore, there is a lack of software tools for performing resilience analyses from start to end in a flexible and programmatic way.

In this dissertation, we present modeling and optimization techniques to address these limitations. In particular, we focus on optimization-based strategies for demonstrably improving the resilience of power systems to extreme weather events through proactive planning and investment. We also present an open source Python package, WNTR, designed to enable water utilities and researches the ability to analyze the resilience of water distribution systems.

1.4 Thesis Outline

This dissertation is divided into two parts. In Part I, we present novel stochastic optimization models to improve power systems resilience to extreme weather events. We begin in Chapter 2 by reviewing power systems models. In Chapter 3, we consider proactive redispatch, transmission line hardening, and transmission line capacity increases as alternatives for mitigating expected load shed due to extreme weather. Our model is based on linearized or "DC" optimal power flow, similar to models in widespread use by independent system operators (ISOs) and regional transmission operators (RTOs). Our computational experiments indicate that proactive redispatch alone can reduce the expected load shed by as much as 25% relative to standard economic dispatch. This resiliency enhancement strategy requires no capital investments and is implementable by ISOs and RTOs solely through operational adjustments. We additionally demonstrate that transmission line hardening and increases in transmission capacity can, in limited quantities, be effective strategies to further enhance power grid resiliency, although at significant capital investment cost. We perform a cross validation analysis to demonstrate the robustness of proposed recommendations.

Our proposed model can be augmented to incorporate a variety of other operational and investment resilience strategies, or combination of such strategies.

We then extend the resilience formulations to investigate the role chemical process facilities, as industrial energy consumers, can play in improving electric grid resilience through demand response. While peak shaving is commonly used to reduce power costs, chemical process facilities that can reduce power consumption on demand during emergencies (e.g., extreme weather events) bring additional value through improved resilience. For process facilities to effectively negotiate demand response (DR) contracts and make investment decisions regarding flexibility, they need to quantify their additional value to the grid. We present a grid-centric mixed-integer stochastic programming framework to determine the value of DR for improving grid resilience in place of capital investments that can be cost prohibitive for system operators.

In Chapter 4, we reformulate the DR problems using the nonlinear alternating current power flow model to investigate the effect of the linear DC approximation. Our numerical results with both models demonstrate that DR can be used to reduce the capital investment necessary for resilience, increasing the value that chemical process facilities bring through DR. However, the linearized model often underestimates the amount of DR needed in our case studies, motivating scalable solution algorithms for Mixed-Integer Nonlinear Programming (MINLP) problems in power systems.

An important step in many MINLP algorithms is the global solution of a Nonlinear Programming (NLP) subproblem. For power systems applications, this involves global solution of NLP's containing the alternating current (AC) power flow model. In Chapter 5, we show that a strong upper bound on the objective of the alternating current optimal power flow (ACOPF) problem can significantly improve the effectiveness of optimization-based bounds tightening (OBBT) on a number of relaxations. We additionally compare the performance of relaxations of the ACOPF problem, including the rectangular form without reference bus constraints, the rectangular form with reference bus constraints, and the polar form. We find that relaxations of the rectangular form significantly strengthen existing relaxations if reference bus constraints

are included. Overall, relaxations of the polar form perform the best. However, neither the rectangular nor the polar form dominates the other. Ultimately, with these strategies, we are able to reduce the optimality gap to less than 0.1% on all but 5 NESTA test cases with up to 300 buses by performing OBBT alone.

The primary disadvantage to OBBT is that it is computationally expensive, especially for large problems. To partially mitigate this computational burden, we propose a decomposition algorithm based on graph partitioning in Chapter 6. The algorithm begins by forming a graph of the problem from the variables (nodes) and constraints (edges). We then partition the graph into two graphs with roughly equal numbers of nodes (variables). This is done recursively with the resulting graphs until the final “leaf” problems are sufficiently small. Each graph produced in this process corresponds to a set of variables and constraints. Many small bounds tightening problems may then be solved with the individual (and small) leaf problems efficiently. In order to retain feasibility information from the constraints contained in other graphs, a few large bounds tightening problems are solved with the larger graphs on variables in the constraints that were removed in order to partition the graph. The algorithm is tested on several instances from version 17.08 of the Power Grid Lib - Optimal Power Flow repository [IEEE PES Task Force on Benchmarks for Validation of Emerging Power System Algorithms, 2017]. Our numerical results demonstrate that the decomposition algorithm produces bounds nearly as tight as those obtained with traditional OBBT but with drastically improved computational performance.

In Part II, we describe the Water Network Tool for Resilience (WNTR), an new open source Python package designed to help water utilities investigate resilience of water distribution systems to hazards and evaluate resilience-enhancing actions. The WNTR modeling framework is presented and a case study is described that uses WNTR to simulate the effects of an earthquake on a water distribution system. The case study illustrates that the severity of damage is not only a function of system integrity and earthquake magnitude, but also of the available resources and repair strategies used to return the system to normal operating conditions. While

earthquakes are particularly concerning since buried water distribution pipelines are highly susceptible to damage, the software framework can be applied to other types of hazards, including power outages and contamination incidents.

Part I

Models and Algorithms for Improving Power Systems Resilience

2. MODELING POWER SYSTEMS

In this Chapter, we briefly review power systems modelling. We first review the alternating current power flow (ACPF) equations and then describe a common linear approximation, typically referred to as the DC power flow (DCPF) equations.

2.1 Alternating Current Power Flow Equations

An electric grid may be modeled as a graph of nodes and edges where the nodes represent buses (e.g., generating stations, substations) and edges represent transmission elements (e.g., lines, transformers). Generators, which may be found at one or more buses, induce sinusoidal voltage waveforms in the network (i.e., the voltage is not constant). The voltages induced by generators produce current waveforms so that power may be transmitted by transmission lines from the generators to energy consumers. At steady state, these waveforms become phasors (i.e., the amplitude and frequency are constant). Thus, the voltage, current, and power phasors may be modeled with complex numbers with either real and imaginary components or magnitudes and angles.

Power flow through a grid is governed by Ohm's Law for AC circuits, Joule's First Law, and Kirchhoff's circuit laws. Kirchhoff's Current Law (KCL) states that the total current entering a bus must equal the total current leaving the bus. Kirchhoff's Voltage Law (KVL) states that the voltage differences around any cycle in the network must sum to zero. Joule's First Law relates the power flow on a transmission line to the current and voltage. The extension of Ohm's Law to AC systems relates the current on a transmission line to the voltage difference along the line and the admittance of the line. The admittance of a transmission line is a complex physical property of the line with real and imaginary components, referred to as the conduc-

tance and susceptance, respectively. The conductance represents the ease at which electric current passes and is determined as:

$$G_l = R_l / (R_l^2 + X_l^2) \quad (2.1)$$

where R_l is the line resistance and X_l is the line reactance. The susceptance represents the ease at which electric current changes and is determined as:

$$B_l = -X_l / (R_l^2 + X_l^2). \quad (2.2)$$

The physical properties of the line may also be expressed in terms of the impedance, which is the inverse of the admittance. Impedance has real and imaginary components resistance and reactance.

With these definitions, the ACPF equations describing power flows on transmission line l from bus b toward bus n are given by,

$$p_{l,b,n} = C_{l,b,n}^{\alpha,p} v_b^2 + C_{l,b,n}^{\beta,p} v_b v_n \cos(\theta_b - \theta_n) + C_{l,b,n}^{\gamma,p} v_b v_n \sin(\theta_b - \theta_n) \quad (2.3.1)$$

$$q_{l,b,n} = C_{l,b,n}^{\alpha,q} v_b^2 + C_{l,b,n}^{\beta,q} v_b v_n \cos(\theta_b - \theta_n) + C_{l,b,n}^{\gamma,q} v_b v_n \sin(\theta_b - \theta_n) \quad (2.3.2)$$

where p is the real component of power flow, q is the imaginary (or reactive) component of power flow, v is the voltage magnitude, and θ is the voltage angle. Let T_l and Θ_l^{shift} be the transformer tap ratio for branch l and the transformer phase shift for branch l , respectively. The coefficients C are

$$C_{l,b,n}^{\alpha,p} \equiv \begin{cases} \frac{G_l}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } b \\ G_l & \text{otherwise} \end{cases} \quad (2.4.1)$$

$$C_{l,b,n}^{\beta,p} \equiv \begin{cases} \frac{B_l T_l \sin(\Theta_l^{\text{shift}}) - G_l T_l \cos(\Theta_l^{\text{shift}})}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } b \\ \frac{-B_l T_l \sin(\Theta_l^{\text{shift}}) - G_l T_l \cos(\Theta_l^{\text{shift}})}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } n \\ -G_l & \text{otherwise} \end{cases} \quad (2.4.2)$$

$$C_{l,b,n}^{\gamma,p} \equiv \begin{cases} \frac{-B_l T_l \cos(\Theta_l^{\text{shift}}) - G_l T_l \sin(\Theta_l^{\text{shift}})}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } b \\ \frac{-B_l T_l \cos(\Theta_l^{\text{shift}}) + G_l T_l \sin(\Theta_l^{\text{shift}})}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } n \\ -B_l & \text{otherwise} \end{cases} \quad (2.4.3)$$

$$C_{l,b,n}^{\alpha,q} \equiv \begin{cases} \frac{-B_l - 0.5B_l^c}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } b \\ -B_l - 0.5B_l^c & \text{otherwise} \end{cases} \quad (2.4.4)$$

$$C_{l,b,n}^{\beta,q} \equiv \begin{cases} \frac{B_l T_l \cos(\Theta_l^{\text{shift}}) + G_l T_l \sin(\Theta_l^{\text{shift}})}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } b \\ \frac{B_l T_l \cos(\Theta_l^{\text{shift}}) - G_l T_l \sin(\Theta_l^{\text{shift}})}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } n \\ B_l & \text{otherwise} \end{cases} \quad (2.4.5)$$

$$C_{l,b,n}^{\gamma,q} \equiv \begin{cases} \frac{B_l T_l \sin(\Theta_l^{\text{shift}}) - G_l T_l \cos(\Theta_l^{\text{shift}})}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } b \\ \frac{-B_l T_l \sin(\Theta_l^{\text{shift}}) - G_l T_l \cos(\Theta_l^{\text{shift}})}{(T_l)^2} & \text{if there is a transformer on branch } l \text{ by bus } n \\ -G_l & \text{otherwise} \end{cases} \quad (2.4.6)$$

$$C_{l,b,n}^T \equiv \begin{cases} -\Theta_l^{\text{shift}} & \text{if there is a transformer on branch } l \text{ by bus } b \\ \Theta_l^{\text{shift}} & \text{otherwise} \end{cases} \quad (2.5)$$

Note that Equations (2.3) could also be written in terms of the real and imaginary components of the voltages. More detailed reviews of power flow formulations may be found in the literature [Zimmerman et al., 2011, Cain et al., 2012].

2.2 The DC Approximation

A linear approximation of the ACPF equations (Equations (2.3)), referred to as the DC approximation, is commonly used. This linearization assumes that the resistance of each transmission line is much less than its reactance, the voltage magnitude at each bus is close to nominal (i.e., 1 in the per unit system), and the voltage angle difference between interconnected buses is small [Zimmerman and Murillo-Sanchez, 2015]. These assumptions result in the following linear angle-to-power relationship to define the real power flow on line l :

$$p_l = \frac{1}{X_l T_l} (\theta_b - \theta_n - \Theta_l^{\text{shift}}) \quad (2.6)$$

where X_l , T_l , and Θ_l^{shift} are the reactance, transformer tap ratio, and transformer phase shift of branch l , respectively, and θ_b and θ_n are the voltage angles at the “from” and “to” buses of branch l , respectively. The real power flow p_l is limited by

$$-S_l^{\max} \leq p_l \leq S_l^{\max} \quad (2.7)$$

where S_l^{\max} is the transfer capacity of line l . Additionally, the power transfer between interconnected buses b and n can be limited by the maximum voltage angle difference $\Theta_{(b,n)}^{\max}$ in

$$-\Theta_{(b,n)}^{\max} \leq \theta_b - \theta_n \leq \Theta_{(b,n)}^{\max} \quad (2.8)$$

Buses may have shunts, loads, and/or generators. An energy balance at each bus b requires

$$\sum_{g \in \mathcal{G}_b} p_g^G + \sum_{l \in \mathcal{L}_b^{in}} p_l - \sum_{l \in \mathcal{L}_b^{out}} p_l = P_b^L + G_b^S \quad (2.9)$$

where \mathcal{G}_b is the set of generators at bus b , \mathcal{L}_b^{in} is the set of all transmission lines with bus b as its “to” bus, \mathcal{L}_b^{out} is the set of all transmission lines with bus b as its “from” bus, p_g^G is the real component of the power generated at generator g , P_b^L is the real component of the power demand at bus b , p_l is the real component of the power flowing along transmission line l , and G_b^S is the shunt conductance at bus b .

3. STOCHASTIC MIXED INTEGER LINEAR PROGRAMS FOR IMPROVING POWER SYSTEMS RESILIENCE TO EXTREME WEATHER¹

Despite ongoing investments in grid modernization, the U.S. electric grid remains vulnerable to a range of events [U.S. Department of Energy, 2015]. More recently, disruptions to the U.S. electric grid have resulted in an estimated annual cost of \$18-70 billion in security, health and safety, and economic consequences [Executive Office of the President, 2013]. Natural disasters, although infrequent, can have major impacts on the electric grid through direct damage of infrastructure, or indirectly through interdependent infrastructures such as gas and water. Also, the increasing reliance on renewable generation contributes to the increasing uncertainty on the “threats to, vulnerabilities of, and potential consequences from all hazards on critical infrastructures” [The White House, 2013]. Depending on the geographical location, such natural disasters include tornadoes, wildfires, hurricanes, and earthquakes, and can result in significant infrastructure damage. The impact of such disasters has demonstrated the need for resiliency, which is the ability to harden the system against – and quickly recover from – low-frequency, high-consequence events.

To improve power systems resilience to extreme weather events, we propose a novel optimization model that explicitly considers a range of possible component outages due to extreme weather events. Specifically, we introduce a two-stage stochastic MILP optimization model – based on linearized or “DC” power flow, of the type widely used in power systems operations – whose solution determines optimal dispatch and / or investment planning in the first stage (pre-storm) to minimize the expected load shed due to system component damage from different realizations of a second-stage

¹Part of this chapter is reprinted from “Evaluating Demand Response Opportunities for Power Systems Resilience Using MILP and MINLP Formulations” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., to appear in AIChE Journal, 2018.

extreme weather event. We demonstrate the utility of our stochastic MILP model on both small (30 bus) and large (2383 bus) benchmark power systems. We also perform a cross validation analysis that demonstrates our model can yield a near-optimal solution (with respect to out-of-sample performance) considering less than 0.001% of all possible outage realizations.

The remainder of this chapter is organized as follows. In the following section, we present our stochastic MILP model for improving power system resiliency to sampled extreme weather outage scenarios by optimally dispatching generators proactively. We then describe, in Section 3.2, strategies for generating extreme weather outage realizations (i.e., scenarios) and briefly discuss data challenges associated with these strategies. Sections 3.3 and 3.4 detail extensions for long term planning and computational results obtained with our stochastic MILP model considering both small and large benchmark power systems. In Section 3.5, we present cross validation results illustrating the robustness of our solutions, which can be obtained using only a small fraction of all probable component outage scenarios. Finally, in Section 3.6, we extend the formulation to consider the role large industrial consumers may play in electric grid resilience through demand response.

3.1 A Resilient Dispatch Formulation

We now present our two-stage stochastic LP model to optimally dispatch generators in preparation for extreme weather events. The first decision stage of the model captures steady state power system operations prior to the event. The second decision stage represents post-event steady-state operation following a number of component outages due to the event. For simplicity, we only consider transmission line outages; the model can be generalized to consider bus and generator outages as well. Any first stage (proactive dispatch) decisions must be made before any uncertainty is revealed, and are therefore *non-anticipative*. Uncertainty is revealed between decision stages one and two. Decisions (reactive dispatch and load shed) in stage two are known as *recourse* decisions, as they are made following realization of any potential uncertain-

ties – in our case, component outages. As described previously, this uncertainty is captured via a finite number of discrete scenarios, each with a (potentially distinct) probability of occurrence.

Generator ramping rates are the key linkage between the first and second decision stages; operational characteristics of thermal generators dictate that they can only feasibly change power output levels at a limited rate. In this context, the optimization objective is to identify a first stage dispatch that minimizes the expected load-shed across the set of extreme event scenarios \mathcal{S} . The second stage decisions can be viewed loosely as a "play book" – if a scenario $s \in \mathcal{S}$ is realized, then the optimal action to minimize load shed is to re-dispatch according to the recourse p_g^G values.

Our model, which we refer to as the Resilient Dispatch Formulation (RDF), is

$$\min \sum_{s \in \mathcal{S}} \Omega_s \left(\sum_{b \in \mathcal{B}} p_{b,s}^{L,shed} + \alpha \sum_{g \in \mathcal{G}} p_{g,s}^{G,cur} \right) \quad (3.1.1)$$

s.t.

$$\sum_{g \in \mathcal{G}_b} p_{g,0}^G + \sum_{k \in \mathcal{K}_b^{\text{in}}} p_{k,0} - \sum_{k \in \mathcal{K}_b^{\text{out}}} p_{k,0} = P_b^L + G_b^S \quad \forall b \in \mathcal{B} \quad (3.1.2)$$

$$P_g^{G,\min} \leq p_{g,0}^G \leq P_g^{G,\max} \quad \forall g \in \mathcal{G} \quad (3.1.3)$$

$$p_{k,0} = \frac{1}{X_k T_k} (\theta_{b,0} - \theta_{n,0} - \Theta_k^{\text{shift}}) \quad \forall (k, b, n) \in \mathcal{L} \quad (3.1.4)$$

$$-S_k^{\max} \leq p_{k,0} \leq S_k^{\max} \quad \forall k \in \mathcal{K} \quad (3.1.5)$$

$$-\Theta_{(b,n)}^{\max} \leq \theta_{b,0} - \theta_{n,0} \leq \Theta_{(b,n)}^{\max} \quad \forall (b, n) \in \mathcal{A} \quad (3.1.6)$$

$$-P_g^{G,R} \leq p_{g,0}^G - p_{g,s}^G \leq P_g^{G,R} \quad \forall g \in \mathcal{G} \quad \forall s \in \mathcal{S} \quad (3.1.7)$$

$$\theta_{\text{ref},0} = 0 \quad (3.1.8)$$

$$\begin{aligned} \sum_{g \in \mathcal{G}_b} (p_{g,s}^G - p_{g,s}^{G,cur}) + \sum_{k \in \mathcal{K}_b^{\text{in}}} p_{k,s} - \sum_{k \in \mathcal{K}_b^{\text{out}}} p_{k,s} \\ = P_b^L - p_{b,s}^{L,shed} + G_b^S \end{aligned} \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (3.1.9)$$

$$P_g^{G,\min} \leq p_{g,s}^G \leq P_g^{G,\max} \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S} \quad (3.1.10)$$

$$p_{k,s} = \frac{1}{X_k T_k} (\theta_{b,s} - \theta_{n,s} - \Theta_k^{\text{shift}}) \quad \forall (k, b, n) \in \mathcal{L} \setminus \mathcal{L}_s, \forall s \in \mathcal{S} \quad (3.1.11)$$

$$p_{k,s} = 0 \quad \forall k \in \mathcal{K}_s, \forall s \in \mathcal{S} \quad (3.1.12)$$

$$-S_k^{\max} \leq p_{k,s} \leq S_k^{\max} \quad \forall k \in \mathcal{K}, \forall s \in \mathcal{S} \quad (3.1.13)$$

$$-\Theta_k^{\max} \leq \theta_{b,s} - \theta_{n,s} \leq \Theta_k^{\max} \quad \forall (b,n) \in \mathcal{A}_s, \forall s \in \mathcal{S} \quad (3.1.14)$$

$$\theta_{\text{ref},s} = 0 \quad \forall s \in \mathcal{S} \quad (3.1.15)$$

$$0 \leq p_{b,s}^{L,\text{shed}} \leq P_b^L \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (3.1.16)$$

$$0 \leq p_{g,s}^{G,\text{cur}} \leq p_{g,s}^G \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S} \quad (3.1.17)$$

where constraints (3.1.2) – (3.1.6) are related to first stage decisions and are denoted by the index 0, constraints (3.1.9) – (3.1.17) are related to second stage decisions and are denoted by scenario s , and constraint (3.1.7) enforces the ramping capability for all generators from the nominal operating state in the first stage to the outage state in the second stage.

The objective function, (3.1.1), is to minimize the sum of the load shed and the over-generation. The over-generation term is needed for feasibility (e.g., generators may become islanded). Constraints (3.1.2) and (3.1.9) represent the nodal real power balance at every bus. Constraints (3.1.3) and (3.1.10) limit the real power output at each generator. Constraints (3.1.4) and (3.1.11) determine the real power flow on each line k between interconnected buses b and n , where $\mathcal{L} \setminus \mathcal{L}_s$ denotes all transmission lines that are in service in scenario s . For the transmission line outages in scenario s , constraint (3.1.12) enforces zero real power flow. Constraints (3.1.5) and (3.1.13) limit the real power flow on each transmission line, and constraints (3.1.6) and (3.1.14) limit the voltage angle difference between connected buses. Since an infinite number of solutions exist for equivalent voltage angle differences $\theta_b - \theta_n$ over all $(b,n) \in \mathcal{A}$, constraints (3.1.8) and (3.1.15) set the voltage angle at the reference bus (denoted with subscript “ref”) to determine a single solution. For each adverse weather event scenario s , constraint (3.1.16) limits the nonnegative load shed below the specified load, P_b^L , and constraint (3.1.17) limits the nonnegative curtailment needed below the generator’s operating level, $p_{g,s}^G$. The generator curtailment constraint approximates

the more realistic model in which generators either shut-down or operate at an output between the minimum operating level (MOL), $P_g^{G,\min}$, and its maximum capacity, $P_g^{G,\max}$.

The RDF is applicable for hours to days before a severe weather event to obtain a resilient dispatch for a given weather forecast that is used to construct scenarios (see Section 3.2). The resilient dispatch may be less efficient than the economic dispatch used in standard operating conditions. However, the RDF operates the generators in a state that enables transition to contingency states with as little load shedding and generator curtailment as possible. Additionally, because the problem is entirely continuous and linear, it may be solved efficiently with commercial solvers such as Gurobi [Gurobi Optimization, 2016] and CPLEX [ILOG, 2009] in the short time frame between weather forecasts and actual weather events.

3.2 Scenario Construction

We construct a set of relevant probabilistic scenarios that capture transmission line outages (e.g., due to damage or de-energization) associated with extreme weather event(s) of interest. These scenarios are used as a key input to our stochastic MILP optimization model, and explicitly represent uncertainty regarding future possible realizations of component damage associated with adverse weather conditions. We note that these scenarios can be associated with a specific operational event (e.g., an oncoming hurricane) or a set of adverse events common to a specific region (e.g., ice storms and high winds), when considering system planning.

Ideally, such scenarios would be obtained from statistical models representing system component outages, constructed (e.g.,) using historical component failure data. For the synthetic test systems we consider here, we instead construct scenarios by first sampling the number of transmission line outages from a negative binomial distribution, and then determine which lines were not in service by sampling from a uniform distribution.

3.3 Extensions for Long Term Investments

The RDF can be extended to consider many possible long term planning investments for improved grid resilience. Possible investments for improving resilience include hardening transmission lines, elevating substations, increasing transmission capacity, placing switches for transmission switching, and many others. All of these could be modeled within a stochastic programming framework. We demonstrate the extensibility of the RDF with the following two examples: (1) transmission line hardening and (2) increasing transmission capacity.

3.3.1 Transmission Line Hardening

Transmission lines can be hardened to severe weather. A few possible hardening strategies include burying transmission lines underground, installing guy wires, and upgrading crossarm materials [Wang et al., 2016b]. However, it is unreasonable to harden every component in the network. Therefore, it is important to understand the priority of lines to be hardened and the impact of this investment on grid resiliency.

Therefore, we augment the RDF in (3.1) to optimally determine the set of branches to harden for a given budget. We introduce a binary variable, δ_k , to indicate whether a transmission line is hardened ($\delta_k = 1$) or not ($\delta_k = 0$). The assumption is that once a line is hardened, it becomes invulnerable to the current weather event. We replace constraints (3.1.11) – (3.1.14) as follows:

$$p_{k,s} = \frac{1}{X_k T_k} (\theta_{b,s} - \theta_{n,s} - \Theta_k^{\text{shift}}) \quad \forall (k, b, n) \in \mathcal{L} \setminus \mathcal{L}_s, \forall s \in \mathcal{S} \quad (3.2.1)$$

$$\begin{aligned} -M(1 - \delta_k) &\leq p_{k,s} X_k T_k - \\ &(\theta_{b,s} - \theta_{n,s} - \Theta_k^{\text{shift}}) \leq M(1 - \delta_k) \end{aligned} \quad \forall (k, b, n) \in \mathcal{L}_s, \forall s \in \mathcal{S} \quad (3.2.2)$$

$$-S_k^{\max} \leq p_{k,s} \leq S_k^{\max} \quad \forall k \in \mathcal{K} \setminus \mathcal{K}_s, \forall s \in \mathcal{S} \quad (3.2.3)$$

$$-S_k^{\max} \delta_k \leq p_{k,s} \leq S_k^{\max} \delta_k \quad \forall k \in \mathcal{K}_s, \forall s \in \mathcal{S} \quad (3.2.4)$$

$$-\Theta_{(b,n)}^{\max} \leq \theta_{b,s} - \theta_{n,s} \leq \Theta_{(b,n)}^{\max} \quad \forall (b, n) \in \mathcal{A}_s, \forall s \in \mathcal{S} \quad (3.2.5)$$

$$-\Theta_{(b,n)}^{\max} - M(1 - \delta_k) \leq \theta_{b,s} - \theta_{n,s} \quad \forall (k, b, n) \in \mathcal{L}_s, \forall s \in \mathcal{S} \quad (3.2.6)$$

$$\theta_{b,s} - \theta_{n,s} \leq \Theta_{(b,n)}^{\max} + M(1 - \delta_k) \quad \forall (k, b, n) \in \mathcal{L}_s, \forall s \in \mathcal{S} \quad (3.2.7)$$

$$\sum_{k \in \mathcal{K}} \delta_k \leq C \quad (3.2.8)$$

$$\delta_k \in \{0, 1\} \quad \forall k \in \mathcal{K} \quad (3.2.9)$$

Constraint (3.2.1) is the real power flow for in-service lines, and constraint (3.2.2) enables real power flow for transmission line outages that have been hardened using Big-M notation. Constraint (3.2.4) enforces the transfer capacity limits on hardened lines, and otherwise prevents power flow on lines that remain in an out of service state. In other words, constraint (3.2.2) results in an unconstrained voltage angle difference when constraint (3.2.4) is inactive. Then constraint (3.2.5) enforces the voltage angle difference limits between connected buses, and constraints (3.2.6) - (3.2.7) only enforce this limit when there exists a hardened line out of the set of outages for interconnections between buses b and n in scenario s . Constraint (3.2.8) limits the number of hardened transmission lines to C . Constraint (3.2.9) defines the first-stage binary variable δ_k .

3.3.2 Increasing Transmission Capacity

An alternative hardening investment is to increase the pre-existing transmission capacity. In this case, we introduce a binary variable, γ_k , to represent whether the capacity of the corresponding transmission line has been increased ($\gamma_k = 1$) or not ($\gamma_k = 0$). Constraint (3.1.13) is replaced by the following:

$$-S_k^{\max}(1 - \gamma_k) - \Lambda S_k^{\max} \gamma_k \leq p_{s,k} \quad \forall k \in \mathcal{K} \quad \forall s \in \mathcal{S} \quad (3.3.1)$$

$$p_{s,k} \leq S_k^{\max}(1 - \gamma_k) + \Lambda S_k^{\max} \gamma_k \quad \forall k \in \mathcal{K} \quad \forall s \in \mathcal{S} \quad (3.3.2)$$

$$\sum_{k \in \mathcal{K}} \gamma_k \leq C \quad (3.3.3)$$

$$\gamma_k \in \{0, 1\} \quad \forall k \in \mathcal{K} \quad (3.3.4)$$

where Λ is the factor by which the transmission capacity will be increased. Constraints (3.3.1)-(3.3.2) enforce the nominal capacity limits when $\gamma_k = 0$, and increases the limits by a factor of Λ when $\gamma_k = 1$. Constraint (3.3.3) limits the number of transmission lines for which the thermal limits can be increased. Constraint (3.3.4) defines the first-stage binary variable γ_k .

3.4 Computational Results

We present results for each of the stochastic programming formulations (RDF, THF, and TCP) for two standard test cases from the publicly available Matpower software package for power flow analysis: case30 and case2383wp [Zimmerman and Murillo-Sanchez, 2015]. The network sizes for the test cases are summarized in Table 3.1. One hundred scenarios were used for case30, and fifty scenarios were used for

Table 3.1.: Summary of network sizes

	case30	case2383wp
# of Buses	30	2383
# of Generators	6	327
# of Transmission Lines	41	2896

case2383wp. The number of transmission line outages in each scenario was drawn from a negative binomial distribution, and the transmission lines that were out of service were drawn from a uniform distribution. The resulting problem sizes for the Resilient Dispatch Formulation (RDF), Transmission Hardening Formulation (THF), and Transmission Capacity Formulation (TCF) are given in Table 3.2.

The stochastic programming formulations were modeled with Pyomo [Hart et al., 2017] and solved with CPLEX [ILOG, 2009]. The THF and TCF formulations for case30 were solved to less than a 0.01% optimality gap. The THF problems for case2383wp were solved to less than a 2% gap, and the TCF formulations for case2383wp were solved to less than a 4% gap. The weight α on the generator curtailment term in the objective was 0.01. The value of Λ used in the TCF formulation

Table 3.2.: Problem size statistics

Test Case	Formulation	Continuous Variables	Binary Variables	Equality Constraints	Inequality Constraints
case30	RDF	9,677	0	6,662	8,664
case30	THF	14,579	41	10,762	11,613
case30	TCF	13,777	41	10,762	10,189
case2383wp	RDF	390,755	0	267,767	341,639
case2383wp	THF	537,401	2,896	412,567	347,950
case2383wp	TCF	535,555	2,896	412,567	344,820

was 2. Due to a lack of ramp rate data for the synthetic test cases, a correlation was developed between the maximum power output of the generator and the ramp rate using real proprietary utility data. The correlation was then used to estimate ramp rates for these cases. A conservative 5-minute ramp rate was used in order to limit the overall change of the state of the grid during the severe weather event because the system dynamics were not modeled.

The results are presented in Figures 3.1 and 3.2; incremental variants variants are shown in Figures 3.3 and 3.4. The left and right panels illustrate results for the THF and TCF formulations, respectively. The y-axis indicates the expected load shed $\mathbb{E}(p^{L,shed})$ across all scenarios, i.e.,

$$\mathbb{E}(p^{L,shed}) = \sum_{s \in \mathcal{S}} \Omega_s \sum_{b \in \mathcal{B}} p_{b,s}^{L,shed}$$

in the objective function (3.1.1). The asterisk represents the *baseline*, which is the case where no preparation is done before the severe weather event. Rather, the generator outputs in the first stage are fixed (as determined by economic dispatch), and the expected load shed is simply the evaluation of the effect of each scenario on the system. Note that when the number of lines hardened is zero, this corresponds

to the RDF solution (redispatch only). The same is true for the TCF figure (i.e., optimal value at zero is the RDF solution).

The RDF results in Figures 3.3 and 3.4 indicate that simply redispatching generators before a severe weather event reduces the expected load shed by 25.0% and 16.6% for case30 and case2383wp, respectively. We note that the incremental reduction in expected load shed is calculated as the percentage with respect to the baseline load shed. This substantial reduction in load shed does not require any capital investments and therefore is immediately implementable by the system operator through operational adjustments.

The THF results indicate that line hardening can be very effective, with higher marginal gains from the initial investments. The left panels of Figures 3.3 and 3.4 show the incremental reduction in expected load shed decreasing with increasing number of lines hardened. For case30, the first two transmission lines hardened each reduce the expected load shed by over 13%. After the first two, there is a significant drop to about 7% followed by a continued gradual decline. Then for case2383wp, hardening just one transmission line reduces the expected load shed by 10%, followed by a significant decrease in which each transmission line hardened after the fifth decreases the expected load shed by less than 1.5%.

The TCF results indicate a slightly less effective strategy, as illustrated in the right panels of Figures 3.3 and 3.4. Our results for case30 indicate that doubling the transmission capacity of three transmission lines results in over 9% total reduction in expected load shed. For case2383wp, doubling the transmission capacity of five transmission lines results in approximately a 4.4% reduction in expected load shed. Increasing the capacity of transmission lines is clearly less effective than hardening per transmission line. Depending on the cost of increasing the transmission capacity of each of these lines, this could be combined with RDF and THF strategies for improved resilience.




Figure 3.1.: The left and right panels illustrate results for THF and TCF, respectively, for case30. The y-axis indicates the expected load shed across all scenarios. The asterisk represents the *baseline*, which is the case where no preparation is done before the severe weather event. The points at zero hardened lines (or zero transmission lines with increased capacity) correspond to the RDF solution.




Figure 3.2.: The left and right panels illustrate results for THF and TCF, respectively, for case2383wp. The y-axis indicates the expected load shed across all scenarios. The asterisk represents the *baseline*, which is the case where no preparation is done before the severe weather event. The points at zero hardened lines (or zero transmission lines with increased capacity) correspond to the RDF solution.

3.5 Out-of-Sample Cross Validation

Even for a small test network such as case30, including all possible scenarios in the stochastic programming formulation is intractable. On average, the scenarios for case30 contained 7 transmission line outages. Over 22 million scenarios would be necessary to consider every combination of 7 outages. To demonstrate that only a




Figure 3.3.: Incremental reduction in expected load shed for each additional hardened line (left) and each additional line with increased capacity (right) for case30.




Figure 3.4.: Incremental reduction in expected load shed for each additional hardened line (left) and each additional line with increased capacity (right) for case2383wp.

small fraction of all possible scenarios is needed to obtain a high quality solution, a cross validation case study was performed with case30.

A set of 10 randomly sampled scenarios (different from the one hundred scenarios used above) was used to solve the RDF and THF problems. The solution was then used to evaluate the performance of the first stage solution on the original 100 scenarios used in Section 3.4. This process was then repeated with 30 and then 100 randomly sampled scenarios. The results are shown in Figure 3.5. Both the line labeled “True Scenarios” and the line in Figure 3.1 detail results when the stochastic programming problems are solved with the original 100 scenarios and the load shed

in the second stage is evaluated on the same 100 scenarios. This quantity represents the best possible performance.

Figure 3.5 shows that when only 10 scenarios are used to determine how to dispatch the generators in the first stage, the performance of the original 100 scenarios in the second stage is actually worse than the baseline. This is still true even after four transmission lines have been hardened, illustrating just how critical scenario selection is. The 30 scenario case performs significantly better, and the 100 scenario case performs almost as well as using the “True Scenarios.” Thus, we obtained a high quality solution with less than 0.001% of all possible scenarios.

3.6 The Role of Energy Consumers

Resilience integrates risk management requirements before, during, and after an event [Preston et al., 2016]. Several papers have investigated various investment and recovery strategies for improving electric grid resilience [Hoffman et al., 2010, Binstock and Mattia, 2007, Carrión et al., 2007, Adibi and Fink, 1994a, Lindemeyer et al., 2001, Castillo, 2014]. However, these actions alone may be insufficient or prohibitively expensive. More recently, utilities have leveraged DR resources to address many electric grid needs including improved system reliability, economic dispatch, wholesale price mitigation, and ancillary services [Henrikson and Brief, 2008]. Therefore, DR can be utilized to defer or offset the need for generation, transmission, and distribution infrastructure investments to improve resilience. In emergency events, system operators can offer pricing benefits to industrial-scale customers if they can absorb loss or reduction of power. Such interruptible load management can improve the reliability of the electric grid by not treating load at each bus as a fixed quantity but a decision variable in the system operator’s optimization routine.

Chemical process facilities, as large industrial electricity consumers, have an opportunity to serve as active participants to improve system resilience through DR. This can reduce the likelihood of system blackout (which could significantly impact the process facility). Furthermore, DR can bring additional value since it may allow




Figure 3.5.: Cross validation results for case30. The line labeled “True Scenarios” shows the expected load shed when the solution to the stochastic programming problem is evaluated on the same 100 scenarios used to solve the stochastic programming problem. The remaining lines show the expected load shed when the stochastic programming problem is solved with a different set of scenarios than those used to evaluate the expected load shed.

the grid operator to forgo significant capital investment in the grid to meet resilience targets. This presents an economic opportunity for process facilities if they can overcome the operational challenges associated with increased flexibility. The additional value they bring needs to be understood and quantified by the process facility operator so they can effectively negotiate DR contracts and make investment decisions regarding potential retrofit and operational changes to support the required flexibility.

Numerous studies have investigated the opportunity for interruptible load during peak operations and for congestion management [Bingyuan et al., 2005, Na et al., 2010]. Aalami et al. [2010] focus on economic incentives for interruptible and curtailable DR programs that incorporate penalty mechanisms for non-compliant participants. Other studies [Mitra et al., 2012, 2013, Zhang et al., 2015a, Ashok and Banerjee, 2001] have formulated detailed industrial customer facility models without

the electric grid physics. These studies generally apply either model-based [Zhang et al., 2009] or data-driven [Zhang et al., 2015b] approaches to define the feasible region for each operating mode as a union of convex subregions in the facility’s production space. For example, Zhang et al. [2016] propose a mode-based plant model formulation for the scheduling of continuous industrial processes to provide interruptible load as operating reserve and solve the problem as a MILP. Vujanic et al. [2012] use robust optimization to operate a batch plant as a DR resource, assuming a fixed capacity of interruptible load. It is clear that there can be financial benefits in operating process facilities to reduce electricity consumption during peak usage periods. However, DR programs also present an opportunity to simultaneously improve grid resilience and provide economic benefit to process facilities that have (or can be retrofitted to provide) the ability to rapidly change electricity consumption in response to requests from power system operators.

In this section, we present a framework to quantify the value of DR for improving system resilience due to weather-related events. Specifically, we formulate a stochastic programming problem to investigate the tradeoff between infrastructure investment (i.e., hardening a line to prevent damage in a weather-related event) and mitigation benefits provided by DR. Weather events are modeled using a set of discrete scenarios with random transmission line outages. The formulations include parameters for the number of allowed demand response buses (selected from a subset of the overall network buses) and the maximum percentage of load reduction. The goal of the optimization problems is to determine the minimum number and selection of hardened lines along with the selection of demand response nodes required to ensure system feasibility over the line-outage scenarios. We demonstrate this framework on several test cases. Our numerical results indicate, for these test cases, that a modest amount of DR can significantly reduce the capital investment required to make the electric grid resilient. This suggests that there may be opportunities for flexible process systems and system operators — i.e., ISO and RTO — or electric utilities to work together toward both improving electric grid resilience and mutual financial benefit.

3.6.1 Problem Formulation

Since the system constraints for ACPF in the transmission network are nonconvex and nonlinear, the resulting MINLP would be too large for efficient solution of realistic grid networks with current off-the-shelf tools. As in the previous sections, we use the DCPF-based model which enables representation of the problem as an MILP, and solutions on large networks with this approximation are tractable. We perform the tradeoff analysis on larger cases using the MILP formulation, and, in the next chapter, compare results from the MILP formulation and the full MINLP formulation on two smaller test cases to assess the impact of the approximation on the conclusions of the analysis.

The stochastic programming formulation based on the linearized DCPF model is shown below in disjunctive form.

$$\min \quad |Y|_{true} \quad (3.4.1)$$

s.t.

$$\sum_{g \in \mathcal{G}_b} p_{s,g}^G - \sum_{\{(l,i,j) \in \mathcal{K} \mid i=b\}} p_{s,l,i,j} = p_{s,b}^L + G_b^{Sh} \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (3.4.2)$$

$$\left[\begin{array}{c} X_{s,l,b,n}^H \\ p_{s,l,b,n} = \frac{1}{X_l^T T_l^m} (\theta_b, s - \theta_n, s + C_{l,b,n}^T) \\ -\Theta_{b,n}^{\max} \leq \theta_{s,b} - \theta_{s,n} \leq \Theta_{b,n}^{\max} \end{array} \right] \vee \left[\begin{array}{c} \neg X_{s,l,b,n}^H \\ p_{s,l,b,n} = 0 \end{array} \right] \quad \forall (l, b, n) \in \mathcal{K}, \forall s \in \mathcal{S} \quad (3.4.3)$$

$$\left[\begin{array}{c} X_{s,b}^D \\ (1 - \Delta_b) P_b^L \leq p_{s,b}^L \leq P_b^L \end{array} \right] \vee \left[\begin{array}{c} \neg X_{s,b}^D \\ p_{s,b}^L = P_b^L \end{array} \right] \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (3.4.4)$$

$$X_{s,l,b,n}^H \text{ is } true \quad \forall \{(l, b, n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, \forall s \in \mathcal{S} \quad (3.4.5)$$

$$Y_l \Leftrightarrow X_{s,l,b,n}^H \quad \forall \{(l, b, n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, \forall s \in \mathcal{S} \quad (3.4.6)$$

$$X_{s,b}^D \text{ is } false, \quad \forall \{b \in \mathcal{B} \mid b \notin B^{DR}\} \quad (3.4.7)$$

$$Z_b \Leftrightarrow X_{s,b}^D \quad \forall b \in \mathcal{B}^{DR}, \forall s \in \mathcal{S} \quad (3.4.8)$$

$$|Z|_{true} \leq N_z \quad (3.4.9)$$

$$P_g^{G,\min} \leq p_{s,g}^G \leq P_g^{G,\max} \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S} \quad (3.4.10)$$

$$-S_l^{\max} \leq p_{s,l,b,n} \leq S_l^{\max} \quad \forall (l, b, n) \in \mathcal{K}, \forall s \in \mathcal{S} \quad (3.4.11)$$

$$\theta_{s,ref} = 0 \quad \forall s \in \mathcal{S} \quad (3.4.12)$$

$$X_{s,l,b,n}^H \in \{true, false\} \quad \forall (l, b, n) \in \mathcal{K}, \forall s \in \mathcal{S} \quad (3.4.13)$$

$$X_{s,b}^D \in \{true, false\} \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (3.4.14)$$

$$Z_b \in \{true, false\} \quad \forall b \in \mathcal{B}^{DR} \quad (3.4.15)$$

$$Y_l \in \{true, false\} \quad \forall l \in \mathcal{L} \quad (3.4.16)$$

The coefficients in the constraints are available from the network data for the particular case study being considered, or computed from equations in Chapter 2. Here Y_l is *true* if line l is selected for hardening. The objective function (3.4.1) seeks to minimize the number of lines selected for hardening, represented by $|Y|_{true}$, which is defined as the count of the number of Y_l variables that are true (i.e., $|\{l \in \mathcal{L} \mid Y_l \text{ is } true\}|$). Equation (3.4.2) is the power balance at the bus nodes, including any connected generators or transmission lines, as well as demand loads and bus shunts.

The set \mathcal{L} refers to all the transmission lines, and \mathcal{L}_s^D is the set of lines that are damaged in scenario s . Unless these lines are selected for hardening, these lines will be removed for scenario s in the model. Following this, the equations in (3.4.3) represent the disjunctions to account for line hardening. The boolean variable $X_{s,l,b,n}^H$ is *true* if line l either was not damaged in scenario s as indicated in (3.4.5), or was damaged in scenario s but selected for hardening, as enforced in (3.4.6). If the line is present in scenario s , then the power flow $p_{s,l,b,n}$ is computed and the phase angle constraints are enforced across the line. Otherwise if the line is not present, then the power flow in the line is forced to be zero (i.e., $p_{s,l,b,n} = 0$) and there is no restriction on phase angle difference bounds between the (unconnected) buses.

The disjunction (3.4.4) along with equations (3.4.7) and (3.4.8) encapsulate the selection of nodes where DR is allowed. Here, Z_b is *false* (no DR) if bus b is not a candidate for DR or not selected for DR, and *true* if bus b is a candidate for DR that is

then selected as an interruptible load. The set $\mathcal{B}^{DR} \subseteq \mathcal{B}$ is the set of candidate buses that can be selected for DR. If the node is selected for DR, then $X_{s,b}^D$ is *true* for the corresponding bus and all scenarios (by equation (3.4.8)), and the system operator is allowed (by disjunction (3.4.4)) to request a load reduction up to a value of $\Delta_b P_b^L$ with $0 \leq \Delta_b \leq 1$. Otherwise, the requested load, P_b^L , must be satisfied. Equation (3.4.9) limits the number of DR nodes selected to be no more than N_z using the same definition for the count as that used in the objective function.

Then equations (3.4.10) and (3.4.11) represent the generator limits and line thermal limits respectively. The line thermal limits can be included for all buses in B since the power flow is forced to be zero if a line is damaged (and not selected for hardening). Equation (3.4.12) fixes the angle of the reference bus to be zero; the selection of reference bus is arbitrary, but necessary to ensure uniqueness of voltage angle.

This optimization problem is converted from disjunctive form to a large-scale MILP using the big-M method for the disjunctions. It can then be used to compute the minimum number of hardened lines required to ensure feasible operation across all weather-event scenarios. By adjusting the values for Δ_b , the allowable fractional reduction in load, and N_z , the allowable number of demand response nodes selected, we can determine the reduction in required hardened lines as a function of these parameters. This enables a tradeoff assessment of the impact of different demand response contracts on the infrastructure investment required to maintain reliability under the scenario set.

3.6.2 Numerical Results

We now evaluate the tradeoff between increased DR capabilities and capital investment in infrastructure for improving resilience to weather-related events. The DCPF-based problem formulation is a stochastic programming problem that seeks to find solutions that provide feasible operation across all of the scenarios. Numerous studies have focused on electric grid resilience to the worst case scenario [Wu

and Conejo, 2017, Salmeron et al., 2004, Alguacil et al., 2014], but this approach is not representative of naturally occurring, severe weather events. Different types of natural disasters have different impact regions, predictability, geographical span, and impact duration [Wang et al., 2016a, Guikema et al., 2006, Hines et al., 2008]. Therefore a range of forecasting models have been proposed in the literature to model grid component damage and outage duration, as well as a range of simulation and optimization models to address service restoration, corrective actions, and preparedness planning. When historical data is available, it can be used to generate representative scenarios regarding transmission element failure due to weather-related events.

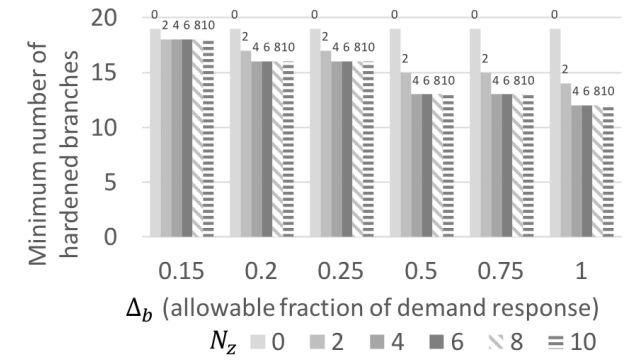
In the absence of system-specific data, we consider failure of transmission lines only, and characterize scenarios according to the set of transmission lines that are damaged (no power transfer occurs across the line - it is essentially removed from the network model). We formulate scenarios by randomly selecting a set of transmission lines to be out of service for each weather-event scenario. Note that the formulation could be adapted to consider failure of other transmission elements, including generators and substations.

We first use the DCPF-based stochastic programming formulation to compute the minimum number of hardened transmission lines that are required to ensure feasible operation across all the scenarios. In this analysis, the maximum required fractional load reduction for DR nodes, Δ_b is varied from 0 to 1, and the number of optimally selected DR node locations N_z is varied from 0 to 10. This provides us with a set of tradeoff curves that provide the minimum required capital investment (in terms of the number of required hardened lines) as a function of the number of demand response contracts and the load reduction fraction. Three test cases are considered from version 17.08 of the Power Grid Lib - Optimal Power Flow repository [IEEE PES Task Force on Benchmarks for Validation of Emerging Power System Algorithms, 2017]: 1) `pglib_opf_case118_ieee`, 2) `pglib_opf_case162_ieee_dtc`, and 3) `pglib_opf_case300_ieee`. For each problem, we considered 100 scenarios with 10 randomly selected branch outages per scenario. Note that branches with only one path between the correspond-

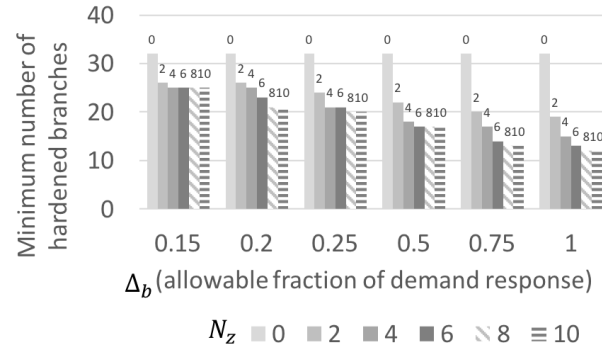
ing pair of buses were excluded since such a branch would likely always need to be hardened. All problem formulations and solution strategies were implemented using Pyomo, a Python-based algebraic modeling language [Hart et al., 2017]. The DC formulation (3.4) is an MILP following big-M transformation of the disjunctions, and the extensive form is solved directly using Gurobi 7.5.2. The results of this trade-off analysis are presented in Figure 3.6. The y-axis shows the minimum number of branches that need to be hardened in order to maintain feasibility. Each group of columns represents a specific value of Δ_b . Within each group of columns, N_z varies from 0 to 10 by 2. The particular value of N_z associated with each column can be seen both in the legend and above the column. The problems were typically solved in under 5 minutes, although most problems took less than 1 minute.

As expected, the minimum number of hardened branches needed to maintain feasibility decreases both as N_z increases and as Δ_b increases. All three test cases show similar trends, indicating that demand response can have a significant impact on resilience. With $\Delta_b=0.25$, only two DR contracts are required to reduce the objective by 11%, 25%, and 21% for cases `pglib_opf_case118_ieee`, `pglib_opf_case162_ieee_dtc`, and `pglib_opf_case300_ieee`, respectively. Increasing Δ_b to 0.5 results in reductions of 21%, 31%, and 25%, respectively. Even more dramatic cost reductions are possible in other cases. For example, with a 50% maximum load reduction, 10 DR contracts can reduce the required number of hardened lines from 32 to 17 (`pglib_opf_case162_ieee_dtc`). If fully interruptable load is a possibility, even more reduction can be realized across these test cases.


The extreme weather events in question typically only occur once every 5-10 years. If chemical, petroleum, and manufacturing plants have the ability to operate flexibly in these rare times, the capital investment needed to make the grid more resilient can be reduced dramatically. As a result, the plants with demand response contracts could gain significant financial benefits in terms of electricity cost reduction with only very rare interruptions in operations. Of course, DR and transmission line hardening are not the only ways of making the electric grid more resilient, and real systems



(a) pglib_opf_case118_ieee



(b) pglib_opf_case162_ieee_dtc



(c) pglib_opf_case300_ieee

Figure 3.6.: The value of DR on resilience as a function of the number of demand response contracts (N_z) and the allowable fraction of DR (Δ_b) with 100 scenarios.

should consider other options along with system specific models and parameters in a more detailed decision making process.

The results we have shown so far only consider the DC model, which is a linear approximation of the true system and does not consider reactive power at all. In the following section, we further validate the value of demand response by both comparing solutions obtained with the DC and AC models and evaluating the the DC solutions on an AC model.

3.7 Summary

In this chapter, we presented stochastic programming formulations for improving power systems resilience to extreme weather events. The key contributions are listed below.

- We presented a stochastic linear program for proactively dispatching generators in order to minimize the expected load shed across a set of probabilistic scenarios. This formulation utilizes operational changes only and can be implemented without any investment cost. For the test cases in this work, redispatching generators proactively can reduce the expected load shed by as much as 25%.
- We presented two extensions to the stochastic programming formulation to consider resilience enhancing investment strategies. We considered both hardening transmission lines and increasing transmission line capacity. Our numerical results demonstrated that both can be very effective, with higher marginal gains from the initial investments. For example, hardening just one transmission line in case2383wp can reduce the expected load shed by over 10% and doubling the capacity of just 5 lines can reduce the expected load shed by approximately 4.4%.
- We performed an out-of-sample cross validation to demonstrate that high quality solutions can be obtained with very small fractions of all possible scenarios.

- We also extended the stochastic programming formulation to consider demand response for improving resilience, demonstrating the value chemical process facilities, as large industrial energy consumers, can bring to electric grid resilience. We found that, for the test cases considered, a modest amount of demand response could significantly reduce the capital investment needed to mitigate the effects of extreme weather events. For example, for the largest test case, the addition of only two demand response contracts (each requiring only a 25% reduction in consumed load) is sufficient to ensure feasible operation across all scenarios while reducing the required number of hardened lines by 21%.

4. THE EFFECT OF THE DC APPROXIMATION¹

The problem formulations presented in Chapter 3 are based on an extensively used linear approximation of the ACPF equations, resulting in a large-scale MILP problem. In this chapter, we investigate the effects of this linear approximation by formulating and solving the demand response problems presented in section 3.6 with the nonconvex ACPF equations. We compare the results obtained with the DC approximation and the full ACPF-based MINLP for a few small test cases. We also evaluate the solutions obtained with the linear DC approximation on the nonlinear AC problem with several larger test cases.

4.1 A Stochastic Mixed Integer Nonlinear Programming Formulation

The formulation below includes the full nonlinear form of the ACPF equations, resulting in a stochastic MINLP. In this formulation, both real and reactive power are represented, and the voltage magnitude is not assumed to be uniform across the network. The following variable transformation, introduced by Jabr [2006], facilitates both the MINLP formulation and its relaxation.

$$\alpha_{s,b} \equiv v_{s,b}^2 \quad (4.1.1)$$

$$\beta_{s,b,n} \equiv v_{s,b}v_{s,n} \cos(\theta_{s,b} - \theta_{s,n}) \quad (4.1.2)$$

$$\gamma_{s,b,n} \equiv v_{s,b}v_{s,n} \sin(\theta_{s,b} - \theta_{s,n}). \quad (4.1.3)$$

¹Part of this chapter is reprinted from “Evaluating Demand Response Opportunities for Power Systems Resilience Using MILP and MINLP Formulations” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., to appear in AICHE Journal, 2018.

With this variable transformation, an MINLP representation of the resilience DR formulation is,

$$\min \quad |Y|_{true} \quad (4.2.1)$$

s.t.

$$\sum_{g \in \mathcal{G}_b} p_{s,g}^G - \sum_{\{(l,i,j) \in \mathcal{K} | i=b\}} p_{s,l,i,j} = p_{s,b}^L + G_b^{Sh} \alpha_{s,b} \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.2.2)$$

$$\sum_{g \in \mathcal{G}_b} q_{s,g}^G - \sum_{\{(l,i,j) \in \mathcal{K} | i=b\}} q_{s,l,i,j} = q_{s,b}^L - B_b^{Sh} \alpha_{s,b} \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.2.3)$$

$$\left[\begin{array}{c} X_{s,l,b,n}^H \\ p_{s,l,b,n} = C_{l,b,n}^{\alpha,p} \alpha_{s,b} + C_{l,b,n}^{\beta,p} \beta_{s,b,n} + C_{l,b,n}^{\gamma,p} \gamma_{s,b,n} \\ q_{s,l,b,n} = C_{l,b,n}^{\alpha,q} \alpha_{s,b} + C_{l,b,n}^{\beta,q} \beta_{s,b,n} + C_{l,b,n}^{\gamma,q} \gamma_{s,b,n} \\ -\tan(\Theta_{b,n}^{\max}) \beta_{s,b,n} \leq \gamma_{s,b,n} \leq \tan(\Theta_{b,n}^{\max}) \beta_{s,b,n} \end{array} \right] \vee \left[\begin{array}{c} \neg X_{s,l,b,n}^H \\ p_{s,l,b,n} = 0 \\ q_{s,l,b,n} = 0 \end{array} \right] \quad \forall (l,b,n) \in \mathcal{K}, \forall s \in \mathcal{S} \quad (4.2.4)$$

$$X_{s,l,b,n}^H \text{ is true} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, \forall s \in \mathcal{S} \quad (4.2.5)$$

$$Y_l \Leftrightarrow X_{s,l,b,n}^H \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, \forall s \in \mathcal{S} \quad (4.2.6)$$

$$\beta_{s,b,n} = \beta_{s,n,b} \quad \forall (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.2.7)$$

$$\gamma_{s,b,n} = -\gamma_{s,n,b} \quad \forall (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.2.8)$$

$$\beta_{s,b,n}^2 + \gamma_{s,b,n}^2 = \alpha_{s,b} \alpha_{s,n} \quad \forall (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.2.9)$$

$$\theta_{s,b} - \theta_{s,n} = \arctan\left(\frac{\gamma_{s,b,n}}{\beta_{s,b,n}}\right) \quad \forall (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.2.10)$$

$$\theta_{s,ref} = 0 \quad \forall s \in \mathcal{S} \quad (4.2.11)$$

$$p_{s,l,b,n}^2 + q_{s,l,b,n}^2 \leq (S_l^{\max})^2 \quad \forall (l,b,n) \in \mathcal{K}, \forall s \in \mathcal{S} \quad (4.2.12)$$

$$(V_b^{\min})^2 \leq \alpha_{s,b} \leq (V_b^{\max})^2 \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.2.13)$$

$$P_g^{G,\min} \leq p_{s,g}^G \leq P_g^{G,\max} \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S} \quad (4.2.14)$$

$$Q_g^{G,\min} \leq q_{s,g}^G \leq Q_g^{G,\max} \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S} \quad (4.2.15)$$

$$\left[\begin{array}{c} X_{s,b}^D \\ (1 - \Delta_b)P_b^L \leq p_{s,b}^L \leq P_b^L \\ (1 - \Delta_b)Q_b^L \leq q_{s,b}^L \leq Q_b^L \end{array} \right] \bigvee \left[\begin{array}{c} \neg X_{s,b}^D \\ p_{s,b}^L = P_b^L \\ q_{s,b}^L = Q_b^L \end{array} \right] \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.2.16)$$

$$X_{s,b}^D \text{ is } false, \quad \forall \{b \in \mathcal{B} \mid b \notin B^{DR}\} \quad (4.2.17)$$

$$Z_b \Leftrightarrow X_{s,b}^D \quad \forall b \in \mathcal{B}^{DR}, \forall s \in \mathcal{S} \quad (4.2.18)$$

$$|Z|_{true} \leq N_z \quad (4.2.19)$$

$$X_{s,l,b,n}^H \in \{true, false\} \quad \forall (l, b, n) \in \mathcal{K}, \forall s \in \mathcal{S} \quad (4.2.20)$$

$$X_{s,b}^D \in \{true, false\} \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.2.21)$$

$$Z_b \in \{true, false\} \quad \forall b \in \mathcal{B}^{DR} \quad (4.2.22)$$

$$Y_l \in \{true, false\} \quad \forall l \in \mathcal{L} \quad (4.2.23)$$

Again the objective function (4.2.1) seeks to minimize the number of lines selected for hardening. Constraints (4.2.2) and (4.2.3) are the real and reactive power balances. If a branch is out of service and not hardened, the real and reactive power flows are forced to zero on the branch by constraints in disjunction (4.2.4). If the branch is hardened or not damaged, then constraint (4.2.4) relates the real and reactive power flows to the new variables $\alpha_{s,b}$, $\beta_{s,b,n}$, and $\gamma_{s,b,n}$ and enforces the angle difference limits through these variables. Equations (4.2.5) and (4.2.6) ensure that the boolean variables $X_{s,l,b,n}^H$ are *true* for all branches without an outage, *true* for all hardened branches, and *false* otherwise. The definitions in (4.1) could be substituted directly into the power flow equations to obtain the standard polar form, however, as discussed in [Kocuk et al., 2016], equations (4.2.7) - (4.2.10) provide an alternate form in the transformed variables. Constraints (4.2.12) - (4.2.15) enforce thermal, voltage, and generator limits, respectively. The disjunction in (4.2.16) is identical to that used in the DC formulation except that it applies to the reactive power load as well. Constraints (4.2.17) - (4.2.23) are also identical to those used in the DC formulation.

4.2 MISOCP Relaxation

The resilience formulation proposed in the previous section is a large-scale stochastic MINLP which is very difficult to solve. We will discuss the algorithm used to solve it in detail in Section 4.3. However, a key component is a convex relaxation of (4.2). There have been a number of relaxations in the literature in the context of ACOPF, including the SOC relaxation [Jabr, 2006, Kocuk et al., 2016], the QC relaxation [Hijazi et al., 2017], and McCormick relaxations of the rectangular form [Bynum et al., 2018b,a]. These relaxations have been used in algorithms for global solution of ACOPF [Liu et al., 2017a, Kocuk et al., 2017b, Lu et al., 2018] and problems including discrete decisions like optimal transmission switching [Kocuk et al., 2017a] and unit commitment [Liu et al., 2018]. The algorithm we use to solve this problem is based on the SOC relaxation, which may be obtained from (4.2) by dropping (4.2.10) - (4.2.11), relaxing (4.2.9) with a convex SOC inequality, and converting the disjunctions to Big-M constraints. The resulting MISOCP relaxation is,

$$\min \sum_{l \in \mathcal{L}} y_l \quad (4.3)$$

s.t.

$$\sum_{g \in \mathcal{G}_b} p_{s,g}^G - \sum_{\{(l,i,j) \in \mathcal{K} \mid i=b\}} p_{s,l,b,n} = p_{s,b}^L + G_b^{Sh} \alpha_{s,b} \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.4)$$

$$\sum_{g \in \mathcal{G}_b} q_{s,g}^G - \sum_{\{(l,i,j) \in \mathcal{K} \mid i=b\}} q_{s,l,b,n} = q_{s,b}^L - B_b^{Sh} \alpha_{s,b} \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.5)$$

$$p_{s,l,b,n} = C_{l,b,n}^{\alpha,p} \zeta_{s,l,b} + C_{l,b,n}^{\beta,p} \eta_{s,l,b,n} + C_{l,b,n}^{\gamma,p} \phi_{s,l,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, \forall s \in \mathcal{S} \quad (4.6)$$

$$q_{s,l,b,n} = C_{l,b,n}^{\alpha,q} \zeta_{s,l,b} + C_{l,b,n}^{\beta,q} \eta_{s,l,b,n} + C_{l,b,n}^{\gamma,q} \phi_{s,l,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, \forall s \in \mathcal{S} \quad (4.7)$$

$$p_{s,l,b,n} = C_{l,b,n}^{\alpha,p} \alpha_{s,b} + C_{l,b,n}^{\beta,p} \beta_{s,b,n} + C_{l,b,n}^{\gamma,p} \gamma_{s,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, \forall s \in \mathcal{S} \quad (4.8)$$

$$q_{s,l,b,n} = C_{l,b,n}^{\alpha,q} \alpha_{s,b} + C_{l,b,n}^{\beta,q} \beta_{s,b,n} + C_{l,b,n}^{\gamma,q} \gamma_{s,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, \forall s \in \mathcal{S} \quad (4.9)$$

$$\underline{\alpha}_{s,b} y_l \leq \zeta_{s,l,b} \leq \overline{\alpha}_{s,b} y_l \quad \forall l \in \mathcal{L}_b \cap \mathcal{L}_s^D, \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.10)$$

$$\overline{\alpha}_{s,b} (y_l - 1) + \alpha_{s,b} \leq \zeta_{s,l,b} \leq \underline{\alpha}_{s,b} (y_l - 1) + \alpha_{s,b} \quad \forall l \in \mathcal{L}_b \cap \mathcal{L}_s^D, \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.11)$$

$$\underline{\beta_{s,b,n}} y_l \leq \eta_{s,l,b,n} \leq \overline{\beta_{s,b,n}} y_l \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.12)$$

$$\overline{\beta_{s,b,n}}(y_l - 1) + \beta_{s,b,n} \leq \eta_{s,l,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.13)$$

$$\eta_{s,l,b,n} \leq \underline{\beta_{s,b,n}}(y_l - 1) + \beta_{s,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.14)$$

$$\underline{\gamma_{s,b,n}} y_l \leq \phi_{s,l,b,n} \leq \overline{\gamma_{s,b,n}} y_l \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.15)$$

$$\overline{\gamma_{s,b,n}}(y_l - 1) + \gamma_{s,b,n} \leq \phi_{s,l,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.16)$$

$$\phi_{s,l,b,n} \leq \underline{\gamma_{s,b,n}}(y_l - 1) + \gamma_{s,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.17)$$

$$\eta_{s,l,b,n} = \eta_{s,l,n,b} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.18)$$

$$\phi_{s,l,b,n} = -\phi_{s,l,n,b} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.19)$$

$$\eta_{s,l,b,n}^2 + \phi_{s,l,n,b}^2 \leq \zeta_{s,l,b} \zeta_{s,l,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.20)$$

$$\beta_{s,b,n} = \beta_{s,n,b} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.21)$$

$$\gamma_{s,b,n} = -\gamma_{s,n,b} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.22)$$

$$\beta_{s,b,n}^2 + \gamma_{s,b,n}^2 \leq \alpha_{s,b} \alpha_{s,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.23)$$

$$p_{s,l,b,n}^2 + q_{s,l,b,n}^2 \leq (S_l^{\max})^2 \quad \forall (l,b,n) \in \mathcal{K}, \forall s \in \mathcal{S} \quad (4.24)$$

$$(V_b^{\min})^2 \leq \alpha_{s,b} \leq (V_b^{\max})^2 \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.25)$$

$$P_g^{G,\min} \leq p_{s,g}^G \leq P_g^{G,\max} \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S} \quad (4.26)$$

$$Q_g^{G,\min} \leq q_{s,g}^G \leq Q_g^{G,\max} \quad \forall g \in \mathcal{G}, \forall s \in \mathcal{S} \quad (4.27)$$

$$-\tan(\Theta_{b,n}^{\max}) \eta_{s,l,b,n} \leq \phi_{s,l,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.28)$$

$$\phi_{s,l,b,n} \leq \tan(\Theta_{b,n}^{\max}) \eta_{s,l,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \in \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.29)$$

$$-\tan(\Theta_{b,n}^{\max}) \beta_{s,b,n} \leq \gamma_{s,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.30)$$

$$\gamma_{s,b,n} \leq \tan(\Theta_{b,n}^{\max}) \beta_{s,b,n} \quad \forall \{(l,b,n) \in \mathcal{K} \mid l \notin \mathcal{L}_s^D\}, (b,n) \in \mathcal{A}, \forall s \in \mathcal{S} \quad (4.31)$$

$$(1 - \Delta_b) P_b^L z_b + P_b^L (1 - z_b) \leq p_{s,b}^L \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.32)$$

$$p_{s,b}^L \leq P_b^L \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.33)$$

$$(1 - \Delta_b) Q_b^L z_b + Q_b^L (1 - z_b) \leq q_{s,b}^L \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.34)$$

$$q_{s,b}^L \leq Q_b^L \quad \forall b \in \mathcal{B}, \forall s \in \mathcal{S} \quad (4.35)$$

$$\sum_{b \in \mathcal{B}} z_b \leq N_z \quad (4.36)$$

$$z_b \in \{0, 1\} \quad \forall b \in \mathcal{B} \quad (4.37)$$

$$y_l \in \{0, 1\} \quad \forall l \in \mathcal{L} \quad (4.38)$$

Rather than applying a Big-M transformation directly to the constraints in (4.2.4), we use a similar approach to that proposed by Kocuk et al. [2017a]. Because we typically have tight bounds on α , β , and γ , we introduce the new variables ζ , η , and ϕ and relate them to α , β , and γ through constraints (4.10) - (4.17). These constraints ensure that ζ , η , and ϕ are all equal to zero for branches that are out of service but not hardened. Thus, the power flows on transmission lines that are out of service but not hardened are also zero by constraints (4.6) and (4.7). For transmission lines that are hardened, constraints (4.10) - (4.17) ensure that ζ , η , and ϕ equal α , β , and γ , respectively. Note that this formulation could be written without η or ϕ . However, including these variables allows extensions which include other relaxations (e.g., the QC relaxation of (4.1)). The next section describes how this MISOCP relaxation is used to solve the original MINLP.

4.3 Solution Approach

All problem formulations and solution strategies were implemented using Pyomo, a Python-based algebraic modeling language [Hart et al., 2017]. The AC formulation (4.2) is a stochastic MINLP that is challenging to solve. Here, we use a tailored multi-tree solution approach similar to those presented by Liu et al. [2018] and Kocuk et al. [2017a]. The algorithm involves three main steps. First, the MISOCP (4.3)-(4.38) relaxation is solved with Gurobi 7.5.2 to obtain a lower bound and a candidate integer solution. Second the integer variables are fixed in the MINLP, and an NLP is solved with Ipopt [Wächter and Biegler, 2006] (using the HSL linear solver MA27 [HSL]) to obtain an upper bound on the objective if the problem is feasible. Third, if the gap between the upper and lower bounds is not small enough, a “no-good” cut [D'Ambrosio et al., 2010] is added to the MISOCP in order to remove already visited integer solutions from the feasible region. Given a candidate integer solution

$\mathbb{X}^k = \{x_1^k, x_2^k, \dots, x_N^k\}$, $\mathcal{I} = \{1, 2, \dots, N\}$, $\mathcal{I}_1 = \{i \in \mathcal{I} | x_i^k = 1\}$, $\mathcal{I}_0 = \{i \in \mathcal{I} | x_i^k = 0\}$, the “no-good” cut

$$\sum_{i \in \mathcal{I}_1} x_i - \sum_{i \in \mathcal{I}_0} x_i \leq |\mathcal{I}_1| - 1 \quad (4.39)$$

ensures that the same integer solution (\mathbb{X}^k) is not obtained in subsequent iterations. The three steps are then repeated. In the worst case, the “no-good” cuts guarantee enumeration of all possible integer solutions. As described by both Liu et al. [2018] and Kocuk et al. [2017a], if the NLP subproblems are solved to global optimality, then the algorithm will converge to global optimality for the original MINLP. In this work, we do not solve the NLP subproblems to global optimality, so the algorithm is a heuristic with regards to optimality of the discrete decisions. Additionally, when the integer variables are fixed in the AC DR formulation, the objective is fixed. Therefore, the NLP subproblem reduces to a feasibility check. Note that proving infeasibility of the NLP subproblem still requires global optimization techniques. We terminate the algorithm when Ipopt finds a feasible solution to one of the NLP subproblems. Note that, in many cases, the algorithm converges in one iteration without any integer cuts, and the globally optimal solution is obtained. Additionally, when the algorithm does not converge on the first iteration, we can still compute an optimality gap using the upper bound from the best feasible solution and the lower bound obtained before any integer cuts are added. In the case study, `pglib_opf_case14_ieee`, [IEEE PES Task Force on Benchmarks for Validation of Emerging Power System Algorithms, 2017] 83 of the 87 problems were provably solved to global optimality. For further algorithm details, see the above references. The following section presents results for both the AC and DC DR formulations and compares the results.

4.4 A Comparison of AC and DC Solutions

We used two approaches to evaluate the accuracy and effectiveness of the DC model on the AC system. First, we solved the demand response problem with both

models and compared the solutions. For this, we used `pglib_opf_case14_ieee` and `pglib_opf_case30_as`, both with 30 scenarios. Figures 4.1 and 4.2 show the minimum number of hardened branches for both the AC and DC models as functions of N_z and Δ_b . In both cases, the DC model tends to underestimate the number of branches that need to be hardened to maintain feasibility, but the overall trends are quite similar. We also compared the integer solutions of the two models for `pglib_opf_case14_ieee`. The detailed results are in Tables 4.1 and 4.2. Table 4.3 summarizes these results. First, we define two sets. The set \mathcal{DR}_{DC}^* is the union of the sets of buses selected for DR by the DC model across all values of Δ_b and N_z . The set \mathcal{DR}_{AC}^* is the union of the sets of buses selected for DR by the AC model across all values of Δ_b and N_z . Table 4.3 shows the intersection of these two sets, the buses unique to \mathcal{DR}_{DC}^* , and the buses unique to \mathcal{DR}_{AC}^* . As the tables show, there is strong agreement between the AC and DC solutions. In fact, all of the buses selected by the DC model were also selected by the AC model. There was only one bus selected by the AC model and not the DC model.

Second, for each integer solution obtained with the DC model, we evaluated the fraction of AC-feasible scenarios (that is, fraction of scenarios that were feasible when evaluated using the AC model). Table 4.4 summarizes the results. The table shows the fraction of AC-feasible scenarios averaged over all values of Δ_b and N_z . Note that we did not include `pglib_opf_case300_ieee` because there were several scenarios for which Ipopt neither converged to an optimal solution nor converged to a point of local infeasibility. Case `pglib_opf_case162_ieee_dtc` also had a few scenarios for which Ipopt obtained neither of these two convergence criteria, but it was less than 0.3% of the scenarios. Therefore, we simply discarded these scenarios when computing the fraction of AC-feasible scenarios. Again, a problem may be feasible even if Ipopt converges to a point of local infeasibility. However, for simplicity, we are using this criterion as an approximation. Future work could use global optimization techniques to prove infeasibility. The average fraction of feasible scenarios was over 85% for all

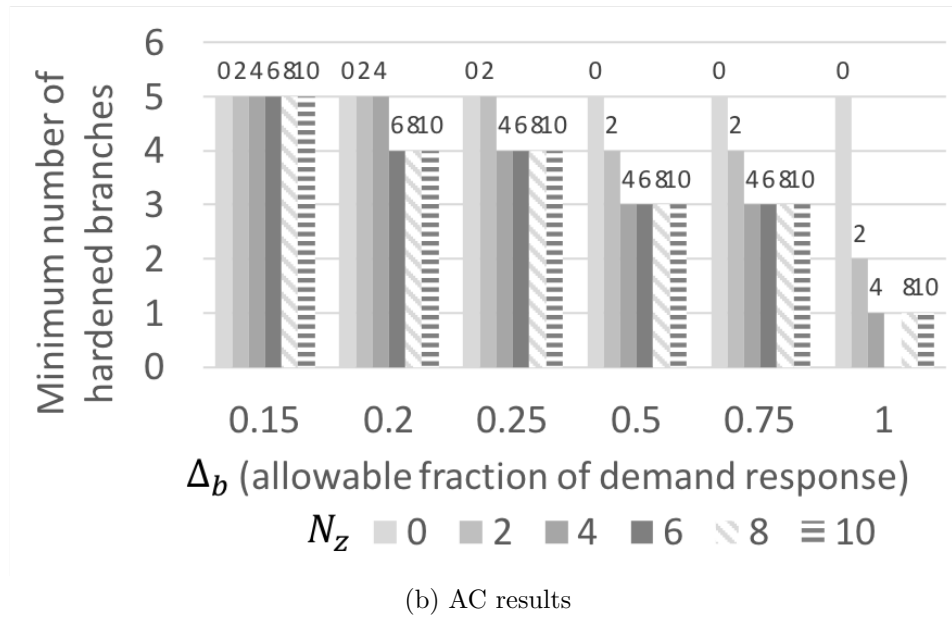



Figure 4.1.: Comparison of DC and AC results for `pglib_opf_case14_ieee`. Each figure shows the minimum number of demand response contracts as a function of Δ_b and N_z .

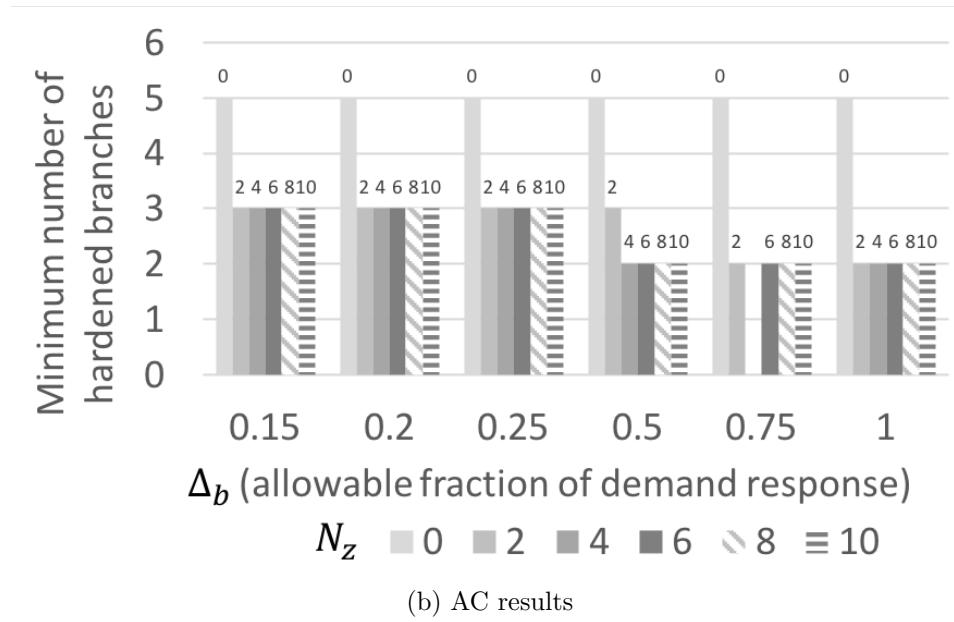



Figure 4.2.: Comparison of DC and AC results for `pglib_opf_case30_as`. Each figure shows the minimum number of demand response contracts as a function of Δ_b and N_z .

Table 4.1.: Integer Solutions DC - pglib_opf_case14_ieee

Δ_b	N_z	Branches Hardened	Demand Response Buses
0	0	1, 6, 18, 19	-
0.05	1-10	1, 6, 18, 19	-
0.1	1-10	1, 6, 18, 19	-
0.15	1-4	1, 6, 18, 19	-
0.15	5-10	3, 18, 19	2, 3, 9, 13, 14
0.2	1	1, 6, 18, 19	-
0.2	2-10	6, 18, 19	3, 9
0.25	1	1, 6, 18, 19	-
0.25	2-10	6, 18, 19	3, 9
0.5	1-10	6, 18, 19	3
0.75	1-10	6, 18, 19	3
1	1-3	18, 19	3
1	4-10	18	3, 12, 13, 14

Table 4.2.: Integer Solutions AC - pglib_opf_case14_ieee

Δ_b	N_z	Branches Hardened	Demand Response Buses
0	0	1, 6, 11, 17, 19	-
0.05	1-10	1, 6, 11, 17, 19	-
0.1	1-10	1, 6, 11, 17, 19	-
0.15	1-10	1, 6, 11, 17, 19	-
0.2	1-5	1, 6, 11, 17, 19	-
0.2	6-10	6, 11, 17, 19	2, 3, 6, 9, 13, 14
0.25	1-2	1, 6, 11, 17, 19	-
0.25	3-10	6, 17, 18, 19	2, 3, 9
0.5	1-2	3, 11, 17, 19	3
0.5	3-10	3, 11, 19	3, 13, 14
0.75	1-2	3, 11, 13, 19	3
0.75	3-10	3, 11, 19	9, 13, 14
1	1	13, 18, 19	3
1	2-3	13, 18	3, 12
1	4-10	18	3, 12, 13, 14

test cases except pglib_opf_case162_ieee_dtc. These are quite high and indicate high quality solutions even for the true AC system.


Table 4.3.: Comparison of DR locations selected by DC and AC models for pglib_opf_case14_ieee across all values of Δ_b and N_z .

Model	DR Buses
$\mathcal{DR}_{DC}^* \cap \mathcal{DR}_{AC}^*$	2, 3, 9, 13, 14, 12
$\mathcal{DR}_{DC}^* \setminus \mathcal{DR}_{AC}^*$	-
$\mathcal{DR}_{AC}^* \setminus \mathcal{DR}_{DC}^*$	6


Table 4.4.: Average fraction of AC-feasible scenarios when implementing DC solution

Test Case	Average fraction of AC-feasible scenarios
pglib_opf_case14_ieee	0.9
pglib_opf_case30_as	0.86
pglib_opf_case118_ieee	0.89
pglib_opf_case162_ieee_dtc	0.67

Finally, Figure 4.3 plots the fraction of AC-feasible scenarios and the minimum number of hardened branches as a function of N_z for the four test cases with $\Delta_b=0.2$. The results for pglib_opf_case118_ieee behave as expected. The fraction of AC-feasible scenarios generally increases as N_z increases. However, there are cases where the fraction of AC-feasible scenarios can decrease with an increase in N_z . The reason for this relates back to Figures 4.1 and 4.2. The DC model often underestimates the amount of DR needed and/or the number of branches that need hardened. In some cases, the DC formulation results in a lower number of hardened branches and smaller fraction of DR than what is actually required to ensure AC-feasibility. An implementation of such a solution in the AC system may decrease capital investment, but it may also not provide the expected resilience. Because of this, it is important to continue development of MINLP strategies that can address AC power flow problems at scale.



(a) pglib_opf_case14_ieee



(b) pglib_opf_case30_as

Figure 4.3.: Fraction of scenarios feasible when integer solutions obtained with the DC model are evaluated with the AC model ($\Delta_b = 0.2$). The solid lines show the minimum number of hardened branches and the dashed lines show the fraction of feasible scenarios, both as a function of N_z .





Figure 4.3.: Fraction of scenarios feasible when integer solutions obtained with the DC model are evaluated with the AC model ($\Delta_b = 0.2$). The solid lines show the minimum number of hardened branches and the dashed lines show the fraction of feasible scenarios, both as a function of N_z .

4.5 Summary

In this chapter, we extended the resilience formulations presented in Chapter 3 by incorporating higher fidelity (but nonconvex) AC power flow constraints in order to investigate the effect of the linear DC power flow approximation. Our key contributions are listed as follows.

- We formulated the demand response problem presented in Section 3.6 with the actual nonlinear AC power flow equations, resulting in a stochastic MINLP.
- We solved the stochastic MINLP using a tailored multi-tree approach involving a sequence of MISOCP relaxations of the MINLP to find candidate integer solutions and NLP subproblems to check feasibility.
- We were able to solve the stochastic MINLP for a few small test cases, and we compared the results with the solution to the linear approximation. Our results with both formulation indicate that there is significant value in the use of demand response for improved resilience. Furthermore, both formulations found similar optimal locations for demand response contracts.
- We also evaluated each integer solution obtained with the DC model on the AC model to compute the fraction of feasible scenarios. This allowed the comparison of results for larger test cases. We found that the linear model often underestimates the amount of demand response necessary to ensure resilience of the AC system, demonstrating the need for improved scalability of algorithms for solving MINLP's based on AC power flow equations.

5. RELAXATIONS AND REFINEMENT TECHNIQUES FOR AC CONSTRAINED POWER FLOW PROBLEMS¹

The alternating current optimal power flow (ACOPF), also referred to as the OPF, is a fundamental problem for reliable and efficient operation of the electric grid and is the basis of more complex operations problems such as unit commitment. Global optimization of the OPF guarantees revenue adequacy of the locational marginal prices (LMPs) by closing the duality gap. Moreover, OPF problems with discrete decision variables, including unit commitment, are modeled as mixed integer nonlinear programming (MINLP) problems, where global solution of the nonlinear ACOPF is a required subproblem for many algorithms. Therefore, efficient global optimization of the ACOPF problem is a critical step towards incorporating higher fidelity models into practical grid operations.

Kocuk et al. [2016] proved that the second order cone (SOC) relaxation is tighter than a linear McCormick relaxation of the rectangular OPF (RM) under certain assumptions on variable bounds. As a result, they initially strengthen the SOC relaxation with arctangent constraints, cycle constraints, and semidefinite programming (SDP) cuts, and later with matrix minor reformulations [Kocuk et al., 2017b]. However, we demonstrate that a quadratic form of the RM relaxation, in combination with explicit reference bus constraints and optimization-based bounds tightening (OBBT) [Caprara and Locatelli, 2010], can also be quite effective in practice for improving

¹©2018 IEEE Part of this chapter is reprinted with permission from “Tightening McCormick Relaxations Toward Global Solution of the ACOPF Problem” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., IEEE Transactions on Power Systems, 2018. In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of Purdue University’s products or services.

Part of this chapter is reprinted from “13th International Symposium on Process Systems Engineering (PSE 2018)”, Volume 44, Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., “Strengthened SOCP Relaxations for ACOPF with McCormick Envelopes and Bounds Tightening”, pages 1555–1560, 2018, with permission from Elsevier.

the performance of the existing SOC [Jabr, 2007] and quadratic convex (QC) [Hijazi et al., 2017] OPF relaxations since this produces tighter bounds on the variables than those assumed in Theorem 3.1 of [Kocuk et al., 2016]. Specifically, we extend our findings in [Bynum et al., 2018a] and demonstrate that neither the RM nor the QC (with McCormick envelopes, reference bus constraints, and OBBT) relaxation dominates the other, and, between the two approaches, we close 91% of the NESTA archive [Coffrin et al., 2014] test cases up to 300 buses. This result requires no branching or piecewise refinement, and therefore is a promising foundation for future work. The remainder of this chapter first reviews these relaxations, and then presents significant improvements to the numerical results reported in [Kocuk et al., 2016] and [Coffrin et al., 2015].

5.1 A Review of Convex Relaxations for the ACOPF Problem

We use lower case symbols to represent variables, upper case symbols to represent parameters, and upper case script symbols to represent index sets. The sets \mathcal{G} , \mathcal{G}_b , and \mathcal{B} represent the sets of all generators, generators connected to bus b , and buses, respectively. The set \mathcal{A} contains all connected bus pairs. The set \mathcal{K} is a set of three-tuples containing both (transmission line index, “from” bus, “to” bus) and (transmission line index, “to” bus, “from” bus) 3-tuples for each transmission line. Finally, $\mathcal{K}_b = \{(l, i, j) \in \mathcal{K} \mid i=b\}$. The superscripts G , D , and Sh represent generation, demand, and shunt, respectively. The parameters $C_{l,b,n}^{\alpha,p}$, $C_{l,b,n}^{\beta,p}$, $C_{l,b,n}^{\gamma,p}$, $C_{l,b,n}^{\alpha,q}$, $C_{l,b,n}^{\beta,q}$, and $C_{l,b,n}^{\gamma,q}$ are functions of branch characteristics [Momoh, 2001].

We explore three relaxations with the following base formulation in (5.1) for co-optimizing real and reactive power, p and q , respectively.

$$\min \sum_{g \in \mathcal{G}} [A_g^2 (p_g^G)^2 + A_g^1 p_g^G + A_g^0] \quad (5.1.1)$$

$$\sum_{g \in \mathcal{G}_b} p_g^G - \sum_{(l,i,j) \in \mathcal{K}_b} p_{l,i,j} = p_b^L + G_b^{Sh} \alpha_b \quad \forall b \in \mathcal{B} \quad (5.1.2)$$

$$\sum_{g \in \mathcal{G}_b} q_g^G - \sum_{(l,i,j) \in \mathcal{K}_b} q_{l,i,j} = q_b^L - B_b^{Sh} \alpha_b \quad \forall b \in \mathcal{B} \quad (5.1.3)$$

$$p_{l,b,n} = C_{l,b,n}^{\alpha,p} \alpha_b + C_{l,b,n}^{\beta,p} \beta_{b,n} + C_{l,b,n}^{\gamma,p} \gamma_{b,n} \quad \forall (l,b,n) \in \mathcal{K} \quad (5.1.4)$$

$$q_{l,b,n} = C_{l,b,n}^{\alpha,q} \alpha_b + C_{l,b,n}^{\beta,q} \beta_{b,n} + C_{l,b,n}^{\gamma,q} \gamma_{b,n} \quad \forall (l,b,n) \in \mathcal{K} \quad (5.1.5)$$

$$p_{l,b,n}^2 + q_{l,b,n}^2 \leq (S_l^{\max})^2 \quad \forall (l,b,n) \in \mathcal{K} \quad (5.1.6)$$

$$(V_b^{\min})^2 \leq \alpha_b \leq (V_b^{\max})^2 \quad \forall b \in \mathcal{B} \quad (5.1.7)$$

$$P_g^{G,\min} \leq p_g^G \leq P_g^{G,\max} \quad \forall g \in \mathcal{G} \quad (5.1.8)$$

$$Q_g^{G,\min} \leq q_g^G \leq Q_g^{G,\max} \quad \forall g \in \mathcal{G} \quad (5.1.9)$$

$$\tan(\Theta_{b,n}^{\min}) \beta_{b,n} \leq \gamma_{b,n} \leq \tan(\Theta_{b,n}^{\max}) \beta_{b,n} \quad \forall (b,n) \in \mathcal{A} \quad (5.1.10)$$

$$\beta_{b,n} = \beta_{n,b}, \gamma_{b,n} = -\gamma_{n,b} \quad \forall (b,n) \in \mathcal{A} \quad (5.1.11)$$

The *rectangular OPF* (i.e., where v^r and v^j represent the real and imaginary components of the nodal voltages, respectively) is defined with the following substitutions for α, β , and γ :

$$\alpha_b \rightarrow (v_b^r)^2 + (v_b^j)^2 \quad \forall b \in \mathcal{B} \quad (5.2.1)$$

$$\beta_{b,n} \rightarrow v_b^r v_n^r + v_b^j v_n^j \quad \forall (b,n) \in \mathcal{A} \quad (5.2.2)$$

$$\gamma_{b,n} \rightarrow v_b^j v_n^r - v_b^r v_n^j \quad \forall (b,n) \in \mathcal{A} \quad (5.2.3)$$

with $v^r, v^j \in [-V^{\max}, V^{\max}]$. The *polar OPF* (i.e., v and θ represent nodal voltage magnitude and angle, respectively) is defined with the following substitutions into (5.1):

$$\alpha_b \rightarrow v_b^2 \quad \forall b \in \mathcal{B} \quad (5.3.1)$$

$$\beta_{b,n} \rightarrow v_b v_n \cos(\theta_b - \theta_n) \quad \forall (b,n) \in \mathcal{A} \quad (5.3.2)$$

$$\gamma_{b,n} \rightarrow v_b v_n \sin(\theta_b - \theta_n) \quad \forall (b,n) \in \mathcal{A}. \quad (5.3.3)$$

Constraints (5.1.2) and (5.1.3) enforce power balances, (5.1.4)-(5.1.5) compute power flows, (5.1.6) enforces thermal limits, (5.1.7) enforces voltage magnitude bounds,

(5.1.8) and (5.1.9) enforce generator limits, and (5.1.10) limits the voltage angle difference for interconnected buses.

The solution to both forms of the OPF problem is non-unique without fixing a reference voltage angle at one of the buses. In the polar form, this reference bus constraint is

$$\theta_{ref} = 0. \quad (5.4)$$

In rectangular form, there is a domain reduction for v_{ref}^r :

$$v_{ref}^j = 0 \quad (5.5.1)$$

$$V_{ref}^{min} \leq v_{ref}^r \leq V_{ref}^{max}. \quad (5.5.2)$$

The choice of reference bus does not impact the optimal solution of the ACOPF problem, but it may impact the quality of the relaxation [Bynum et al., 2018a, Bienstock and Munoz, 2014]. For comparison purposes, we used the reference bus location specified in the test cases.

5.1.1 Second-Order Cone Programming Relaxation

There is a SOC equality relationship amongst α, β , and γ that may be relaxed to the following convex inequality:

$$\beta_{b,n}^2 + \gamma_{b,n}^2 \leq \alpha_b \alpha_n \quad \forall (b, n) \in \mathcal{A}. \quad (5.6)$$

The SOC problem [Jabr, 2007] is given by (5.1) and (5.6).

5.1.2 McCormick Relaxations of the Rectangular Form

McCormick envelopes can be applied to (5.2), yielding

$$\alpha_b = \widehat{v_b^r} + \widehat{v_b^j}, \quad \beta_{b,n} = \widehat{v_b^r v_n^r} + \widehat{v_b^j v_n^j}, \quad \gamma_{b,n} = \widehat{v_b^j v_n^r} - \widehat{v_b^r v_n^j} \quad (5.7.1)$$

$$(v_b^r)^2 \leq \widehat{v_b^r} \leq (V_b^{r,max} + V_b^{r,min})v_b^r - V_b^{r,max}V_b^{r,min} \quad (5.7.2)$$

$$(v_b^j)^2 \leq \widehat{v_b^j} \leq (V_b^{j,max} + V_b^{j,min})v_b^j - V_b^{j,max}V_b^{j,min} \quad (5.7.3)$$

$$\widehat{v_b^r v_n^r} \in MCC(v_b^r, v_n^r), \quad \widehat{v_b^j v_n^j} \in MCC(v_b^j, v_n^j) \quad (5.7.4)$$

$$\widehat{v_b^r v_n^j} \in MCC(v_b^j, v_n^r), \quad \widehat{v_b^r v_n^r} \in MCC(v_b^r, v_n^j) \quad (5.7.5)$$

where $\widehat{xy} \in MCC(x, y)$ denotes the McCormick envelopes for the bilinear term xy . The RM problem is given by (5.1) and (5.5)-(5.7).

5.1.3 Quadratic Relaxations of the Polar Form

The QC formulation [Hijazi et al., 2017] is a quadratic convex relaxation of the polar form in (5.3):

$$v_b^2 \leq \alpha_b \leq (V_b^{max} + V_b^{min})v_b - V_b^{max}V_b^{min} \quad (5.8.1)$$

$$\beta_{b,n} \in MCC(\widehat{v_b v_n}, \widehat{\cos\theta_{b,n}}), \quad \gamma_{b,n} \in MCC(\widehat{v_b v_n}, \widehat{\sin\theta_{b,n}}) \quad (5.8.2)$$

$$\widehat{v_b v_n} \in MCC(v_b, v_n) \quad (5.8.3)$$

$$\widehat{\cos\theta_{b,n}} \in CR(\theta_{b,n}), \quad \widehat{\sin\theta_{b,n}} \in SR(\theta_{b,n}) \quad (5.8.4)$$

$$\theta_{b,n} = \theta_b - \theta_n \quad (5.8.5)$$

$$\Theta_{b,n}^{min} \leq \theta_{b,n} \leq \Theta_{b,n}^{max}, \quad V_b^{min} \leq v_b \leq V_b^{max}. \quad (5.8.6)$$

Here, $\widehat{\cos x} \in CR(x)$ and $\widehat{\sin x} \in SR(x)$ denote relaxations of the cosine and sine functions, respectively [Hijazi et al., 2017]. Additionally, when $\Theta_{b,n}^{min} \geq 0$ ($\Theta_{b,n}^{max} \leq 0$), the sine function is concave (convex) and requires two linear over (under) estimators; this is the linear variant of S-CONV in [Coffrin et al., 2015]. The QC problem is given by (5.1), (5.4), (5.6), and (5.8).

5.2 Numerical Results

To analyze the strength of the different relaxations, we iteratively solve the specified relaxation and perform OBBT to compute valid bounds for the specified convex relaxation \mathcal{R} of the ACOPF feasible region by minimizing (for the lower bound) and maximizing (for the upper bound) each variable, i.e.,

$$\min / \max \{x_i | x \in \mathcal{R}, f(x) \leq UB\}. \quad (5.9)$$

Here, $f(x)$ denotes the objective function in (5.1.1) and UB denotes the objective value of the best-known solution to the NLP. Note that any of the relaxations may define \mathcal{R} . A more efficient approach to OBBT is to optimize only over a subset of the constraints in \mathcal{R} (e.g., corresponding to a subset of the network [Kocuk et al., 2016]). However, this can lead to weaker bounds. Our bounds tightening procedure is similar to the *minimal continuous constraint relaxation network algorithm* presented in [Coffrin et al., 2015] but includes the UB constraint.

To evaluate the impact of the UB constraint, we implement OBBT with and without the UB constraint, denoted as **OBBT(UB)** and **OBBT**, respectively. All problems were modeled with Pyomo [Hart et al., 2017] and solved with IPOPT [Wächter and Biegler, 2006] using the linear solver MA27 [HSL, 2007]. OBBT was performed in parallel on a cluster with 24 64-GB-RAM nodes and 16 2.6 GHz Intel Sandy Bridge cores per node. We used 12 processes per node. The relaxations compared and summary results are as follows:

Relaxation	Description	Cases Closed ²
SOC	Eqs (5.1), (5.6)	16%
RM ^o	Eqs (5.1), (5.6)-(5.7) & OBBT(UB)	16%
RM ^r	Eqs (5.1), (5.5)-(5.7) & OBBT	52%
RM ^{ro}	Eqs (5.1), (5.5)-(5.7) & OBBT(UB)	67%
QC ^r	Eqs (5.1), (5.4), (5.6), (5.8) & OBBT	67%
QC ^{ro}	Eqs (5.1), (5.4), (5.6), (5.8) & OBBT(UB)	90%

We note that because the reference bus selection does not result in further domain reduction for polar OPF, QC° was omitted.

Detailed results are reported in Tables 5.1 - 5.6 for the NESTA archive [Coffrin et al., 2014], including the optimality gap, wallclock time, and total iterations. We define a single iteration as performing OBBT once on all appropriate variables. Our stopping criteria are as follows: (1) optimality gap less than 0.1%, (2) wallclock time exceeds one hour, or (3) optimality gap improved less than 0.1% in 20 iterations of OBBT. For (3), the gap and wallclock time reported were obtained by the last iteration prior to stalling.

Our major findings are that the $\text{OBBT}(\text{UB})$ significantly improves performance across relaxations, and neither the $\text{RM}^{\text{r}\circ}$ nor the $\text{QC}^{\text{r}\circ}$ dominates the other. Overall, the $\text{QC}^{\text{r}\circ}$ formulation performs best. However, the $\text{RM}^{\text{r}\circ}$ formulation can significantly

²Percentage of 58 NESTA cases with less than 0.1% optimality gap.

²Initializing the NLP from the updated solution to the relaxation and recomputing the UB between iterations of OBBT results in an optimality gap reduction from 0.5% to less than 0.1% for both $\text{RM}^{\text{r}\circ}$ and $\text{QC}^{\text{r}\circ}$.

Table 5.1.: Optimality Gaps for relaxations of NESTA archive standard operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE

Case	Optimality Gap(%)					
	SOC	RM°	RM^{r}	$\text{RM}^{\text{r}\circ}$	QC^{r}	$\text{QC}^{\text{r}\circ}$
3_lmbd	1.3	1.3	0.5	0.0	0.2	0.0
5_pjm	14.5	14.5	5.0	0.1	9.3	5.7
6_c	0.3	0.3	0.1	0.0	0.1	0.1
6_ww	0.6	0.6	0.1	0.1	0.0	0.0
14_ieee	0.1	0.1	0.1	0.1	0.0	0.0
29_edin	0.1	0.1	0.1	0.1	0.1	0.1
30_fsr	0.4	0.4	0.4	0.1	0.1	0.1
30_ieee	15.9	15.9	0.1	0.0	0.0	0.1
118_ieee	1.8	1.8	1.6	0.7	0.5	0.1
162_ieee_dtc	4.0	4.0	4.0	4.0	0.7	0.1
189_edin	0.2	0.2	0.2	0.2	0.2	0.1
300_ieee	1.2	1.2	1.2	1.2	0.2	0.0

Table 5.2.: Optimality Gaps for relaxations of NESTA archive congested operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE

Case	Optimality Gap(%)					
	SOC	RM ^o	RM ^r	RM ^{ro}	QC ^r	QC ^{ro}
3_lmbd__api	3.3	3.3	0.1	0.1	0.1	0.0
4_gs__api	0.7	0.7	0.1	0.0	0.0	0.1
5_pjm__api	0.3	0.3	0.1	0.0	0.0	0.0
6_c__api	0.3	0.3	0.1	0.0	0.1	0.0
14_ieee__api	1.3	1.3	0.2	0.0	0.3	0.1
24_ieee_rts__api	20.8	20.8	2.0	0.6	0.3	0.1
29_edin__api	0.4	0.4	0.4	0.4	0.1	0.1
30_as__api	4.8	4.8	0.3	0.1	0.0	0.0
30_fsr__api	46.0	46.0	41.6	41.2	2.4	0.1
30_ieee__api	1.0	1.0	0.1	0.1	0.1	0.0
39_epri__api	3.0	3.0	0.3	0.1	0.1	0.0
57_ieee__api	0.2	0.2	0.1	0.1	0.0	0.1
73_ieee_rts__api	14.4	14.4	14.4	14.4	0.2	0.0
89_pegase__api	20.4	20.4	20.2	20.2	18.9	9.1
118_ieee__api	43.9	43.9	26.9	26.4	9.3	8.7
162_ieee_dtc__api	1.3	1.3	1.0	0.9	0.1	0.1
189_edin__api	5.7	5.7	5.5	2.9	0.3	0.1
300_ieee__api	0.7	0.7	0.7	0.7	0.1	0.0

Table 5.3.: Optimality Gaps for relaxations of NESTA archive small angle difference operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%.
©2018 IEEE

Case	Optimality Gap(%)					
	SOC	RM ^o	RM ^r	RM ^{ro}	QC ^r	QC ^{ro}
3_lmbd__sad	4.3	4.3	0.1	0.0	0.0	0.0
4_gs__sad	4.9	4.9	0.0	0.0	0.0	0.0
5_pjm__sad	3.6	3.6	0.0	0.0	0.0	0.0
6_c__sad	1.4	1.4	0.0	0.0	0.0	0.0
6_ww__sad	0.8	0.8	0.0	0.0	0.0	0.0
9_wsc__sad	1.5	1.5	0.0	0.0	0.0	0.0
24_ieee_rts__sad	11.4	11.4	0.1	0.0	0.1	0.0
29_edin__sad	34.7	34.7	2.4	0.5 ²	1.2	0.5 ²
30_as__sad	9.2	9.2	0.2	0.1	0.1	0.0
30_fsr__sad	0.6	0.6	0.2	0.0	0.1	0.0
30_ieee__sad	5.8	5.8	0.1	0.1	0.0	0.0
39_epri__sad	0.1	0.1	0.0	0.0	0.0	0.0
57_ieee__sad	0.1	0.1	0.1	0.1	0.1	0.1
73_ieee_rts__sad	8.4	8.4	2.6	2.0	0.1	0.0
89_pegase__sad	0.3	0.3	0.3	0.1	0.1	0.1
118_ieee__sad	12.8	12.8	5.1	1.5	1.4	0.1
162_ieee_dtc__sad	7.1	7.1	7.1	7.1	0.5	0.0
189_edin__sad	2.3	2.3	2.2	1.8	1.0	0.8
300_ieee__sad	1.3	1.3	1.3	1.3	0.2	0.0

Table 5.4.: Computational Performance for relaxations of NESTA archive standard operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%.
©2018 IEEE

Case	Wallclock Time(s) / # Iterations					
	SOC	RM ^o	RM ^r	RM ^{ro}	QC ^r	QC ^{ro}
3_lmbd	0/0	0/0	1/4	1/3	1/3	1/3
5_pjm	0/0	0/0	3/6	35/84	3/6	7/15
6_c	0/0	0/0	4/9	2/3	1/2	1/1
6_ww	0/0	0/0	1/1	1/1	1/1	1/1
14_ieee	1/0	2/0	3/1	3/1	4/1	3/1
29_edin	3/0	4/0	38/9	16/3	4/0	4/0
30_fsr	1/0	2/0	2/0	16/7	8/3	6/2
30_ieee	1/0	1/0	18/8	12/5	8/3	6/2
118_ieee	4/0	7/0	217/10	253/18	59/2	83/3
162_ieee_dtc	7/0	12/0	13/0	13/0	430/7	488/8
189_edin	5/0	9/0	9/0	9/0	10/0	46/1
300_ieee	14/0	23/0	22/0	23/0	290/2	284/2

Table 5.5.: Computational Performance for relaxations of NESTA archive congested operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%.
©2018 IEEE

Case	Wallclock Time(s) / # Iterations					
	SOC	RM ^o	RM ^r	RM ^{ro}	QC ^r	QC ^{ro}
3_lmbd__api	0/0	0/0	2/7	1/3	2/6	1/3
4_gs__api	0/0	0/0	1/3	1/3	1/3	1/2
5_pjm__api	0/0	0/0	1/2	1/2	1/2	1/2
6_c__api	0/0	0/0	3/5	1/2	2/3	1/2
14_ieee__api	0/0	1/0	9/8	8/7	4/3	5/4
24_ieee_rts__api	1/0	1/0	21/10	27/13	8/3	8/3
29_edin__api	7/0	5/0	5/0	5/0	25/3	35/3
30_as__api	1/0	2/0	25/11	25/11	9/3	8/3
30_fsr__api	1/0	1/0	57/25	63/28	27/11	27/11
30_ieee__api	1/0	1/0	27/12	18/8	11/4	11/4
39_epri__api	1/0	2/0	63/24	45/17	7/2	7/2
57_ieee__api	1/0	3/0	61/13	25/5	16/2	9/1
73_ieee_rts__api	3/0	5/0	5/0	5/0	29/2	40/3
89_pegase__api	5/0	11/0	179/10	208/11	1663/24	3604/55
118_ieee__api	4/0	8/0	1102/75	1074/74	318/10	301/11
162_ieee_dtc__api	7/0	12/0	824/24	785/23	799/13	433/7
189_edin__api	5/0	9/0	276/9	746/28	83/2	85/2
300_ieee__api	14/0	23/0	23/0	23/0	346/2	291/2

Table 5.6.: Computational Performance for relaxations of NESTA archive small angle difference operating conditions cases up to 300 buses. Highlighting indicates cases not closed to less than 0.1% gap. Cases are not shown if the SOC relaxation gap is below 0.1%. ©2018 IEEE

Case	Wallclock Time(s) / # Iterations					
	SOC	RM ^o	RM ^r	RM ^{ro}	QC ^r	QC ^{ro}
3_lmbd__sad	0/0	0/0	1/3	1/3	1/1	1/1
4_gs__sad	0/0	0/0	1/2	1/2	1/1	1/1
5_pjm__sad	0/0	0/0	2/3	2/3	1/1	1/1
6_c__sad	0/0	0/0	2/4	2/3	1/1	1/1
6_ww__sad	0/0	0/0	1/1	1/1	1/1	1/1
9_wscs__sad	0/0	0/0	3/3	2/3	2/1	1/1
24_ieee_rts__sad	1/0	2/0	17/8	15/7	8/3	6/2
29_edin__sad	3/0	5/0	72/14	78/14	40/4	33/4
30_as__sad	1/0	2/0	19/8	17/7	6/2	6/2
30_fsr__sad	1/0	1/0	20/8	14/6	6/2	6/2
30_ieee__sad	1/0	1/0	12/5	10/4	6/2	6/2
39_epri__sad	1/0	2/0	25/9	17/6	2/0	2/0
57_ieee__sad	1/0	3/0	25/5	16/3	3/0	3/0
73_ieee_rts__sad	3/0	4/0	124/16	169/21	28/2	28/2
89_pegase__sad	5/0	9/0	57/3	162/10	91/2	43/1
118_ieee__sad	4/0	8/0	516/34	520/33	113/4	248/9
162_ieee_dtc__sad	7/0	13/0	13/0	13/0	674/11	563/9
189_edin__sad	5/0	9/0	155/5	214/8	46/1	46/1
300_ieee__sad	14/0	23/0	23/0	23/0	162/1	281/2

tighten the SOC relaxation and may be improved by adjusting the location of the reference bus [Bienstock and Munoz, 2014]. This is left for future work. Across all relaxations tested, only 5 cases did not solve to less than a 0.1% optimality gap. This is a *significant improvement* to the results reported in [Kocuk et al., 2016] and [Coffrin et al., 2015]. Note that this advancement requires no branching or piecewise refinement, which is a promising foundation for future work on the global optimization of the OPF. We recognize that the full OBBT approach considered here will not scale well to larger cases, and integration of these techniques (and those described in [Kocuk et al., 2016] and [Coffrin et al., 2015]) within a scalable global optimization framework is left for future work.

5.3 The Impact of the Reference Bus

Figure 5.1 shows the importance of the reference bus. The figure shows the variable bound ranges for the real and imaginary components of the voltages averaged over all buses of the same shortest path in terms of the number of transmission lines from the reference bus. The iteration count for OBBT is shown on the x-axis. The figure shows the first four iterations of bounds tightening for `nesta_case24_ieee_rts_api`. As shown in the figure, before any bounds tightening is performed, the reference bus is the only bus with good bounds on the voltage. After one iteration of bounds tightening, the bounds on the voltages at buses connected directly to the reference bus are tightened dramatically. On the second iteration of bounds tightening, the bounds on the voltages at buses a distance of 2 from the reference bus are tightened more than the bounds on variables at other buses. The trend continues through all four iterations of bounds tightening shown. In this figure, it is clear that the effectiveness of bounds tightening at a given iteration is directly related to the distance of the corresponding bus from the reference bus.

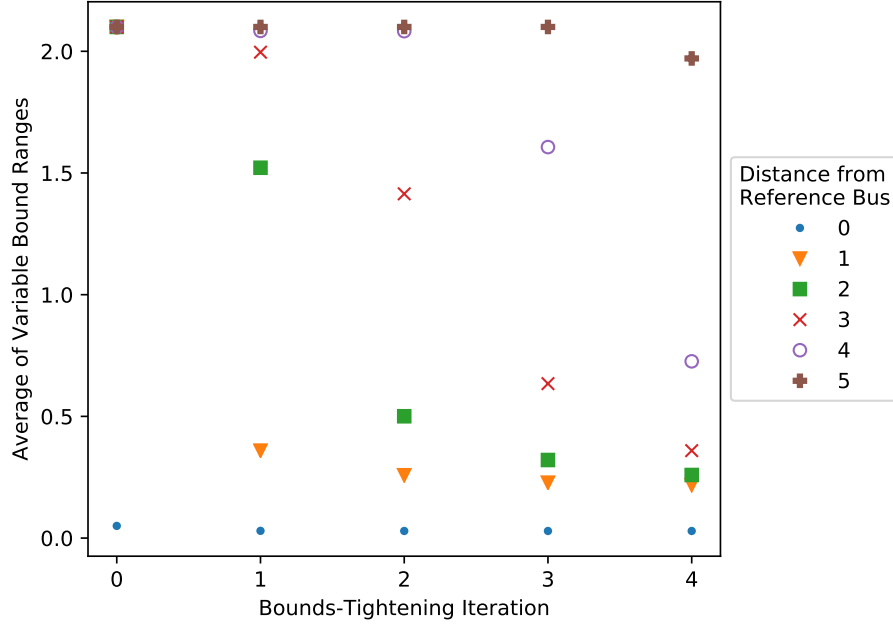


Figure 5.1.: Impact of the reference bus on bounds tightening for `nesta_case24_ieee_rts_api`. The abscissa shows the iteration count, and the ordinate shows the average range between the upper and lower variable bounds.

5.4 Summary

In this chapter, we focused on relaxations and refinement techniques for AC-constrained power flow problems. Our key contributions are listed below.

- We demonstrate that relaxations to the rectangular form of the ACOPF problem can significantly strengthen existing relaxations if reference bus constraints are properly incorporated.
- We illustrate that, for relaxations of the rectangular form, the effectiveness of optimization-based bounds tightening at a given iteration is directly related to the distance of the corresponding bus from the reference bus.
- We are able to reduce the optimality gap to less than 0.1% on all but 5 NESTA test cases with up to 300 buses with optimization-based bounds tightening alone.

6. DECOMPOSING OPTIMIZATION-BASED BOUNDS TIGHTENING PROBLEMS VIA GRAPH PARTITIONING

Mixed-integer nonlinear programming (MINLP) encompasses a wide range of important optimization problems [Puranik and Sahinidis, 2017]. Common algorithms for global solution of both nonlinear programming (NLP) problems and MINLPs include branch-and-bound (B&B) methods and multitree methods (e.g., outer-approximation) [Dakin, 1965, Land and Doig, 1960, Belotti et al., 2013]. B&B techniques involve the solution of convex relaxations of the original problem to obtain lower bounds on the objective and local solutions to obtain upper bounds on the objective. The algorithm progressively refines these bounds by branching on variables to divide the feasible region into smaller sub-regions [Puranik and Sahinidis, 2017]. Outer-approximation (OA) methods typically solve mixed-integer linear programs (MILPs) globally to obtain lower bounds, pushing the branching requirements to an efficient MILP solver [Liu et al., 2017a].

Both of these algorithms rely heavily on bounds tightening (BT) for accelerating convergence. Also referred to as domain reduction, BT can strengthen convex relaxations for the lower bounding problems, reduce the number of nodes visited in B&B trees, and reduce the required number of binary variables needed in MILP relaxations for OA. For example, turning off domain reduction techniques in BARON, an MINLP B&B solver [Tawarmalani and Sahinidis, 2004], results in a 47% increase in computational time and an 802% increase in number of nodes explored for problems in MINLPLib [Bussieck et al., 2003, Puranik and Sahinidis, 2017].

Many BT techniques for MINLPs exist. Consider the following MINLP:

$$\min f(x) \tag{6.1.1}$$

$$s.t. \tag{6.1.2}$$

$$h_c(x, y) \leq 0 \quad \forall c \in \mathbb{C} \tag{6.1.3}$$

$$\underline{x}_v \leq x_v \leq \overline{x}_v \quad \forall v \in \mathbb{V}_x \tag{6.1.4}$$

$$y_v \in \{0, 1\} \quad \forall v \in \mathbb{V}_y \tag{6.1.5}$$

where \mathbb{C} is the set of constraints, \mathbb{V}_x is the set of continuous variables, and \mathbb{V}_y is the set of binary variables. For simplicity, we assume $f(x)$ is convex. Also consider the following convex relaxation of Problem (6.1):

$$\min f(x) \tag{6.2.1}$$

$$s.t. \tag{6.2.2}$$

$$g_c(x, y) \leq 0 \quad \forall c \in \mathcal{C} \tag{6.2.3}$$

$$\underline{x}_v \leq x_v \leq \overline{x}_v \quad \forall v \in \mathcal{V}_x \tag{6.2.4}$$

$$0 \leq y_v \leq 1 \quad \forall v \in \mathcal{V}_y \tag{6.2.5}$$

where \mathcal{C} is the set of constraints comprising a convex relaxation of the original set of constraints. Here, each constraint, g_c , is convex. We have allowed for the introduction of auxiliary variables through the new set of continuous variables, \mathcal{V}_x .

An effective but computationally expensive approach to bounds tightening is optimization-based bounds tightening (OBBT) [Puranik and Sahinidis, 2017]. OBBT typically involves solving two convex optimization problem for each variable (or a subset of the variables):

$$\min / \max x_i \tag{6.3.1}$$

$$s.t. \tag{6.3.2}$$

$$(6.2.3), (6.2.4), (6.2.5) \tag{6.3.3}$$

$$f(x) \leq U \tag{6.3.4}$$

Here, U is the best-known feasible solution to Problem (6.1). If the optimal objective of the minimization problem is larger than the current lower bound, \underline{x}_i , for the corresponding variable, then the lower bound can be updated. A similar argument holds for the upper bound and the maximization problem.

A very efficient approach to BT is feasibility-based bounds tightening (FBBT) and extensions [Puranik and Sahinidis, 2017, Belotti et al., 2010]. FBBT is typically performed by forming a directed acyclic graph (DAG) from the constraints of Problem (6.1) and using interval arithmetic to propagate bounds through the graph [Moore et al., 2009]. FBBT is usually less effective than OBBT because only one constraint is considered at a time [Puranik and Sahinidis, 2017]. Belotti et al. [2010] present a linear program (LP) which converges to the fixed-point of the FBBT algorithm if the DAG were formed from a linear relaxation of Problem (6.1), enabling bounds updates for all variables with the solution of one LP.

Ryoo and Sahinidis [1995] utilize the lagrange multipliers from the solution of Problem (6.2) to tighten bounds. Let L and U be the optimal objective value of Problem (6.2) and a valid upper bound for Problem (6.1), respectively. If a variable x_i is at its upper bound, \overline{x}_i , with an optimal multiplier value $\lambda^* > 0$ at the solution of Problem (6.2), then the lower bound may be updated as follows

$$\underline{x}_i = \overline{x}_i - \frac{U - L}{\lambda^*} \quad (6.4)$$

A similar argument holds if the variable is at its lower bound at the solution of Problem (6.2) [Ryoo and Sahinidis, 1995]. This approach is also very efficient because it only requires the solution of Problem (6.2), which is typically solved at every node of the B&B tree. For a more thorough review of BT methods, see [Puranik and Sahinidis, 2017].

OBBT has been shown to be very effective for refinement of optimal power flow (OPF) problems. Coffrin et al. [2015] demonstrate that OBBT is very effective on polar relaxations of the optimal power flow problem. Relaxations of the rectangular

form, tightened with OBBT, can also be quite effective if reference bus constraints are incorporated [Bynum et al., 2018b]. In fact, Bynum et al. [2018a] demonstrate that the effectiveness of OBBT at a given iteration is directly related to the topological distance of the corresponding variable from the reference bus. Multiple global optimization algorithms have been written for the OPF problem which rely heavily on OBBT [Liu et al., 2017a, Kocuk et al., 2017b, Lu et al., 2018].

Despite the effectiveness of OBBT on OPF problems, the approach is very computationally expensive, especially for large problems. To mitigate the high computational cost, several more efficient approaches to BT for OPF problems have been proposed. Chen et al. [2016] present closed-form bound tightening methods by writing the first order necessary conditions for minimizing or maximizing a variable subject to a small subset of constraints. Kocuk et al. [2016] use an efficient form of OBBT for the OPF problem where, for each variable, a small optimization problem is solved by only considering a small portion of the electric grid and the corresponding constraints. In other words, a valid form of Problem (6.3) is solved for each variable, but only a carefully selected subset of the constraints are included. There is a tradeoff between the strength of the bounds obtained and the computational effort required because, as the size of the network considered is reduced, information on feasibility of the remainder of the network is lost.

In this chapter, we propose a decomposed bounds tightening (DBT) algorithm based on graph partitioning, extending the BT ideas presented by Kocuk et al. [2016]. The main idea is to solve a few large optimization problems (corresponding to Problem (3)) and many small optimization problems. As discussed in more detail below, by solving a few large bounds tightening problems, we retain feasibility information from parts of the network not considered when solving the small bounds tightening problems. Although the algorithm is general, we demonstrate the effectiveness of the algorithm on OPF test cases because OBBT is known to work well for these problems and the power systems networks have a graph structure which is favorable for our graph-partitioning-based algorithm.

The remainder of this chapter is outlined as follows. In the following section we present our proposed DBT algorithm in detail. We then present a mixed-integer linear programming (MILP) formulation that can be used to perform the graph partitioning step of the DBT algorithm. We then describe the application of our algorithm to the OPF problem and the details of our implementation. Finally, we present numerical results and summarize the chapter.

6.1 Decomposed Bounds Tightening Algorithm

In this section, we present our decomposed bounds tightening (DBT) algorithm. The basis of the algorithm is that small optimization problems can be solved efficiently. One way to solve small BT problems is to solve the OBBT problems (6.3) but with most of the constraints discarded. Although this does improve computational performance, it produces weaker bounds because the feasible region has been enlarged. The DBT algorithm attempts to retain information from the removed constraints by first solving a few large OBBT problems with all of the constraints. This idea will be explained further after a few definitions and a more formal description of the algorithm.

Consider a graph, G_0 , defined by Eq. (6.2.3) where the set of nodes, \mathcal{V}_0 , represents variables. The set of edges, \mathcal{E}_0 , is defined by the set of constraints. For each constraint, there is an edge in G_0 for every pair of variables in the constraint. The algorithm begins by partitioning G_0 into two sub-graphs, G_a and G_b , of approximately equal numbers of variables (nodes) by removing a set of constraints. At this point, bounds tightening can be performed by solving optimization problems defined by G_a or G_b . These optimization problems will be approximately half of the size of Problem (6.3) and, therefore, can be solved more efficiently. For example, G_a will have a set of nodes (variables), \mathcal{V}_a , and a set of edges (constraints), \mathcal{E}_a . Let \mathcal{C}_a be the set of

constraints defining the set of edges \mathcal{E}_a . For each variable, i , in \mathcal{V}_{ax} , the following bounds tightening problems can be solved.

$$\min / \max x_i \tag{6.5.1}$$

$$s.t. \tag{6.5.2}$$

$$g_c(x, y) \leq 0 \quad \forall c \in \mathcal{C}_a \tag{6.5.3}$$

$$\underline{x}_v \leq x_v \leq \overline{x}_v \quad \forall v \in \mathcal{V}_{ax} \tag{6.5.4}$$

$$0 \leq y_v \leq 1 \quad \forall v \in \mathcal{V}_{ay} \tag{6.5.5}$$

However, the resulting bounds will not be as tight as solving Problem (6.3) because Problem (6.5) contains a subset of the constraints from Problem (6.3). In order to retain information from the constraints in G_b , we first perform bounds tightening on the variables associated with the constraints removed in order to partition G_0 . These bounds tightening problems, hopefully very few in number, are solved with all of the constraints in Eq. (6.2.3).

Depending on the original size of the problem, bounds tightening problems associated with G_a and G_b may still be large. To address this, the graph partitioning may be continued recursively until the bounds tightening problems are sufficiently small. This process produces a binary tree of graphs. Each graph represents a set of edges (constraints) and nodes (variables). A graph is at stage s if s partitions were needed to generate the graph. A graph is a leaf if it is not partitioned. Finally, G_0 is referred to as the root graph.

Algorithm 1 presents the DBT algorithm in more detail. The algorithm begins by defining G_0 , s , T_s , V_{max} , and N which represent the graph defined by the original convex relaxation (Eq. (6.2.3)), the partitioning stage, the set of graphs at stage s , the maximum number of nodes (variables) in the leaf graphs, and the maximum number of bounds tightening iterations, respectively. The purpose of the first while loop is to partition the graphs until the leaf graphs have less than V_{max} nodes. The

Algorithm 1: Decomposed Bounds Tightening

Initialize:

1. $G_0 = (\mathcal{V}_0, \mathcal{E}_0)$ is the graph of the convex relaxation of the original optimization problem. The set of nodes, \mathcal{V}_0 , represents variables. For each constraint, there is an edge for every pair of variables in the constraint.
2. $s = 0$, $\mathcal{T}_s = \{G_0\}$
3. Maximum number of variables in leaves, V_{max}
4. Maximum number of bounds tightening iterations, N

while $\mathcal{T}_s \neq \emptyset$ **do**

$\mathcal{T}_{s+1} = \emptyset$;

foreach $G_i = (\mathcal{V}_i, \mathcal{E}_i) \in \mathcal{T}_s$ **do**

if $|\mathcal{V}_i| > N$ **then**

 Partition G_i into two graphs, G_a and G_b , by removing a set of constraints;

$\mathcal{T}_{s+1} = \mathcal{T}_{s+1} \cup \{G_a, G_b\}$;

else

G_i is a leaf;

end

end

$s = s + 1$;

end

$S = s$;

$n = 0$;

while $n < N$ **do**

$s = 0$;

while $s < S$ **do**

foreach $G_i = (\mathcal{V}_i, \mathcal{E}_i) \in \mathcal{T}_s$ **do**

if G_i is a leaf **then**

 Perform OBBT on each relevant variable in \mathcal{V}_i subject to the constraints in G_i ;

else

 Perform OBBT on each variable involved in the constraints removed in order to partition G_i ;

end

end

 Update any relaxations that depend on the bounds of the tightened variables;

$s = s + 1$;

end

$n = n + 1$;

end

purpose of the second while loop is to perform bounds tightening. Bounds tightening starts at stage 0 and proceeds with increasing stage. At each stage s and for each graph in T_s , bounds tightening problems are defined by the variables and constraints in the graph. If the graph is not a leaf, BT problems are solved for each variable involved in the constraints removed in order to partition the graph. If the graph is a leaf, then BT problems are solved for any variables whose bounds affect the tightness of the relaxation. After performing bounds tightening, any relaxations that depend on variable bounds should be updated with the new bounds.

6.2 Graph Partitioning

We formulated a MIP to partition the graphs. Balanced graph partitioning is a well-studied topic with many well-established algorithms. We formulate a custom MILP to account for the fact that the partitioning is done by removing constraints but there may be multiple edges in the graph for a single constraint. The binary variable z_c takes a value of 0 if constraint c is being removed in order to partition the graph and 1 otherwise. The binary variable y_i is 1 if node i (variable i from the convex relaxation of the original global optimization problem) is connected to an arbitrary reference node, ref , and 0 otherwise. Any node connected to the reference node will be placed in graph G_a , and any node not connected to the reference node will be placed in graph G_b . The value of y_{ref} is fixed to 1. The binary variable $x_{i,j}$ is 1 if nodes i and j are directly connected through a constraint (edge). The set \mathcal{K} is the set of all pairs of nodes connected directly through a constraint (i.e., pairs of nodes with an edge between them). The set \mathcal{K}_c is the set of pairs of nodes in constraint c . The set \mathcal{V}_i is the set of nodes directly connected to node i through a constraint. The MILP is presented below.

$$\max \sum_{c \in \mathcal{C}} z_c \tag{6.6.1}$$

$$s.t. \tag{6.6.2}$$

$$x_{i,j} \geq z_c \quad \forall (i,j) \in \mathcal{K}_c \quad \forall c \in \mathcal{C} \quad (6.6.3)$$

$$y_i \geq x_{i,j} y_j \quad \forall j \in \mathcal{V}_i \quad \forall i \in \mathcal{V} \quad (6.6.4)$$

$$y_i \leq \sum_{\{j \in \mathcal{V}_i | (i,j) \in \mathcal{K}\}} x_{i,j} y_j + \sum_{\{j \in \mathcal{V}_i | (j,i) \in \mathcal{K}\}} x_{j,i} y_j \quad \forall i \in \mathcal{V} \quad (6.6.5)$$

$$\sum_{i \in \mathcal{V}} y_i \geq 0.375 |\mathcal{V}| \quad (6.6.6)$$

$$\sum_{i \in \mathcal{V}} y_i \leq 0.625 |\mathcal{V}| \quad (6.6.7)$$

$$y_{ref} = 1 \quad (6.6.8)$$

$$z_c \in \{0, 1\} \quad \forall c \in \mathcal{C} \quad (6.6.9)$$

$$x_{i,j} \in \{0, 1\} \quad \forall (i,j) \in \mathcal{K} \quad (6.6.10)$$

$$y_i \in \{0, 1\} \quad \forall i \in \mathcal{V} \quad (6.6.11)$$

For convenience, the formulation is nonlinear as written. However, an exact transformation to an MILP is possible by linearization of the bilinear terms with McCormick Envelopes. The objective is to minimize the number of constraints removed in order to partition the graph. Constraint (6.6.3) ensures that two nodes are connected if any constraint containing both of the corresponding variables is not removed. Constraints (6.6.4) - (6.6.5) serve to identify whether or not a node is connected to the reference node. The idea is that node i is connected to the reference node if any node directly connected to node i is connected to the reference node. In a sense, this is a recursive constraint. Constraints (6.6.6) and (6.6.7) ensure that the two resulting graphs are roughly equal in size.

We note that a more effective formulation would minimize the number of variables involved in removed constraints. For simplicity of implementation (as described in detail later), we use Problem (6.6).

6.3 Application to Optimal Power Flow

We test the DBT algorithm on several instances of the alternating current optimal power flow (ACOPF) problem (a continuous, nonconvex NLP) from version 17.08 of the Power Grid Lib - Optimal Power Flow repository [IEEE PES Task Force on Benchmarks for Validation of Emerging Power System Algorithms, 2017]. The ACOPF problem seeks to minimize the operating cost of the system while satisfying customer power demand. The problem is defined by a set of buses connected by a set of transmission lines. Each bus may have a generator and/or load. The constraints describing power flow on the transmission lines are nonlinear and nonconvex. We utilize the Quadratic Convex (QC) relaxation [Hijazi et al., 2017] of the polar form of the ACOPF problem strengthened with a second-order cone (SOC) constraint [Jabr, 2006]. For ease of implementation, we solve the partitioning problems with graphs defined by the buses and transmission lines rather than the variables and constraints. We partitioned the graphs by removing transmission lines. We subsequently removed the constraints associated with those removed transmission lines and performed bounds tightening on the variables in those constraints.

All models were implemented with Pyomo [Hart et al., 2017]. Pyomo Blocks allowed the model to be constructed with the same binary tree structure as the tree resulting from partitioning the graphs. The NLPs were solved with Ipopt [Wächter and Biegler, 2006] to obtain an upper bound. The SOC relaxations were solved with Gurobi [Gurobi Optimization, 2016] to obtain lower bounds. The bounds tightening problems were also solved with Gurobi.

6.4 Results

We compare three versions of bounds tightening for the ACOPF Problem, including traditional OBBT with all variables and constraints (Full Space), our proposed decomposed bounds tightening algorithm (Decomposed), and a version where we only perform BT with the leaf graphs (Leaves). The computational results are summarized

in Figure 6.1. Figure 6.1 shows the time required to complete one iteration of BT for each of the test cases. The figure shows that there is not a significant computational difference between the three methods for small cases. However, as the number of buses grows, the computational expense of the Full Space method grows very quickly. In comparison, the Decomposed and Leaves methods both scale remarkably well.

Figures 6.2 - 6.8 show detailed results for each of the test cases. The figures show time on the x-axis and percent gap on the y-axis. The figures are ordered with increasing problem size. For the smaller problems, all three methods typically perform similarly. However, as the problem size increases, the differences become dramatic. Typically, the Decomposed and Leaves methods initially perform similarly, both far more effective than the Full Space method. However, as BT proceeds, the Leaves method stalls at optimality gaps much larger than either the Decomposed or Full Space Methods. In contrast, the DBT algorithm obtains optimality gaps nearly as small as the Full Space method.

One exception to the trends observed above is `pglib_opf_case162_ieee_dtc`. For this test problem, the Full Space method clearly outperforms both the Decomposed

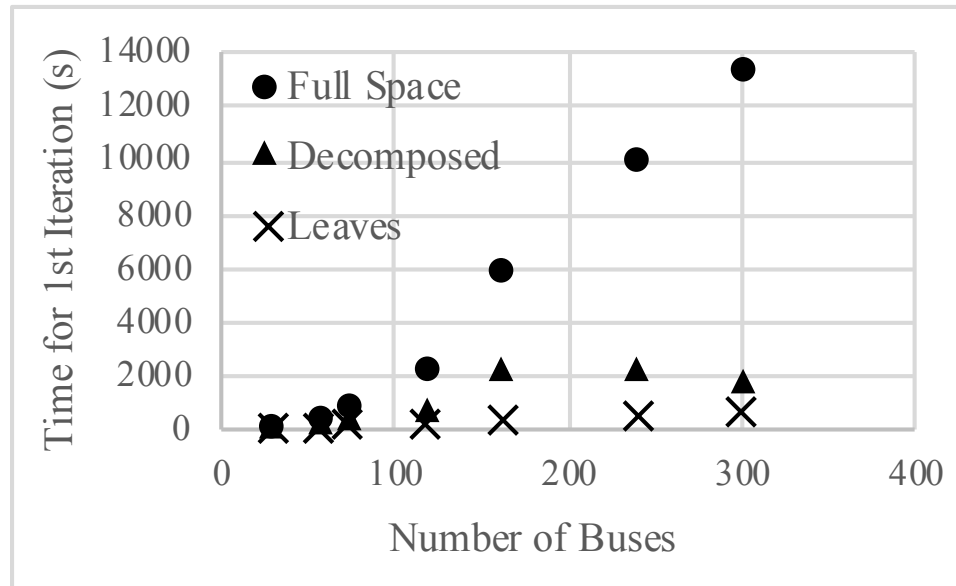


Figure 6.1.: Wallclock time for one iteration of BT

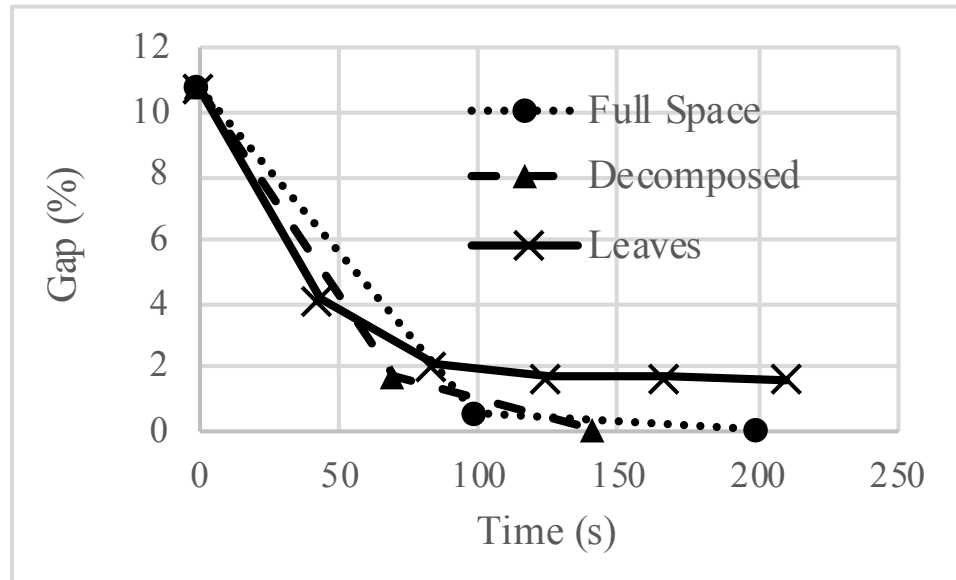


Figure 6.2.: pglib_opf_case30_ieee




Figure 6.3.: pglib_opf_case57_ieee_sad

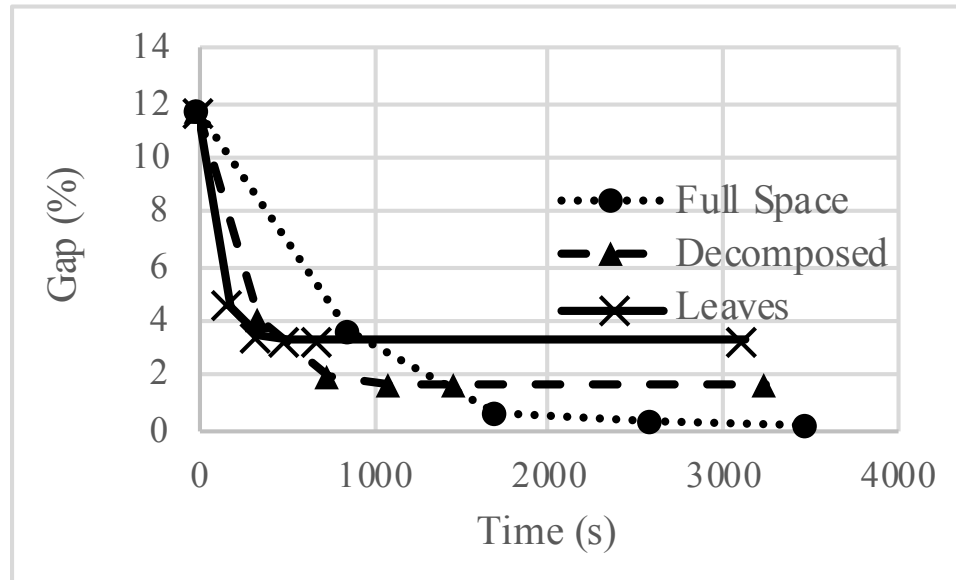


Figure 6.4.: pglb_opf_case73_ieee_rts_api




Figure 6.5.: pglb_opf_case118_ieee




Figure 6.6.: pglib_opf_case162_ieee_dtc

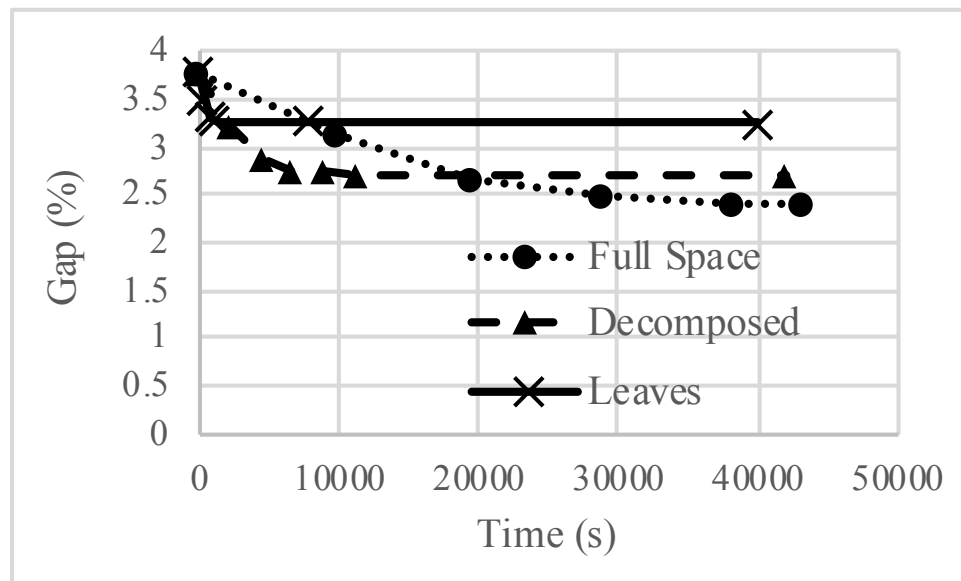


Figure 6.7.: pglib_opf_case240_pserc




Figure 6.8.: pglb_opf_case300_ieee

and Leaves methods. One possible explanation for this is presented in Table 6.1. Table 6.1 shows the number of branches required to partition the root graph for each test case. The number of branches required to partition the root graph for `pglib_opf_case162_ieee_dtc` is drastically higher than that for the other test problems. For the DBT algorithm this leads to significant increases in computation expense due to an increase in the number of large BT problems that must be solved. Additionally, an increase in the number of branches removed likely corresponds to less effective bounds tightening. This highlights the importance of effective partitioning for the DBT algorithm.

6.5 Summary

In this chapter, we proposed a decomposition algorithm for bounds tightening based on graph partitioning. The algorithm solves many small optimization problems to tighten bounds and a few large problems in order to retain feasibility information lost through the graph partitioning. Our numerical results on several instances of the ACOPF problem demonstrate that our algorithm produces variable bounds nearly as tight as those obtained with traditional optimization-based bounds tightening. Moreover, our algorithm scales much more favorably with problem size, resulting in drastically reduced computational expense.

Table 6.1.: Number of branches removed in order to partition the root graph.

Test Case	Number of Branches Removed
<code>pglib_opf_case30_ieee</code>	4
<code>pglib_opf_case57_ieee__sad</code>	5
<code>pglib_opf_case73_ieee_rts__api</code>	4
<code>pglib_opf_case118_ieee</code>	4
<code>pglib_opf_case162_ieee_dtc</code>	16
<code>pglib_opf_case240_pserc</code>	7
<code>pglib_opf_case300_ieee</code>	4

Part II

Software Tools for Water Distribution System Resilience

7. WATER NETWORK TOOL FOR RESILIENCE¹

7.1 Modeling Framework

This chapter presents a comprehensive software framework for assessing the resilience of drinking water systems to disasters, including earthquakes. The software improves upon currently available capabilities by fully integrating hydraulic and water quality simulation, damage estimates and response actions, and resilience metrics into a single platform. This software is available as an open source software package called the Water Network Tool for Resilience (WNTR, pronounced winter). In the rest of this chapter, the relevant literature is reviewed, the WNTR software framework is described in detail and then an earthquake case study is presented to demonstrate the capabilities in WNTR. While the case study focuses on earthquakes, the software framework is general and can be applied to a wide range of disruptive incidents and repair strategies.

The United States Environmental Protection Agency, in partnership with Sandia National Laboratories, developed WNTR to integrate critical aspects of resilience modeling for water distribution systems into a single software framework. WNTR provides a flexible platform for modeling both disruptive incidents and repair strategies in water distribution systems. WNTR advances previous research by making major improvements to hydraulic simulations including PDD, the ability to add and repair leaks, and the ability to change operations and/or response strategies in the middle of a simulation. Additionally, WNTR combines damage and response models into a single framework, allowing for a seamless evaluation of response and/or

¹Part of this chapter is reprinted with permission from “A Software Framework for Assessing the Resilience of Drinking Water Systems to Disasters with an Example Earthquake Case Study” by Klise, K.A., Bynum, M., Moriarty, D., and Murray, R., Environmental Modelling & Software, 2017. Note that this work was done in collaboration with both Sandia National Laboratories and the U.S. Environmental Protection Agency.

mitigation strategies to minimize overall consequences to the water utility. The primary modeling components in WNTR include: (1) disaster models (e.g., attenuation models to predict ground movement after an earthquake), (2) fragility curves used to assign the probability of damage to network components, (3) flexible controls to change the status and operation of network components, (4) models to estimate leaks in the network, (5) PDD hydraulic simulation to model the network during low pressure conditions, (6) resilience metrics to evaluate the effect of the disruption and repair strategies, and (7) the ability to perform Monte Carlo simulations.

WNTR is a Python package, which requires Python (2.7, 3.4, or 3.5) along with several Python package dependencies, including NetworkX, Pandas, Matplotlib, NumPy, and SciPy. A water network model can be constructed within WNTR or created from an EPANET formatted water network model input file. Connectivity of the water network (i.e., the network layout) is stored as a NetworkX data object [Hagberg et al., 2008]. The NetworkX package facilitates use of numerous methods to analyze the structure of complex networks. Results from hydraulic and water quality simulations are stored as Pandas data objects [McKinney, 2012], which allows for powerful time series analysis of node and link attributes, such as pressure and flow. WNTR includes high-quality graphics capabilities, including network graphics and animation, using Matplotlib [Hunter, 2007]. A dependence on NumPy and SciPy [Walt et al., 2011] enables efficient numerical computation, and working within the Python environment facilitates simulations run in loops or in parallel using standard Python methods [Python Software Foundation, 2016]. In addition to the features listed above, WNTR is compatible with all unit conventions of EPANET formatted water network model input files.

WNTR integrates these components into a single easy-to-use software platform for evaluating resilience of water distribution systems. Leveraging the object-oriented programming capabilities of Python, WNTR easily performs many complex analyses. Note that the object-oriented EPANET Python package, OOPNET, makes use of similar Python packages to facilitate scientific computing but currently only uses

EPANET for hydraulic simulations, limiting the tool to DD hydraulic simulations [Steffelbauer and Fuchs-Hanusch, 2015]. WNTR can perform DD simulations using EPANET or PDD simulations using WNTRs own hydraulic solver. The program interface is flexible and allows users to make changes to the network structure and operations, and add disruptive events and recovery actions.

7.1.1 Software Availability

WNTR can be installed through the U.S. Environmental Protection Agency GitHub organization at <https://github.com/USEPA/WNTR>. The software was made publicly available in October 2016. The GitHub site includes links to software documentation, software testing results, and contact information. The software documentation includes Python code to demonstrate the application program interface (API) and code structure. Python distributions, such as Anaconda, are recommended to help users manage the Python environment.

7.1.2 Earthquake Attenuation Models

WNTR includes attenuation models that describe how seismic waves diminish as they travel away from the epicenter. The energy associated with a seismic wave is commonly measured in terms of PGA and PGV. Attenuation models are often developed using data from a particular earthquake and depend on the local geology. Many models exist and choosing an appropriate model for a particular study area can be difficult. Infrastructure damage is associated with both PGA and PGV. Pump and tank damage is often estimated using PGA, whereas PGV is frequently used to estimate repair rate for pipes. Repair rate is defined as the number of repairs needed per km of pipe. Correction factors can be applied to the repair rate to account for variation in pipe material and soil type. The PGA attenuation models in WNTR include:

$$PGA = 403.8 \times 10^{0.265M} (R + 30)^{-1.218} \quad (7.1)$$

$$\ln(PGA) = 0.4 + 1.2M - 0.76 \ln(D) + 0.0094D \quad (7.2)$$

$$\log(PGA) = -1.83 + 0.386M - \log(R) - 0.0015R \quad (7.3)$$

where M is the earthquake magnitude (unitless), R is the distance to the epicenter (km), and D is the distance to the focus (km) [Kawashima et al., 1984, Baag et al., 1998, Lee and Cho, 2002]. These are the same PGA models used in REVAS.NET. The PGV attenuation models in WNTR include:

$$PGV = 10^{0.848+0.775M+1.834(R+17)} \quad (7.4)$$

$$PGV = 10^{-0.285+0.711M-1.85(R+17)} \quad (7.5)$$

where Eq (7.4) is for rock and Eq (7.5) is for soil [Yu and Jin, 2008]. Since attenuation models are commonly developed for a particular earthquake and site, analysts often use the average of several models to capture general trends [American Lifelines Alliance (ALA), 2001a, Yoo et al., 2015]. PGV is used to compute a pipe repair rate. The following linear and power law repair rate (RR) models from [American Lifelines Alliance (ALA), 2001a] are included in WNTR:

$$RR = 0.00187PGV \quad (7.6)$$

$$RR = 0.00108PGV^{1.173} \quad (7.7)$$

To include pipe and soil characteristics in the calculation, RR is multiplied by a correction factor [Isoyama et al., 2000], which is a function of pipe diameter, pipe material, topography, and liquefaction potential. A weight is assigned to different categories. Correction factor (C) is computed as follows:

$$C = C_d C_m C_t C_l \quad (7.8)$$

where C_d is the correction factor for pipe diameter (unitless), C_m is the correction factor for pipe material (unitless), C_t is the correction factor for topography (unit-

less), and C_l is the correction factor for liquefaction potential (unitless). By default, WNTR uses the categories and weights supplied in [Isoyama et al., 2000] but users can customize categories and weights by modifying input to the code. Additional models for PGA, PGV, RR, and C can be added by the user or included in a future release of the software. Users can also import data from external sources, such as USGS ShakeMap.

7.1.3 Fragility Curves

Fragility curves define the probability of exceeding a damage state as a function of environmental change. Fragility curves are commonly used to predict network damage after an earthquake caused by ground motion. Fragility curves are closely related to survival curves, which are used to define the probability of component failure due to age. Network components can have multiple damage states, for example, a pipe can be subject to major or minor leaks. To estimate damage caused by an earthquake, fragility curves are defined as a function of PGA, PGV, or RR. The American Lifelines Alliance reports [American Lifelines Alliance (ALA), 2001a,b] include seismic fragility curves for water system components. Damage to above ground pumps, tanks, and tunnels is often estimated using PGA, while damage to below ground pipes is often estimated using PGV and RR [American Lifelines Alliance (ALA), 2001a,b]. WNTR includes methods to define fragility curves (or survival curves) with multiple damage states using a wide array of statistical distributions. The fragility curves are then used to assign a probability of damage for each network component. The damage state is determined stochastically by selecting a uniform random variable for each network component and the water network model is adjusted accordingly to reflect the damage state that is selected. Fragility curves can also be used to estimate damage from other types of disasters, for example, fragility curves can be defined as a function of flood stage, wind speed, or temperature. FEMA's HAZUS software can also be used to estimate damage to water distribution systems from earthquakes and floods [Federal

Emergency Management Agency (FEMA), 2003a,b]. The output from HAZUS can be used in WNTR to adjust model components and reflect damage states.

7.1.4 Flexible Controls

Similar to EPANET, network controls in WNTR include time-based and conditional rules that govern how the network is operated. In EPANET, these controls can be used to change link status (open, closed, or active) or settings (pump speed or control valve setting). WNTR extends the controls available in EPANET to include a wide range of component status and settings. For example, conditional controls can be defined to simulate cascading failures where pipe leaks start when pressures exceed a critical threshold. Time-based and conditional controls can also be used to start and stop leaks and power outages or to stop the hydraulic simulation. Hydraulic simulations can also be paused while controls are changed and the simulation can then be restarted.

7.1.5 Leak Model

Leaks can cause large changes in network hydraulics. WNTR can be used to explicitly model water lost due to a leak between the time when it starts and the time when repair crews can isolate the leak by using available valves. WNTR can simulate leaks at junctions and tanks or at any location along a pipe by splitting the pipe into two sections and adding a junction connecting the two new pipe sections. Pipe breaks can be modeled by removing the existing pipe, adding two new junctions with leaks, and adding two new pipes connecting the new junctions to the network. In this case, the two new junctions are not connected to each other (preventing flow between the two junctions that were connected by the original pipe). In WNTR, leaks are modeled with a general form of the equation proposed by Crowl and Louvar [2001] where the mass flow rate of fluid through the hole, d^{leak} , is expressed as:

$$d^{leak} = C_d A \sqrt{2\rho p}^\alpha \quad (7.9)$$

where C_d is the discharge coefficient (unitless), A is area of the hole (m^2), ρ is the density of the fluid (kg/m^3), and p is the gauge pressure inside the pipe (Pa). The default discharge coefficient is 0.75 (assuming turbulent flow [Lambert, 2000]), but the user may specify other values if needed. The value of α is set to 0.5 (assuming large leaks out of steel pipes [Lambert, 2000]). As described by Lambert [2000], the primary factors affecting C_d and α are flow regime, pipe material, and orifice shape (longitudinal, round, circumferential). Additionally, both C_d and A may vary with pressure. Future development of WNTR could include extension of Eq (7.9) to account for these factors. For example, a linear model may be used to describe the area as a function of pressure as in [Cassa et al., 2010].

7.1.6 Pressure-Driven Demand Hydraulics

WNTR uses a PDD model proposed by Wagner et al. Wagner et al. [1988]:

$$d = \begin{cases} 0 & p \leq P_0 \\ D_f \sqrt{\frac{p-P_0}{P_f-P_0}} & P_0 \leq p \leq P_f \\ D_f & p \geq P_f \end{cases} \quad (7.10)$$

where d is the actual demand delivered to customers (m^3/s), D_f is the customers expected demand (m^3/s), p is the pressure (Pa), P_f is the pressure above which the customer should receive the expected demand (Pa), and P_0 is the pressure below which the customers cannot receive any water (Pa). P_f and P_0 can be defined at each node and can be modified by the user. The set of nonlinear equations describing network pressures and flows comprising the hydraulic model, including (7.10), is solved directly using a Newton-Raphson algorithm. However, Newton-Raphson algorithms do not guarantee convergence if the derivatives of the equations of interest are not Lipschitz continuous [Kelley, 2003]. The derivative of the PDD equation above is not Lipschitz continuous or even continuous; therefore, a cubic Hermite spline is applied when p is near P_0 and P_f to enforce continuity.

7.1.7 Resilience Metrics

Resilience is the ability of a system to minimize disruption and return to normal function after a disaster. Metrics that quantify resilience need to account for system function before, during, and after the disruption. Numerous resilience metrics have been suggested [US Environmental Protection Agency (USEPA), 2015c], including metrics that compute redundancy, robustness, reliability, rapidity, adaptability, and resourcefulness. For water distribution systems, these metrics generally fall into four categories: topographic, hydraulic, water quality/security, and economic. While some metrics define resilience as a single quantity, other metrics are a function of time, space, or both. For this reason, state transition plots [Barker et al., 2013], network graphics, and network animation are useful ways to visualize resilience metrics. When quantifying resilience, it is important to understand which metric best defines resilience for a particular disaster or response scenario. For example, topographic metrics are useful when studying the effect of adding redundant pipes while hydraulic metrics are useful when studying the effect of pipe breaks. Numerous resilience metrics can be computed using WNTR. These metrics include shortest path lengths [Hagberg et al., 2008], articulation points [Hagberg et al., 2008], Todini resilience index [Todini, 2000], and water service availability [Davis, 2014], among others. Some metrics can also be converted to capture statistics about the population served. For example, WNTR can be used to compute the number of people that are impacted by water service availability less than a critical threshold.

The metrics used in the case study below include water service availability and population impacted. Water service availability is computed as follows:

$$WSA_t = \left(\sum_{n \in N} V_{nt} \right) / \left(\sum_{n \in N} \hat{V}_{nt} \right) \quad (7.11)$$

where WSA_t is the water service availability of the network (unitless) at time t , N is the set of network nodes, V_{nt} is the actual water volume (m^3) received at node n at time t , \hat{V}_{nt} is the expected water volume (m^3) received at node n at time t .

Resilience metrics, like *WSA*, that depend on expected water demand can be influenced by changes in customer behavior after a disaster. Water conservation efforts, for example, change the expected demand of customers. It is important to model expected customer behavior after disasters to account for the impact this might have on resilience.

Population at each node is computed as follows:

$$pop_n = \frac{q_n}{R} \quad (7.12)$$

where pop_n is the population at node n , q_n is the average water volume consumed per day (m^3/day) at node n under normal conditions, R is the average volume of water consumed per capita per day (m^3/day). For this study, R is set to $0.76 m^3/day$ (200 gallons/day). This calculation does not account for population movement throughout the day.

The population at a given node is impacted by the disruption if they receive less than a fixed percentage of their expected water volume. Population impacted is computed as follows:

$$PI_t = \sum_{n \in N} pop_n \delta_{nt} \quad (7.13)$$

where

$$\delta_{nt} = \begin{cases} 1 & \text{if } V_{nt}/\hat{V}_{nt} < \tau \\ 0 & \text{otherwise} \end{cases} \quad (7.14)$$

where PI_t is the population impacted at time t , d_{nt} is a binary variable that is set to 1 if V_{nt}/\hat{V}_{nt} is less than a threshold, τ and set to zero otherwise. The threshold is set to 80% in the case study below.

7.1.8 Monte Carlo Simulation

To address the uncertainty associated with modeling potential disasters, Monte Carlo techniques can be used to simulate multiple realizations or outcomes of a given scenario. WNTR has the capability to allow characteristics of a disaster scenario to be drawn from statistical distributions. As mentioned above, fragility curves provide probabilities of the location and severity of a disruption. The duration of the disruption can also be drawn from statistical distributions. Distributions can be a function of component properties (e.g., pipe age, material, or joint type), environmental change (e.g., PGA, PVG, or RR), or available resources (e.g., number of repair crews available). WNTR is compatible with many statistical distributions and random selection methods that can be used for stochastic simulation. WNTR also includes the ability to pause the hydraulic simulation, change network operations, and then restart the simulation. The water network model and simulation results can also be saved to files and reloaded for future analysis. These features are helpful when evaluating various response action plans and when simulating long periods with different time resolution. Standard Python tools can also be used to run simulations in parallel [Python Software Foundation, 2016].

7.2 Earthquake Case Study

The following case study demonstrates multiple capabilities of WNTR that can be used to evaluate disruptive incidents and repair strategies after an earthquake. Figure 7.1 shows the network model of a real water distribution system that is used for the case study [Watson et al., 2009]. The system serves approximately 152,000 customers and has one reservoir, two valves, 34 tanks, 61 pumps, 3,323 nodes, and 3,829 pipes. The pipes are made of cast iron, ductile case iron, and polyethylene and the soil type in the region has various degrees of liquefaction potential (Figure 7.1). The spatial distribution of these properties was used to determine pipe fragility, as described below.

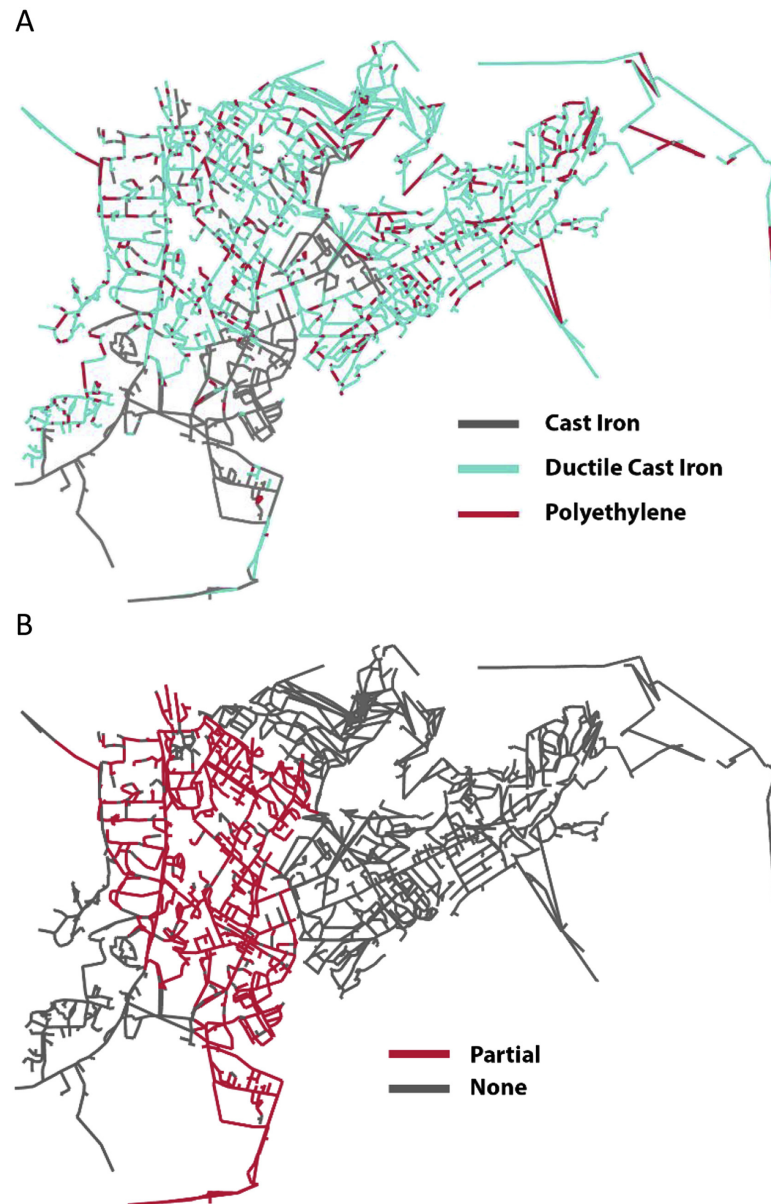


Figure 7.1.: Water distribution network model showing A) pipe material and B) liquefaction potential.

An earthquake was assumed to have occurred at 6:00 AM along an N-S fault that parallels a valley on the west side of the network (Figure 7.2). In the model, this was 24 hours after the start of the simulation. The earthquake epicenter was assumed to have occurred at the south, central, or north part of the city, with magnitudes of 6.0, 6.5, or 7.0. The hydraulic response of the network was simulated for 14 days after the earthquake. PDD hydraulic simulations were run using a minimum and nominal pressure of zero and 25 psi, respectively. In the simulations, ground movement from the earthquake damaged pipes and tanks, pumps stopped operating due to either direct damage or power outages, and damaged natural gas pipelines cause fires. Given the disorder that ensues after an earthquake, especially after large earthquakes, it was assumed that repairing network components and fighting fires was not possible for the first 12 hours after the earthquake occurs (note that 12 hours might be too long in some cases and this value can be adjusted). After that period, fire crews were scheduled to fight fires and repair crews were scheduled to fix damaged components in the network. Three repair strategies including water conservation and seismic-resistant pipes were simulated. Network resilience was quantified using two metrics, water service availability and population impacted, computed using Eqs. (7.11) and (7.13), respectively. To incorporate uncertainty in outcomes, 50 realizations of each of the 27 scenarios were run (three earthquake locations, three earthquake magnitudes, and three repair strategies).

7.2.1 Earthquake Damage

For this case study, network component damage includes tank and pipe leaks, and pump failure. For each tank, pump, and pipe, fragility curves were used to define the probability of damage as shown in Figure 7.3. For tanks, the fragility curves in Figure 7.3A included four damage states (DS): DS1 resulted in a leak with diameter 0.05 m, 0.25 m for DS2, 0.5 m for DS3, and 1 m for DS4. For pumps, there was a single damage state, DS1, associated with pump shut off (Figure 7.3B). For pipes, there were two damage states: DS1 resulted in a minor leak for which the leak diameter




Figure 7.2.: Water distribution network model showing the fault line (dashed red line), node elevation, and epicenter of three earthquakes. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

was drawn from a uniform distribution with a minimum of 0.01 m and a maximum of 0.05 m, while DS2 represented a major leak with a diameter from 0.05 m to 0.15 m (Figure 7.3C). Leaks were added to the midpoint of the pipe, and were constrained to not to exceed 75% of the pipe diameter.

For each scenario and realization, PGA, PGV, and RR was computed for each network component. Here, PGA was computed using the average of Eqs. (7.1), (7.2), and (7.3). As an example, Figure 7.4A shows the spatial distribution of PGA for a magnitude 6.5 earthquake at the central location. PGV was computed using the average of Eqs. (7.4) and (7.5). RR was computed using Eq. (7.6) from PGV and a correction factor that was computed using Eq. (7.8) with the default weights defined in [Isoyama et al., 2000].

The damage state for each component was then determined stochastically based on the PGA or RR calculated for the component and a uniform random variable selected for that component. The fragility curve incorporates both of these values: PGA or RR determines the x-axis value while the random variable determines the y-axis value. The intersection point of these values determines the damage state for a particular component for a single realization. The water network model was then adjusted to reflect the selected damage: if a pump was damaged, a control was added that changed its status to off; new nodes were added to the network model to represent leaks. As an example, the location of pipe and tank damage resulting from a magnitude 6.5 earthquake at the central location is shown in Figure 7.4B.

While damage to other network components such as reservoirs, turbines, and treatment plants were not included in this study, similar methods could be used to define the probability of damage to these components and the network could be modified accordingly.

Fires are important to include when modeling the effects of an earthquake because water utilities are responsible for maintaining enough capacity to fight fires. For this case study, the locations of fires were randomly distributed throughout the network. Typical firefighting demands and durations [International Code Council (ICC), 2012]

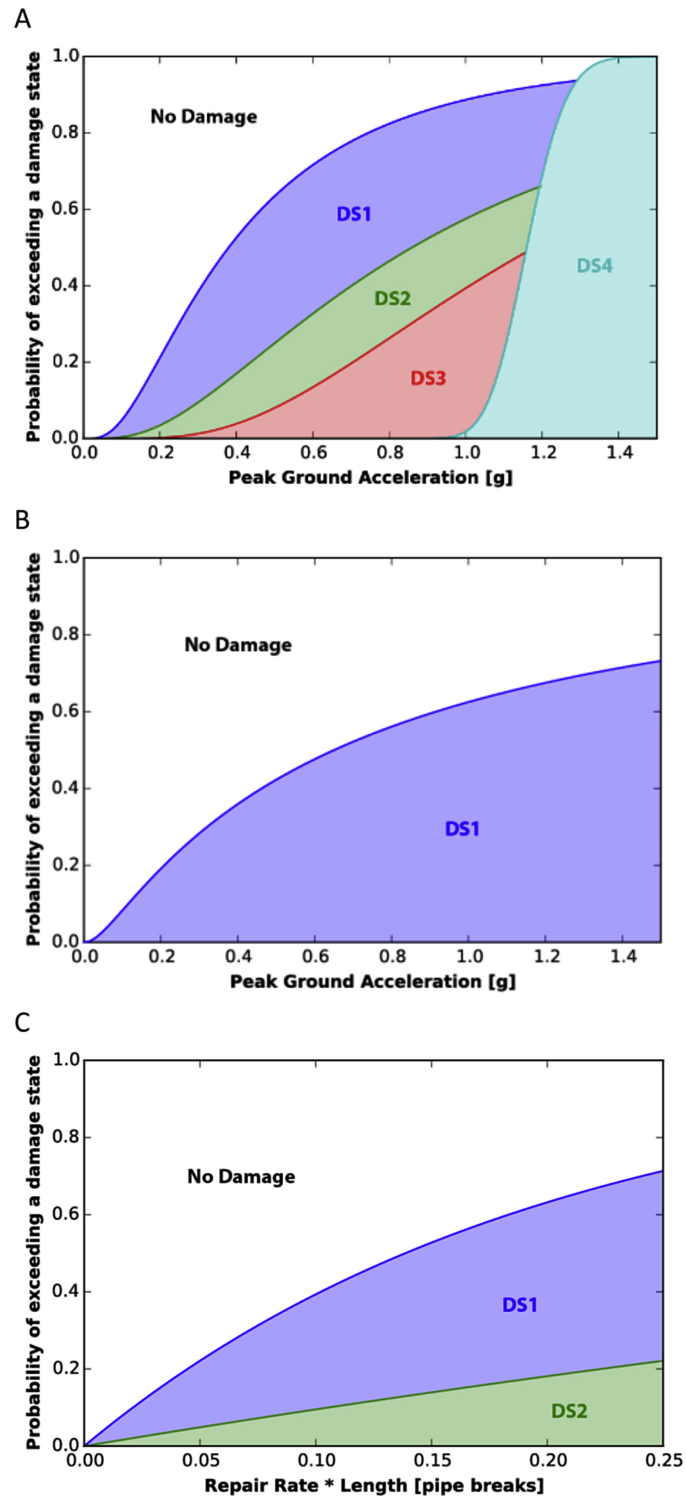


Figure 7.3.: Fragility curves for A) tank, B) pump, and C) pipe damage.

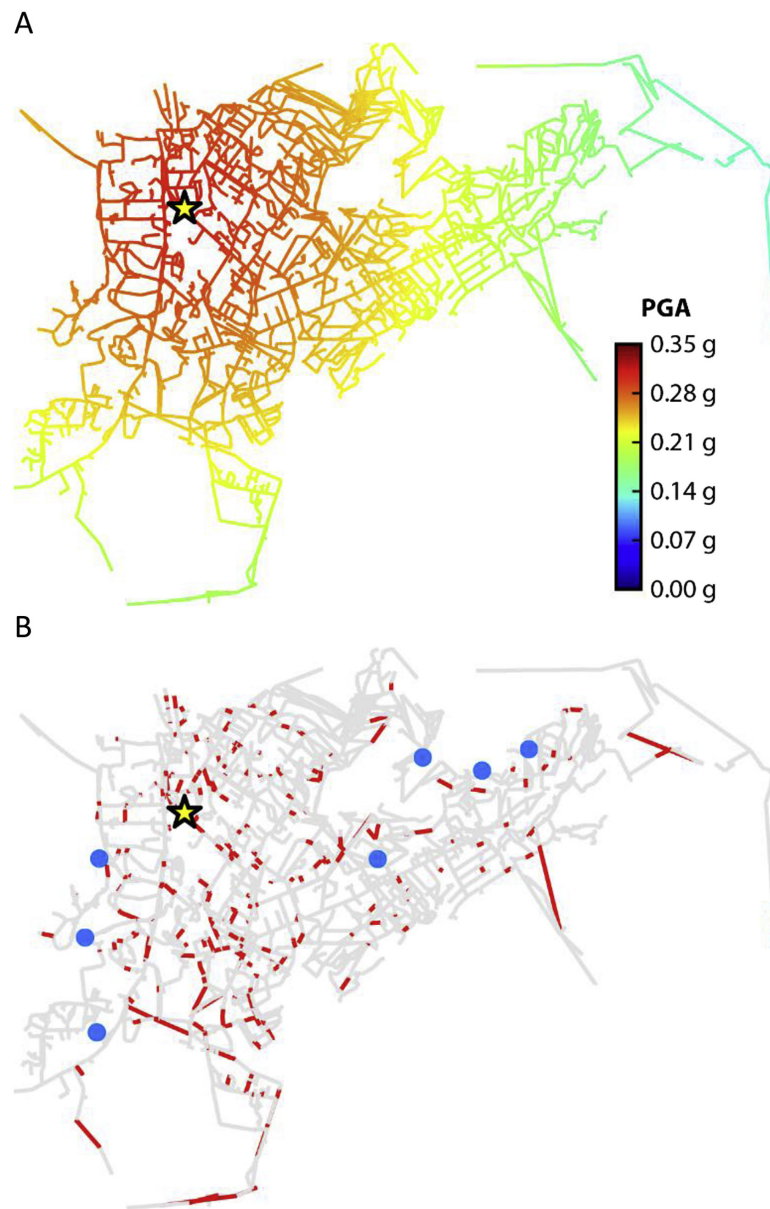


Figure 7.4.: A) PGA and B) location of pipe damage (in red) and tank damage (in blue) for a magnitude of 6.5 earthquake at the central epicenter location. The yellow star indicates the location of the central epicenter. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

were added to the network model to simulate the water withdrawal needed to fight one fire every 12 hours over the full duration of the simulation. For each fire, demand was selected from a uniform distribution with a minimum of 1500 gallons/min and a maximum of 8000 gallons/min. Duration was selected from a uniform distribution with a minimum of two hours and a maximum of four hours.

7.2.2 Repair Strategies

Three repair strategies were considered in this case study. For all strategies, repairs began 12 hours after the earthquake. There were five repair crews to fix pipe damage, two repair crews to fix tank damage, one repair crew to fix pump damage, and one firefighting crew. In the first repair strategy (RS1), customer expected demand continued at their normal rate after the earthquake. In the second repair strategy (RS2), customer expected demand was reduced by 40% for 14 days after the earthquake. This was done to simulate the impact of water conservation after the disaster. In the third strategy (RS3), the utility retrofitted pipes near the fault zone to be seismic-resistant and customer expected demand continued at their normal rate after the earthquake. In this case, pipes in the partial liquefaction zone (Figure 7.1), did not sustain damage after the earthquake.

The pipe and tank repair crews fixed one leak every 12 hours. Within the first six hours, the leaks were isolated using valves at the nearest junctions, and the leak was fixed within the next six hours. After the leak was repaired, the pipe was opened. To stagger repairs, the exact time of isolation and repair were drawn from a uniform distribution within the first or second six-hour period. Repairs were prioritized based on the largest cumulative leak volume at the time of each repair. As leak rates can increase or decrease through time as the pressure fluctuates, this prioritization changed throughout the simulation. Within WNTR, the simulation was paused every 12 hours, the cumulative leak rates were computed, and the largest five pipe leaks and largest two tank leaks were scheduled for repair.

The single pump repair crew fixed one pump every eight hours. In this network, the reservoir is near the lowest point in the network, and pumps are critical to deliver water to customers in the upper sections of the network. For this reason, pump repairs were prioritized based on pipe distance to the reservoir, with the closest pump fixed first. Since this prioritization does not include feedback from the simulation, the simulation does not have to be paused to schedule pump repairs.

For this case study, there was one fire crew to fight a fire every 12 hours, and the fires were fought in a random order. If the pressure was not sufficient to pull the required water volume, the hydraulic simulation continued to try to extract the expected demand for the duration of the fire. The difference between the expected and actual demand was recorded for each fire to measure water service availability for firefighting.

The earthquake damage and repair strategies used in this case study were based on reasonable assumptions. These parameters can all be adjusted within WNTR to customize analysis for a specific water utility and disaster scenario.

7.2.3 Case Study Results

Figure 7.5 shows the results of a single realization of a magnitude 6.5 earthquake at the central location with RS1. For this realization, 239 pipes were damaged, 7 tanks were damaged, and 14 pumps lost power. Figure 7.5 shows the leak rates on all 239 damaged pipes and the tank levels for all 34 tanks as a function of time. As Figure 7.5A shows, the leak rates were high until the pipes were repaired. However, as the network was repaired and pressures increased, new leaks were sometimes discovered and subsequently repaired, as shown with the leak highlighted in red in Figure 7.5A. That particular leak was fixed 9.1 days after the earthquake. Figure 7.5B shows how tank levels (from both damaged and non-damaged tanks) varied as components were repaired. In that figure, tank levels from damaged tanks are highlighted in red. Once the tank was repaired, the tank level is shown in gray. Figure 7.5B also shows that

some tanks that were not damaged by the earthquake can drain due to damage at other locations in the network.

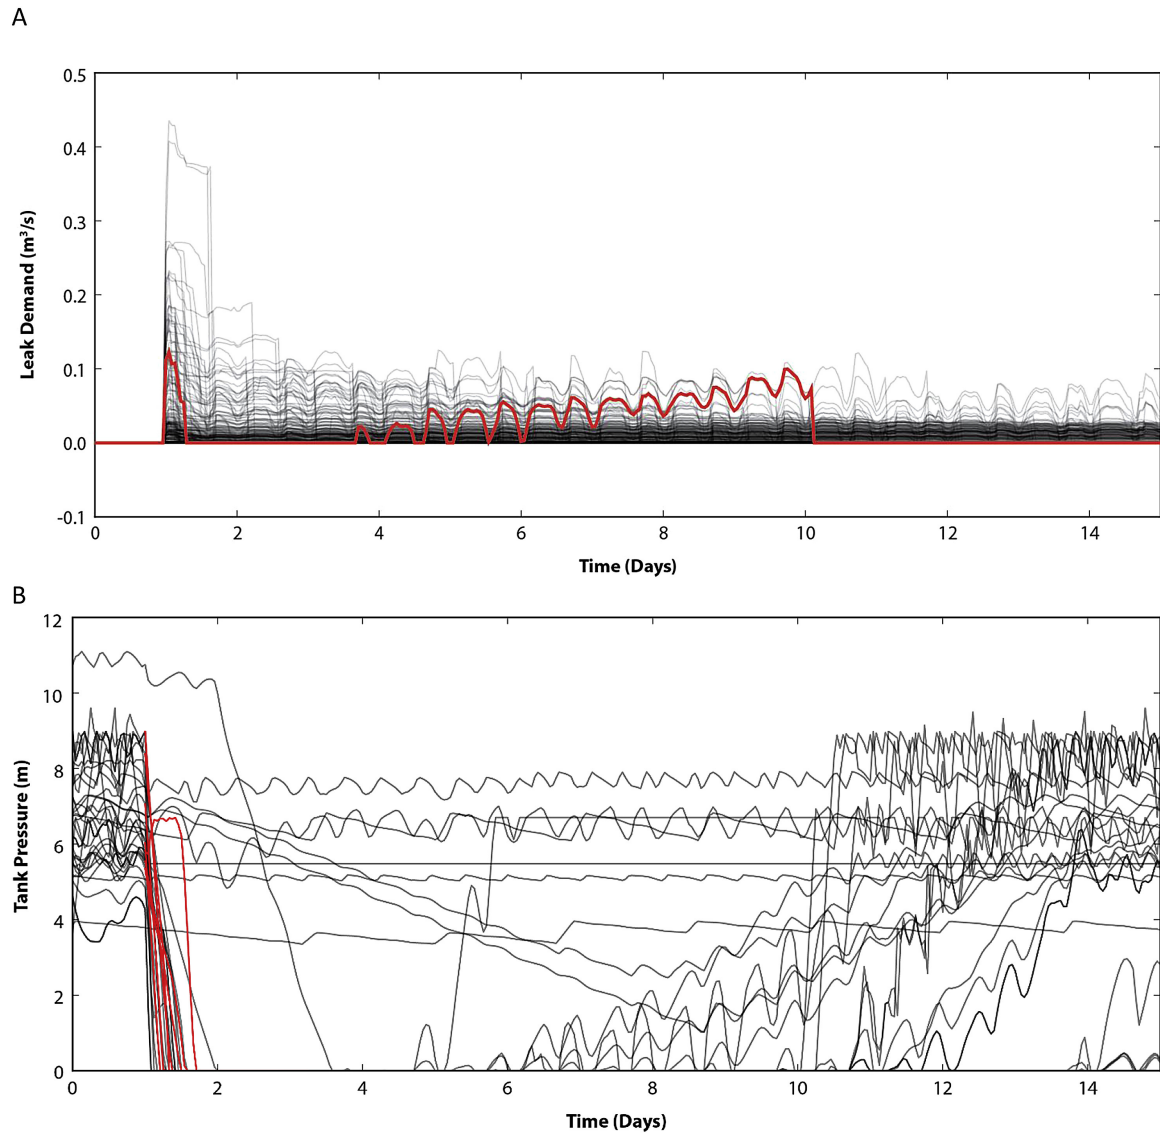


Figure 7.5.: (A) Leak rate and (B) tank level as a function of time for a single realization of a magnitude 6.5 earthquake at the central location using no conservation (RS1). In A), a particular leak is shown in red illustrating the possibility of a leak becoming more prominent as the system is restored. In B), damaged tanks are shown in red up until the time they are repaired, tanks are refilled as the system is restored. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Figure 7.6, Figure 7.7, and Figure 7.8A show the results for 50 realizations of the same earthquake magnitude (6.5), location (central), and repair strategy (RS1). For each realization, different network components were damaged as defined by the probability distribution functions in the fragility curves and by different fire demands assigned to different locations. Similarly, pipes and tanks were repaired in a different order according to the prioritization for leak repair and fires were fought in a random order. Figure 7.6 shows the water service availability of the network over time, reflecting the proportion of customer expected demand that was actually delivered for the entire network. As the customer expected demand vary on a diurnal cycle, so does the water service availability. A day after the earthquake, water service availability varied from about 45% to more than 60%, showing the potential range in outcomes for the same scenario. The median value (in black) reached a low of 54% and slowly returned to full service at 6.58 days after the earthquake. The median value was used here to compute a recovery time, defined as the time when the system could deliver 90% of the pre-earthquake water volume to customers.

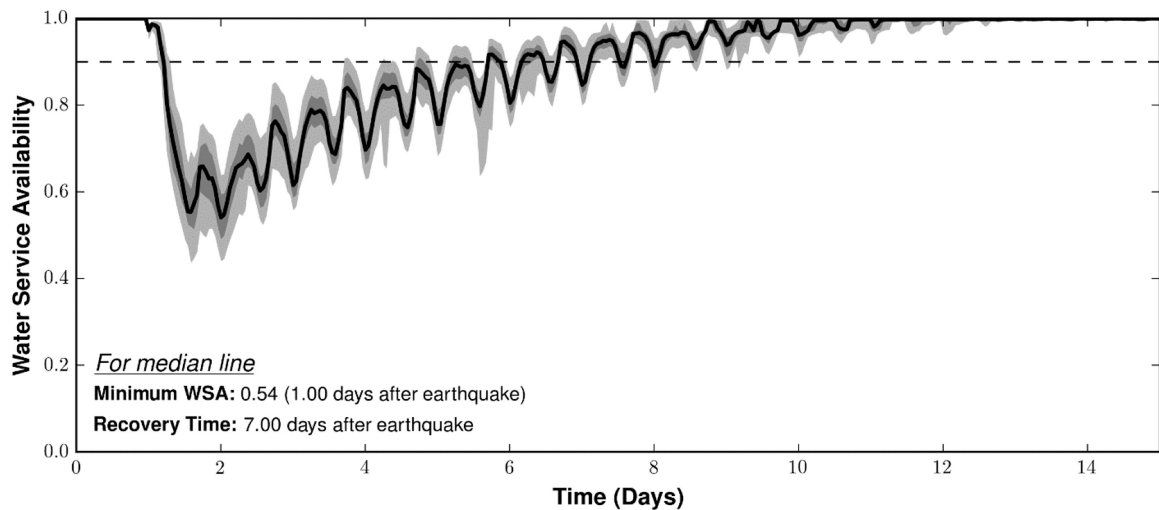


Figure 7.6.: Water service availability for a magnitude 6.5 earthquake at the central location using no conservation (RS1). Light gray is the 5th-95th percentile, dark gray is the 25th-75th percentile, and the black line is the median. Threshold for recovery time shown as a dotted line.




Figure 7.7.: Firefighting capacity for a magnitude 6.5 earthquake at the central location using no conservation (RS1). Light gray is the 5th-95th percentile, dark gray is the 25th-75th percentile, and the black line is the median. Firefighting capacity is computed every 12 h (indicated by x's) at the location of a particular fire.

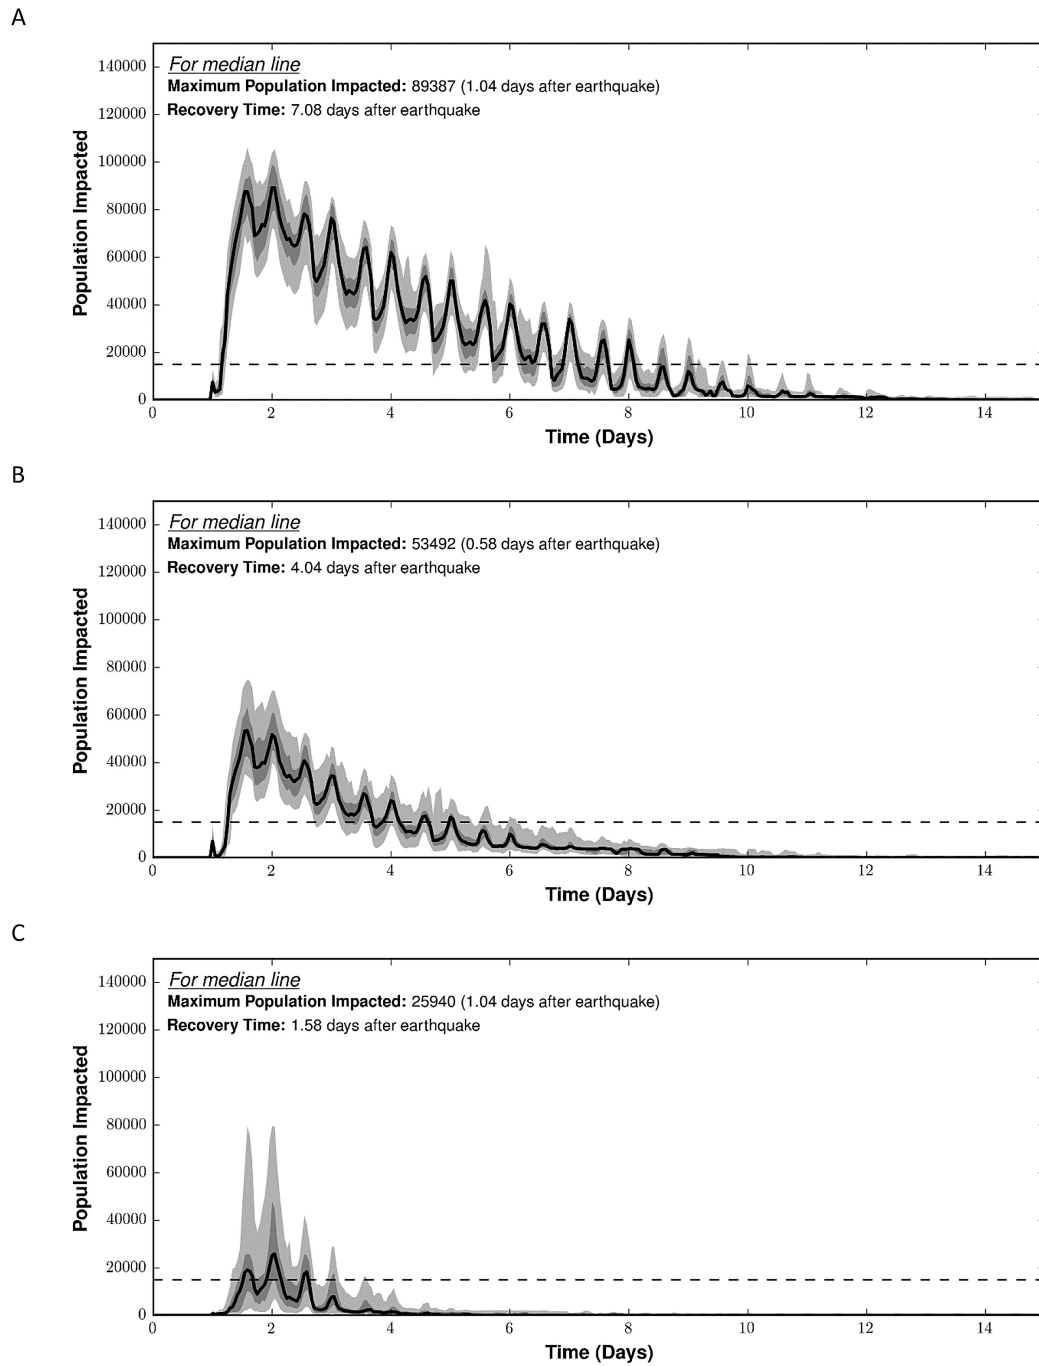


Figure 7.8.: Population impacted (A) without conservation (RS1), (B) with conservation (RS2), and (C) with seismic-resistant pipes (RS3) for 50 realizations of a magnitude 6.5 earthquake at the central location. Light gray is the 5th-95th percentile, dark gray is the 25th-75th percentile, and the black line is the median. Threshold for recovery time shown as a dotted line.

Figure 7.7 shows the firefighting capacity for the same 50 realizations. Firefighting capacity is defined as the water service availability at the nodes where fire demand was added to the system. The quantity is computed at the same frequency as the fires occur, which is 12 hours in this case. The results show that firefighting capacity varies widely and slowly recovered over the 14 day simulation. After one day, firefighting capacity varied from about 15%-85% for the 25th to 75th percentile scenarios, but ranged from about 75%-100% by the end of the 14 days.

Figure 7.8 shows the population impacted for each of the 50 realizations. With no conservation efforts (RS1) between 70,000 and 100,000 people were impacted a day after the earthquake, receiving less than 80% of their expected water volume. Based on the median value of the 50 realizations, almost 90,000 people were impacted. The recovery time was defined as the time when less than 10% of the population was impacted. For RS1, it took 7.08 days to meet this level. Figure 7.8B shows the same results when water conservation efforts were in place (RS2) immediately following the earthquake. This shows that the population impacted was much lower with conservation strategies in place. Based on the median value, approximately 53,000 people were impacted, and the time to recovery was reduced to 4.04 days. Figure 7.8C shows results when pipes were retrofitted to be seismic-resistant near the fault line. These results show dramatic reduction in the number of people impacted and the time to recovery, median result indicate that less than 26,000 people were impacted and the time to recovery was just 1.58 days after the earthquake.

Table 7.1 and Table 7.2 summarize the minimum water service availability, maximum population impacted, and time to recovery for each magnitude, location, and repair strategy. These values were computed using the median value for each scenario over all 50 realizations. As expected, water service availability decreased and population impacted increased as the earthquake magnitude increases. Unsurprisingly, the location of the earthquake also had an impact on water service availability and population impacted. Earthquakes at the northern location, further from the city, had less of an impact than earthquakes occurring in the city. Given RS1 conditions, earth-

quakes of magnitudes 6.5 or greater at the south and central locations impacted the network almost equally. To further illustrate this point, Figure 7.9 shows the distribution of maximum population impacted over all 50 realizations when the earthquake was located at the north, central, and southern epicenter and RS1 is used. This comparison shows that while location can have a large influence on population impacted, the results can vary widely depending on the specific network damage and response strategy. Table 7.1 and Table 7.2 also illustrate that both water conservation and earthquake resistant pipes in the liquefaction area helped to increase overall water service availability and decrease population impacted. Under most scenarios the use of earthquake resistant pipes was slightly more effective than water conservation, however, earthquake resistant pipes were significantly better for earthquakes occurring the central location near the retrofitted pipes.




Figure 7.9.: Effect of location of the earthquake on maximum population impacted for a 6.5 magnitude earthquake at the central location and no conservation (RS1).

Table 7.1.: Minimum water service availability (WSA) and recovery time for each earthquake magnitude, location, and repair strategy. Values are computed using the median of the 50 realizations. A recovery time of 0 means that the median value never went below 90% water service availability.

Repair Strategy	Magnitude	South Location		Central Location		North Location	
		Minimum WSA	Recovery Time in Days	Minimum WSA	Recovery Time in Days	Minimum WSA	Recovery Time in Days
RS1 (Pipe, tank, pump repair & firefighting)	6.0	0.88	1.04	0.85	1.08	0.98	0.00
	6.5	0.53	6.04	0.54	7.00	0.65	4.04
	7.0	0.28	14.00+	0.27	14.00+	0.35	14.00+
RS2 (RS1 with 40% water conservation)	6.0	0.98	0.00	0.97	0.00	1.00	0.00
	6.5	0.66	3.08	0.67	4.00	0.79	1.58
	7.0	0.37	14.00+	0.35	14.00+	0.45	9.04
RS3 (RS1 with seismic-resistant pipes)	6.0	1.00	0.00	1.00	0.00	1.00	0.00
	6.5	0.83	1.08	0.87	1.08	0.93	0.00
	7.0	0.45	6.04	0.50	6.04	0.60	4.04

Table 7.2.: Maximum population impacted and recovery time for each earthquake magnitude, location, and repair strategy. Values are computed using the median of the 50 realizations. A recovery time of 0 means that the median value never went above 10% population impacted.

Repair Strategy	Magnitude	South Location		Central Location		North Location	
		Max Pop. Impacted	Recovery Time in Days	Max Pop. Impacted	Recovery Time in Days	Max Pop. Impacted	Recovery Time in Days
RS1 (Pipe, tank, pump repair & firefighting)	6.0	27,142	1.04	33,902	1.08	2,548	0.00
	6.5	90,736	6.04	89,387	7.00	71,520	4.04
	7.0	127,398	14.00+	129,232	14.00+	118,908	14.00+
RS2 (RS1 with 40% water conservation)	6.0	4,378	0.00	4,260	0.00	302	0.00
	6.5	57,167	3.08	53,492	4.00	32,260	1.58
	7.0	104,194	14.00+	105,814	14.00+	89,560	9.04
RS3 (RS1 with seismic-resistant pipes)	6.0	551	0.00	1,060	0.00	646	0.00
	6.5	36,984	1.08	25,940	1.08	13,971	0.00
	7.0	104,518	6.04	93,146	6.04	73,782	4.04

Overall, these results indicate that this network could expect to withstand a magnitude 6.0 earthquake with only 2-15% reduction in water service availability, and less than 34,000 customers impacted. In contrast, a magnitude 7.0 earthquake could be expected to affect more than 85% of customers, result in a minimum water service availability of 0.27-0.35, and require more than 14 days to recover. This case study is intended to illustrate how water utilities could use WNTR to evaluate disruption and resilience enhancing actions following an earthquake. Similar analysis could be used to evaluate other types of disruptions and utility specific data could be used to validate the model.

7.3 Summary

In this chapter, we presented an open-source software package, WNTR, for evaluating and improving the resilience of water distribution systems. Our key contributions are listed below.

- WNTR integrates hydraulic and water quality simulation, a wide range of damage and response options, and resilience metrics into a single framework, allowing for a seamless evaluation of water network resilience.
- WNTR can help water utilities estimate potential damages, understand how damage to infrastructure would occur over time, evaluate preparedness strategies, prioritize response actions, and identify worst-case scenarios, efficient repair strategies, and best practices for maintenance and operations.
- We presented an earthquake case study to demonstrate how WNTR can be used to perform a resilience analysis from start to end. The case study included estimating physical damage to the water distribution system components, evaluation of hydraulic performance, and implementation and comparison of repair strategies.

8. SUMMARY

Critical infrastructure resilience is vital for both the economy and public safety and health. In this thesis, we focus on developing effective models, algorithms, and software for evaluating and improving the resilience of power systems and water distribution systems. In this chapter, we summarize the major contributions of this work and propose future work.

8.1 Thesis Summary and Contributions

8.1.1 Stochastic Programming Formulations for Improving Power Systems Resilience to Extreme Weather

Chapter 3, focuses on stochastic programming formulations for improving power systems resilience to extreme weather events. We first propose a stochastic linear programming problem which seeks to dispatch generators proactively in order to minimize the expected load shed given a set of probabilistic scenarios involving transmission line failures. We then extend the formulation to consider hardening investments (introducing discrete decisions) and perform out-of-sample cross validation to demonstrate the robustness of the formulations. Finally, we consider the role chemical process facilities, as large industrial consumers, can play in electric grid resilience through demand response.

Our proposed stochastic programming models for enhancing power systems resilience to extreme weather events show promise, based on results obtained using synthetic and real world test systems. Specifically, relatively simple operational changes can significantly reduce the impact of extreme weather events. For the test cases considered in this work, redispatching generators before a severe weather event can reduce the expected load shed by as much as 25%. Transmission line hardening and

increasing transmission capacity can also be very effective, with higher marginal gains from the initial investments. Hardening just one transmission line in case2383wp can reduce the expected load shed by over 10%, and doubling the capacity of just 5 lines can reduce the expected load shed by approximately 4.4%.

Our out-of-sample cross validation demonstrated that high quality solutions can be obtained with very small fractions of all possible outage scenarios. Specifically, we solved the stochastic resilience problem with one set of scenarios and evaluated the solution on an independent set of “true” scenarios. For case30, only 100 scenarios were needed to obtain solutions nearly identical to those obtained when solving the problem directly with the “true” scenarios. With an average of 7 transmission line outages per scenario, this is less than 0.001% of all possible scenarios.

By extending the stochastic programming framework to estimate the value of demand response capabilities we were able to consider the role chemical process facilities can play in electric grid resilience. We formulated a stochastic MILP that seeks to determine the minimum number of hardened transmission lines required to ensure feasible operation across a set of weather-related failure scenarios. For all three test cases considered, we found that a modest amount of demand response could significantly reduce the capital investment needed to mitigate the effects of extreme weather events. On the largest test case, the addition of only two demand response contracts (each requiring only a 25% drop in consumed load) is sufficient ensure feasible operation across all scenarios while reducing the required number of hardened lines (e.g., the required capital investment) by 21%.

8.1.2 The Effect of the DC Power Flow Approximation

All of the results presented in Chapter 3 were obtained with an MILP formulation based on the linear DC approximation for the alternating current (AC) power flow equations. While this approximation is widely used in power systems analysis, it can underestimate the impact of thermal limits and other system constraints. In Chapter

4, we investigate the effect this approximation has on the solution of our resilience formulations.

First, we formulate and solve the demand response problem presented in section 3.6 with the actual nonlinear AC power flow equations. Because this formulation (an MINLP) cannot be efficiently solved for large-scale problems, we compare the results from the DC and AC formulations on a set of smaller test cases. We use a tailored multi-tree approach that solves a sequence of MISOCP relaxations of the MINLP to find candidate integer solutions and NLP subproblems to check feasibility. Using this approach, results with the smaller test cases show that both the DC approximation and the nonlinear AC formulation have similar trends, indicating that there is value in the use of demand response for improved resilience. Furthermore, they also find similar optimal locations for demand response contracts. For case `pglib_opf_case14_ieee`, every bus selected for demand response by the DC model was also selected for demand response by the AC model (cumulative across all values of Δ_b and N_z). Only one bus was selected for demand response by the AC model and not the DC model.

We also evaluated each integer solution obtained with the DC model on the AC model to compute the fraction of AC-feasible scenarios. We found that the linear model often underestimates the amount of demand response or hardened lines necessary to ensure resilience of the AC system, demonstrating a need for improved scalability of algorithms for solving MINLPs based on AC power flow equations.

8.1.3 Relaxations and Refinement Techniques for AC Constrained Power Flow Problems

Many algorithms for the solution of mixed-integer nonlinear programming problems require the global solution of nonlinear programming subproblems. Therefore, Chapter 5 focuses on relaxations and refinement techniques for AC constrained power flow problems. We demonstrate that, by properly incorporating reference bus constraints, relaxations of the rectangular form of the ACOPF problem can significantly strengthen existing relaxations. We also illustrate that the effectiveness of

optimization-based bounds tightening (OBBT) at a given iteration is directly related to the distance of the corresponding bus from the reference bus. Ultimately, we are able to reduce the optimality gap to less than 0.1% on all but 5 NESTA test cases with up to 300 buses.

8.1.4 Decomposed Bounds Tightening

Although OBBT is extremely effective for tightening ACOPF relaxations, the computational expense of OBBT grows quickly with the number of buses in the network. Therefore, we propose a decomposed bounds tightening algorithm based on graph partitioning. Our numerical results demonstrate that our algorithm produces variable bounds nearly as tight as those obtained with traditional OBBT but with a significantly reduced computational expense, especially for large problems.

8.1.5 Water Network Tool for Resilience: A Python Package

Drinking water systems are subject to a wide range of hazardous incidents that can disrupt service to customers and damage critical infrastructure. Chapter 7 introduces WNTR, a new open source Python package developed in collaboration with Sandia National Laboratories and the U.S. Environmental Protection Agency. WNTR is designed to help water utilities and researchers investigate resilience of water distribution systems to a wide range of hazardous scenarios and to evaluate resilience-enhancing actions. Unlike current modeling tools, WNTR integrates hydraulic and water quality simulation, a wide range of damage and response options, and resilience metrics into a single framework, allowing for a seamless evaluation of water network resilience. The software can help water utilities estimate potential damages, understand how damage to infrastructure would occur over time, evaluate preparedness strategies, prioritize response actions, and identify worse case scenarios, efficient repair strategies, and best practices for maintenance and operations.

While a few modeling tools have been developed that focus on evaluating the resilience of water networks after an earthquake, WNTR was designed to simulate a wide range of disaster scenarios, including earthquakes, contamination incidents, and power outages. The software includes methods to predict infrastructure damage using fragility or survival functions, and other approaches can be added using the flexible Python environment. WNTR allows the user to change features of the network model to reflect damage, and initiates leaks in pipes or tanks as necessary. By simulating hydraulics using PDD the effects of pipe breaks or damaged tanks on system pressures and delivered water volumes is more accurately predicted. Moreover, WNTR allows for simulation to continue even when parts of the water system are isolated and unable to provide water to customers. In this way, WNTR is able to estimate the performance of a drinking water system during extreme failures which is a significant advancement from current modeling tools. WNTR also provides more flexibility by allowing the user to change system controls and model components mid-simulation.

In this work, several response strategies were demonstrated including fixing pipe leaks, repairing tanks, restoring power to pumps, fighting fires, and implementing conservation strategies. WNTR can also be used to simulate the response to contamination events by, for example, flushing contaminated water from hydrants or instituting do-not-drink orders. Mitigation strategies, such as installing online monitoring early warning systems, can also be modeled within WNTR. Optimization of response and mitigation strategies is an area of ongoing research that could be added to WNTR in the future.

WNTR contains a wide variety of metrics that can serve as indicators of resilience, including topographic, hydraulic, and water quality measures, some of which are measured at the node or pipe level and some at the system-level. In this work, two metrics were used: water service availability and population impacted. These are system-level metrics reflecting the ability of the entire water system to withstand and recover from disasters. These types of metrics can be used to compare a water systems resilience across different types of disasters, and to measure improvements

in resilience based upon newly implemented mitigation actions or planned response strategies. In this context, it is important to consider how different system functions recover (i.e. firefighting capacity, water availability, water quality) when customers are asked to conserve resources. Further research is needed to assess the most useful resilience metrics for water systems and to fully understand how systems-level metrics embody resilience characteristics.

8.2 Recommendations for Future Work

Designing and retrofitting critical infrastructure systems to be resilient while minimizing costs remains a challenging problem. This dissertation was dedicated to modeling and solution methods for this problem. However, we also developed the following recommendations for future work.

8.2.1 Stochastic Programming Formulations for Improving Power Systems Resilience to Extreme Weather

We have seen that several resilience enhancement strategies can, independently, significantly improve power systems resilience to extreme weather events. However, integrating several strategies and optimizing over them simultaneously presents both data and modeling challenges. First, more accurate cost models are needed for the various hardening approaches. Additionally, as more resilience enhancement strategies are integrated into a single stochastic MILP, the number of binary variables increases, increasing the need for formulations with tight continuous relaxations.

There are many resilience enhancement strategies that were not considered in this dissertation. Such actions include transmission switching (i.e., network reconfiguration), investments in hardening current substations and generators, and development of new facilities including storage and microgrids. For example, recent studies have demonstrated that microgrids can prevent cascading outages (Phase I) [Chen et al., 2010], mitigate performance degradation (Phase II) [Wang and Wang, 2015, Liu et al., 2017b], and improve restoration (Phase III) [Castillo, 2013].

Furthermore, the size of the stochastic programming problem increases with the product of the network size and the number of scenarios, making the solution of large scale systems computationally challenging. When the extensive form cannot be solved directly in the time allotted, stage-based or scenario-based decomposition methods such as progressive hedging [Rockafellar and Wets, 1991] or Benders decomposition [Benders, 1962] can be applied.

8.2.2 Solution Techniques for ACPF-based MINLP's

The results in Chapter 4 demonstrated the need for scalable algorithms for solving MINLP's based on the AC power flow equations. The results in Chapter 5, along with other literature, demonstrate the effectiveness of optimization based bounds tightening (OBBT) for refining relaxations of the ACOPF problem. However, traditional OBBT is a very expensive refinement technique, especially for large problems. Thus, in Chapter 6, we proposed a decomposed bounds tightening (DBT) algorithm which is far more efficient and nearly as effective as traditional OBBT. Despite this, there is a great deal of room for improvement. First, bounds tightening alone does not guarantee convergence. The DBT algorithm needs to be integrated into a branch and bound or an outer-approximation algorithm with finite convergence guarantees. Additionally, the DBT algorithm should be integrated with other domain reduction such as feasibility-based bounds tightening which is more efficient but typically less effective. Integrating these domain reduction techniques in effective ways is a challenging problem.

8.2.3 Water Distribution System Resilience

In Chapter 7, we presented a Python package, WNTR, for evaluating and improving the resilience of water distribution systems (WDS). However, WNTR does not do any optimization to design or retrofit WDS's for improved resilience. Utilizing stochastic optimization techniques (as was done for power systems resilience) is challenging

because WDS's typically do not operate at steady state. Thus, dynamics must be introduced into the optimal resilience formulations. Additional complications involve discrete switches that are part of typical WDS operation. For example, pumps may be turned on if a tank level gets too low or off if the suction pressure gets too low. Finally, there is no well-established linear approximation to the hydraulic equations in WDS models (such as the linear DC approximation for power systems). Thus, determining how to improve WDS resilience is a challenging problem. Research needs to be done to develop models of the appropriate fidelity for obtaining robust resilience enhancement strategies.

REFERENCES

REFERENCES

- H.A. Aalami, M. Parsa Moghaddam, and G.R. Yousefi. Demand response modeling considering interruptible/curtailable loads and capacity market programs. *Applied Energy*, 87:243–250, 2010.
- M.M. Adibi and L.H. Fink. Power system restoration planning. *IEEE Transactions on Power Systems*, 9(1):22–28, 1994a.
- MM Adibi and LH Fink. Power system restoration planning. *IEEE Transactions on Power Systems*, 9(1):22–28, 1994b.
- Natalia Alguacil, Andrs Delgadillo, and Jos M. Arroyo. A trilevel programming approach for electric grid defense planning. *Computers & Operations Research*, 41:282 – 290, 2014. ISSN 0305-0548. doi: <https://doi.org/10.1016/j.cor.2013.06.009>. URL <http://www.sciencedirect.com/science/article/pii/S0305054813001664>.
- American Lifelines Alliance (ALA). Seismic fragility formulations for water systems, part 1. *American Lifelines Alliance*, 2001a.
- American Lifelines Alliance (ALA). Seismic fragility formulations for water systems, part 2. *American Lifelines Alliance*, 2001b.
- American Lifelines Alliance (ALA). Seismic guidelines for water pipelines. *American Lifelines Alliance*, 2005.
- American National Standards Institute (ANSI). Risk analysis and management for critical asset protection (ramcap) standard for risk and resilience management of water and wastewater systems. *American Water Works Association, Denver, CO*, 2010.
- S. Ashok and R. Banerjee. An optimization mode for industrial load management. *IEEE Transactions on Power Systems*, 16(4):879–884, 2001.
- Chang-Eob Baag, Sung-Joon Chang, Nam-Dae Jo, and Jin-Soo Shin. Evaluation of seismic hazard in the southern part of korea. In *Proceedings of 2nd International Symposium on Seismic Hazards and Ground Motion in the Region of Moderate Seismicity*, pages 31–50, 1998.
- R. Baldick, B. Chowdhury, I. Dobson, Zhaoyang Dong, Bei Gou, D. Hawkins, H. Huang, M. Joung, D. Kirschen, Fangxing Li, Juan Li, Zuyi Li, Chen-Ching Liu, L. Mili, S. Miller, R. Podmore, K. Schneider, Kai Sun, D. Wang, Zhigang Wu, Pei Zhang, Wenjie Zhang, and Xiaoping Zhang. Initial review of methods for cascading failure analysis in electric power transmission systems iee pes cams task force on understanding, prediction, mitigation and restoration of cascading failures. In *2008 IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century*, pages 1–8, July 2008. doi: 10.1109/PES.2008.4596430.

Kash Barker, Jose Emmanuel Ramirez-Marquez, and Claudio M Rocco. Resilience-based network component importance measures. *Reliability Engineering & System Safety*, 117:89–97, 2013.

Pietro Belotti, Sonia Cafieri, Jon Lee, and Leo Liberti. Feasibility-based bounds tightening via fixed points. In *International Conference on Combinatorial Optimization and Applications*, pages 65–76. Springer, 2010.

Pietro Belotti, Christian Kirches, Sven Leyffer, Jeff Linderoth, James Luedtke, and Ashutosh Mahajan. Mixed-integer nonlinear optimization. *Acta Numerica*, 22:1–131, 2013.

J. F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4(1):238–252, Dec 1962. ISSN 0945-3245. doi: 10.1007/BF01386316. URL <https://doi.org/10.1007/BF01386316>.

D. Bienstock and S. Mattia. Using mixed-integer programming to solve power grid blackout problems. *Discrete Optimization*, 4:115–141, 2007.

Daniel Bienstock and Gonzalo Munoz. On linear relaxations of opf problems. *arXiv preprint arXiv:1411.1120*, 2014.

Yang Bingyuan, Wu Jiguang, and Liu Junyong. Research on pricing model of congestion management considering influence of partitioned interruptible load. *Power System Technology*, 29(9):41–45, 2005.

Michael R Bussieck, Arne Stolbjerg Drud, and Alexander Meeraus. Minlplib collection of test models for mixed-integer nonlinear programming. *INFORMS Journal on Computing*, 15(1):114–119, 2003.

Michael Bynum, Anya Castillo, Jean-Paul Watson, and Carl D. Laird. Strengthened SOCP relaxations for ACOPF with mccormick envelopes and bounds tightening. In *13th International Symposium on Process Systems Engineering (PSE 2018)*, volume 44 of *Computer Aided Chemical Engineering*, pages 1555 – 1560. Elsevier, 2018a. doi: <https://doi.org/10.1016/B978-0-444-64241-7.50254-8>.

Michael Bynum, Anya Castillo, Jean-Paul Watson, and Carl D Laird. Tightening mccormick relaxations toward global solution of the acopf problem. *IEEE Transactions on Power Systems*, 2018b.

Mary B Cain, Richard P Oneill, and Anya Castillo. History of optimal power flow and formulations. Technical report, Technical report, US FERC, 2012.

A. Caprara and M. Locatelli. Global optimization problems and domain reduction strategies. *Mathematical Programming*, (125):123–137, 2010.

Miguel Carrión, José Manuel Arroyo, and Natalia Alguacil. Vulnerability-constrained transmission expansion planning: A stochastic programming approach. *IEEE Transactions on Power Systems*, 22(4):1436–1445, 2007.

AM Cassa, JE Van Zyl, and RF Laubscher. A numerical investigation into the effect of pressure on holes and cracks in water supply pipes. *Urban Water Journal*, 7(2): 109–120, 2010.

A. Castillo. Microgrid provision of blackstart in disaster recovery for power system restoration. In *2013 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 534–539, Oct 2013. doi: 10.1109/SmartGridComm.2013.6688013.

A. Castillo. Risk analysis and management in power outage and restoration: A literature survey. *Electric Power Systems Research*, 107:9–15, 2014.

Chen Chen, Alper Atamtürk, and Shmuel S Oren. Bound tightening for the alternating current optimal power flow problem. *IEEE Trans. Power Syst.*, 31(5): 3729–3736, 2016.

Qiming Chen, Kun Zhu, and J. D. McCalley. Dynamic decision-event trees for rapid response to unfolding events in bulk transmission systems. In *2001 IEEE Porto Power Tech Proceedings (Cat. No.01EX502)*, volume 2, pages 5 pp. vol.2–, 2001. doi: 10.1109/PTC.2001.964752.

X. Chen, H. Dinh, and B. Wang. Cascading failures in smart grid - benefits of distributed generation. In *2010 First IEEE International Conference on Smart Grid Communications*, pages 73–78, Oct 2010. doi: 10.1109/SMARTGRID.2010.5622022.

Carleton Coffrin, Dan Gordon, and Paul Scott. NESTA, the NICTA energy system test case archive. *CoRR*, abs/1411.0359, 2014.

Carleton Coffrin, Hassan L Hijazi, and Pascal Van Hentenryck. Strengthening convex relaxations with bound tightening for power network optimization. In *International Conference on Principles and Practice of Constraint Programming*, pages 39–57. Springer, 2015.

Critical Infrastructure Partnership Advisory Council (CIPAC). All-hazard consequence management planning for the water sector: Preparedness, emergency response, and recovery cipac workgroup. *US Department of Homeland Security, Washington DC*, 2009.

Daniel A Crowl and Joseph F Louvar. *Chemical process safety: fundamentals with applications*. Pearson Education, 2001.

Robert J Dakin. A tree-search algorithm for mixed integer programming problems. *The computer journal*, 8(3):250–255, 1965.

Craig A Davis. Water system service categories, post-earthquake interaction, and restoration strategies. *Earthquake Spectra*, 30(4):1487–1509, 2014.

Claudia D Ambrosio, Antonio Frangioni, Leo Liberti, and Andrea Lodi. On interval-subgradient and no-good cuts. *Operations Research Letters*, 38(5):341–345, 2010.

John Eidinger and Craig A. Davis. Recent earthquakes: Implications for u.s. water utilities. *Water Research Foundation, Report #4408, Denver, CO*, 2012.

Executive Office of the President. Economic benefits of increasing electric grid resilience to weather outages. Technical report, Presidents Council of Economic Advisers, U.S. Department of Energys Office of Electricity Delivery and Energy Reliability, with assistance from the White House Office of Science and Technology, 2013.

Neng Fan, David Izraelevitz, Feng Pan, Panos M Pardalos, and Jianhui Wang. A mixed integer programming approach for optimal power grid intentional islanding. *Energy Systems*, 3(1):77–93, 2012.

Federal Emergency Management Agency (FEMA). Multi-hazard loss estimation methodology, earthquake model, hazus-mh 2.1, technical manual. *Washington, DC, USA: Federal Emergency Management Agency*, 2003a.

Federal Emergency Management Agency (FEMA). Multi-hazard loss estimation methodology, flood model, hazus-mh, technical manual. *Washington, DC, USA: Federal Emergency Management Agency*, 2003b.

Roberto Guidotti, Hana Chmielewski, Vipin Unnikrishnan, Paolo Gardoni, Therese McAllister, and John van de Lindt. Modeling the resilience of critical infrastructure: The role of network dependencies. *Sustainable and resilient infrastructure*, 1(3-4): 153–168, 2016.

S. Guikema, R. Davidson, and H. Liu. Statistical models of the effects of tree trimming on power system outages. *IEEE Transactions on Power Delivery*, 21(3): 1549–1557, 2006.

Seth D Guikema, Steven M Quiring, and Seung-Ryong Han. Prestorm estimation of hurricane damage to electric power distribution systems. *Risk analysis*, 30(12): 1744–1752, 2010.

Inc. Gurobi Optimization. Gurobi optimizer reference manual, 2016. URL <http://www.gurobi.com>.

Aric Hagberg, Pieter Swart, and Daniel S Chult. Exploring network structure, dynamics, and function using networkx. In *7th Python in Science Conference (SciPy2008)*, Pasadena, CA, 2008.

William E Hart, Carl D Laird, Jean-Paul Watson, David L Woodruff, Gabriel A Hackebeil, Bethany L Nicholson, and John D Sirola. *Pyomo - optimization modeling in python*, volume 67. Springer, 2017.

Carmen Baskette Henrikson and Kristin Brief. Designing a successful demand response program: Its not your grandfathers load control program. Technical report, Technical report, ACEE Summer Study on Energy Efficiency in Buildings, 2008.

Hassan Hijazi, Carleton Coffrin, and Pascal Van Hentenryck. Convex quadratic relaxations for mixed-integer nonlinear programs in power systems. *Mathematical Programming Computation*, 9(3):321–367, 2017.

P. Hines, J. Apt, and S Talukdar. Trends in the history of large blackouts in the united states. In *2008 IEEE Power Energy Soc. General Meeting-Convers. Del. Electr. Energy 21st Century*, pages 1–8, 2008.

Patricia Hoffman, William Bryan, Alic Lippert, Mindi Farber-DeAnda, Matthew Cleaver, Carleen Lewandowski, and Kateri Young. Hardening and resiliency: U.s. energy industry response to recent hurricane seasons. Technical report, Technical report, US DOE, 2010.

HSL. A collection of fortran codes for large scale scientific computation. URL <http://www.hsl.rl.ac.uk/>.

HSL. A collection of Fortran codes for large scale scientific computation. 2007. URL <http://www.hsl.rl.ac.uk>.

John D Hunter. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(3):90–95, 2007.

IEEE PES Task Force on Benchmarks for Validation of Emerging Power System Algorithms. Power grid lib - optimal power flow. <https://github.com/power-grid-lib/pglib-opf>, 2017.

IBM ILOG. *IBM ILOG CPLEX V12.1: User's manual for CPLEX*. 01 2009.

International Code Council (ICC). International fire code, appendix b - fire-flow requirements for buildings. *International Code Council*, ISBN: 978-1-60983-046-5, 2012.

XJ Irias, W Cain, Y Prashar, and R McMullin. Rapid modeling of seismic damage to water infrastructure. In *7th US-Japan Workshop on Water System Seismic Practices in Niigata, Japan*, 2011.

Ryoji Isoyama, Eisuke Ishida, Kiyoji Yune, and Toru Shirozu. Seismic damage estimation procedure for water supply pipelines. *Water supply*, 18(3):63–68, 2000.

R.A Jabr. A conic quadratic format for the load flow equations of meshed networks. *IEEE Trans. on Power Syst.*, 22(4):2285–2286, 2007.

Rabih A Jabr. Radial distribution load flow using conic programming. *IEEE transactions on power systems*, 21(3):1458–1459, 2006.

Kazuhiko Kawashima, Koh Aizawa, and Kazuyuki Takahashi. Attenuation of peak ground motion and absolute acceleration response spectra. In *Proceedings, Eighth World Conference on Earthquake Engineering*, volume 2, pages 257–264, 1984.

Carl T Kelley. *Solving nonlinear equations with Newton's method*, volume 1. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2003.

B. Kocuk, S. S. Dey, and X. A. Sun. New formulation and strong misocp relaxations for ac optimal transmission switching problem. *IEEE Transactions on Power Systems*, 32(6):4161–4170, Nov 2017a. ISSN 0885-8950. doi: 10.1109/TPWRS.2017.2666718.

Burak Kocuk, Santanu S Dey, and X Andy Sun. Strong socp relaxations for the optimal power flow problem. *Operations Research*, 64(6):1177–1196, 2016.

Burak Kocuk, Santanu S Dey, and X Andy Sun. Matrix minor reformulation and socp-based spatial branch-and-cut method for the ac optimal power flow problem. *arXiv preprint arXiv:1703.03050*, 2017b.

Allan Lambert. What do we know about pressure-leakage relationships in distribution systems. In *IWA Conference: Systems Approach to Leakage Control and Water Distribution System Management*, 2000.

Ailsa H Land and Alison G Doig. An automatic method of solving discrete programming problems. *Econometrica: Journal of the Econometric Society*, pages 497–520, 1960.

Daniele Laucelli, Luigi Berardi, and Orazio Giustolisi. Assessing climate change and asset deterioration impacts on water distribution networks: Demand-driven or pressure-driven network modeling? *Environmental Modelling & Software*, 37:206–216, 2012.

K Lee and KH Cho. Attenuation of peak horizontal acceleration in the sino-korea craton. In *Proceedings of the Annual Fall Conference of Earthquake Engineering Society of Korea, Cheonan, Korea*, pages 3–10, 2002.

D. Lindemeyer, H.W. Dommel, and M.M. Adibi. Power system restorationa bibliographical survey. *Int. J. Electr. Power Energy Syst.*, 23(3):219–227, 2001.

Jianfeng Liu, Michael Bynum, Anya Castillo, Jean-Paul Watson, and Carl D Laird. A multitree approach for global solution of acopf problems using piecewise outer approximations. *Computers & Chemical Engineering*, 2017a.

Jianfeng Liu, Anya Castillo, Jean-Paul Watson, and Carl D. Laird. Global solution strategies for the network-constrained unit commitment problem with ac transmission constraints. *Under Revision, IEEE Transactions on Power Systems*, 2018.

Xindong Liu, Mohammad Shahidehpour, Zuyi Li, Xuan Liu, Yijia Cao, and Zhao-hong Bie. Microgrids for enhancing the power grid resilience in extreme conditions. *IEEE Transactions on Smart Grid*, 8(2):589–597, 2017b.

Mowen Lu, Harsha Nagarajan, Russell Bent, Sandra D Eksioglu, and Scott J Mason. Tight piecewise convex relaxations for global optimization of optimal power flow. *arXiv preprint arXiv:1803.04633*, 2018.

A Mani, M Tabesh, and MR Zolfaghari. Hydraulic performance of post-earthquake water distribution networks based on head driven simulation method. *Water Science & Technology*, 13(5):1281, 2013.

I Markov, Mircea Grigoriu, and TD O’Rourke. An evaluation of seismic serviceability of water supply networks with application to the san francisco auxiliary supply system. In *Technical Report NCEER*, number 94-0001. 1994.

Wes McKinney. *Python for data analysis: Data wrangling with Pandas, NumPy, and IPython*. 2012.

S. Mitra, I.E. Grossmann, J.M. Pinto, and N. Arora. Optimal production planning under time-sensitive electricity prices for continuous power-intensive processes. *Computers & Chemical Engineering*, 38:171–184, 2012.

S. Mitra, L. Sun, and I.E. Grossmann. Optimal scheduling of industrial combined heat and power plants under time-sensitive electricity prices. *Energy*, 54:194–211, 2013.

J. A. Momoh. *Electric power system applications of optimization*. Markel Dekker, 2001.

Ramon E Moore, R Baker Kearfott, and Michael J Cloud. *Introduction to interval analysis*, volume 110. Siam, 2009.

J Muranho, A Ferreira, J Sousa, A Gomes, and A Sá Marques. Pressure-dependent demand and leakage modelling with an epanet extension–waternetgen. *Procedia Engineering*, 89:632–639, 2014.

João Muranho, Ana Ferreira, Joaquim Sousa, Abel Gomes, and Alfeu Sá Marques. Waternetgen: an epanet extension for automatic water distribution network models generation and pipe sizing. *Water science and technology: water supply*, 12(1): 117–123, 2012.

Yu Na, Furong Wei, and Li Ya. Optimization model of purchasing interruptible load for network congestion constrained dispatch during the peak load periods. In *Power and Energy Engineering Conference*, pages 496–500. SciRes, 2010.

National Academy of Sciences (NAS). Disaster resilience: A national imperative. *The National Academies Press, Washington, DC*, 2012. prepared by the NAS Committee on Science, Engineering, and Public Policy.

Min Ouyang and Leonardo Duenas-Osorio. Multi-dimensional hurricane resilience assessment of electric power systems. *Structural Safety*, 48:15–24, 2014.

Mathaios Panteli and Pierluigi Mancarella. Influence of extreme weather and climate change on the resilience of power systems: Impacts and possible mitigation strategies. *Electric Power Systems Research*, 127:259–270, 2015.

Mathaios Panteli, Cassandra Pickering, Sean Wilkinson, Richard Dawson, and Pierluigi Mancarella. Power system resilience to extreme weather: Fragility modeling, probabilistic impact assessment, and adaptation measures. *IEEE Transactions on Power Systems*, 32(5):3747–3757, 2017a.

Mathaios Panteli, Dimitris N Trakas, Pierluigi Mancarella, and Nikos D Hatziairgiyriou. Power systems resilience assessment: hardening and smart operational enhancement strategies. *Proceedings of the IEEE*, 105(7):1202–1213, 2017b.

Assela Pathirana. Epanet2 desktop application for pressure driven demand modeling. In *Water Distribution Systems Analysis 2010*, pages 65–74. 2010.

Benjamin L. Preston, Scott N. Backhaus, Mary Ewers, Julia A. Phillips, Cesar A. Silva-Monroy, Jeffrey E. Dagle, Alfonso G. Tarditi, John Looney, and Thomas J. King, Jr. Resilience of the U.S. electricity system: A multi-hazard perspective. Technical report, Technical report, US DOE, 2016.

Yash Puranik and Nikolaos V Sahinidis. Domain reduction techniques for global nlp and minlp optimization. *Constraints*, 22(3):338–376, 2017.

Python Software Foundation. Python, 2016. URL <https://www.python.org/>.

R. T. Rockafellar and Roger J.-B. Wets. Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research*, 16(1):119–147, 1991. doi: 10.1287/moor.16.1.119. URL <https://doi.org/10.1287/moor.16.1.119>.

Natalia Romero, TD O’Rourke, LK Nozick, and CA Davis. Seismic hazards and water supply performance. *Journal of Earthquake Engineering*, 14(7):1022–1043, 2010.

Lewis A Rossman. Epanet 2: Users manual. *US Environmental Protection Agency, EPA 600/R00/057, Cincinnati, OH*, 2000.

Hong S Ryoo and Nikolaos V Sahinidis. Global optimization of nonconvex nlps and minlps with applications in process design. *Computers & Chemical Engineering*, 19(5):551–566, 1995.

J. Salmeron, K. Wood, and R. Baldick. Analysis of electric grid security under terrorist threat. *IEEE Transactions on Power Systems*, 19(2):905–912, May 2004. ISSN 0885-8950. doi: 10.1109/TPWRS.2004.825888.

Chengcheng Shao, Mohammad Shahidehpour, Xifan Wang, Xiuli Wang, and Biyang Wang. Integrated planning of electricity and natural gas transportation systems for enhancing the power grid resilience. *IEEE Transactions on Power Systems*, 32(6):4418–4429, 2017.

Peixin Shi and Thomas D O’Rourke. Seismic response modeling of water supply systems. *National Center for Earthquake Engineering Research, Technical Report MCEER-08-0016, Buffalo, NY*, 2008.

Hongbiao Song and Mladen Kezunovic. A new analysis method for early detection and prevention of cascading events. *Electric Power Systems Research*, 77(8):1132 – 1142, 2007. ISSN 0378-7796. doi: <https://doi.org/10.1016/j.epsr.2006.09.010>. URL <http://www.sciencedirect.com/science/article/pii/S0378779606002343>.

David Steffelbauer and Daniela Fuchs-Hanusch. Oopnet: An object-oriented epanet in python. *Procedia Engineering*, 119:710–718, 2015.

Wei Sun, Chen-Ching Liu, and Li Zhang. Optimal generator start-up strategy for bulk power system restoration. *IEEE Transactions on Power Systems*, 26(3):1357–1366, 2011.

Taronne HP Tabucchi. Modeling post-earthquake restoration of the los angeles water supply system. *M.S. Thesis submitted to Cornell University, Ithaca, NY*, 2007.

J. C. Tan, P. A. Crossley, P. G. McLaren, P. F. Gale, I. Hall, and J. Farrell. Application of a wide area backup protection expert system to prevent cascading outages. *IEEE Transactions on Power Delivery*, 17(2):375–380, Apr 2002. ISSN 0885-8977. doi: 10.1109/61.997902.

Mohit Tawarmalani and Nikolaos V Sahinidis. Global optimization of mixed-integer nonlinear programs: A theoretical and computational study. *Mathematical programming*, 99(3):563–591, 2004.

The White House. Presidential policy directive – critical infrastructure security and resilience. <https://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil>, February 2013. Accessed: 2016-09-19.

Ezio Todini. Looped water distribution networks design using a resilience index based heuristic approach. *Urban water*, 2(2):115–122, 2000.

Nemanja Trifunovic. *Pattern Recognition for Reliability Assessment of Water Distribution Networks*. 2012.

U.S. Department of Energy. Quadrennial energy review: Energy transmission, storage, and distribution infrastructure. Technical report, U.S. Department of Energy, 2015.

US Environmental Protection Agency (USEPA). Flood resilience: a basic guide for water and wastewater utilities. *US Environmental Protection Agency, Office of Water, EPA 817/B-14/006, Washington DC*, 2014.

US Environmental Protection Agency (USEPA). Planning for an emergency drinking water supply. *US Environmental Protection Agency, Office of Research and Development, National Homeland Security Research Center, EPA 600/R-11/054*, 2015a.

US Environmental Protection Agency (USEPA). Power resilience guide for water and wastewater utilities. *US Environmental Protection Agency, Office of Water, EPA 600/R15/004, Washington DC*, 2015b.

US Environmental Protection Agency (USEPA). Systems measures of water distribution system resilience. *US Environmental Protection Agency, Office of Research and Development, EPA 600/R 14/383, Washington DC*, 2015c.

US Environmental Protection Agency (USEPA). Drought response and recovery: A basic guide for water utilities. *US Environmental Protection Agency, Office of Water, EPA 810/B-16/001., Washington DC*, 2016.

Marianna Vaiman, P Hines, J Jiang, S Norris, M Papic, A Pitto, Y Wang, and G Zweigle. Mitigation and prevention of cascading outages: Methodologies and practical applications. In *Power and Energy Society General Meeting (PES), 2013 IEEE*, pages 1–5. IEEE, 2013.

R. Vujanic, S. Mariethos, P. Goulart, and M. Morari. Robust integer optimization and scheduling problems for large electricity consumers. In *Proceedings of the 2012 American Control Conference*, pages 3108–3113, 2012.

Andreas Wächter and Lorenz T Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical programming*, 106(1):25–57, 2006.

Janet M Wagner, Uri Shamir, and David H Marks. Water distribution reliability: simulation methods. *Journal of water resources planning and management*, 114(3): 276–294, 1988.

David J Wald, Bruce C Worden, Vincent Quitoriano, and Kris L Pankow. Shakemap manual: technical manual, user’s guide, and software guide. Technical report, 2006.

Stéfan van der Walt, S Chris Colbert, and Gael Varoquaux. The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011.

Q. Wang, J. P. Watson, and Y. Guan. Two-stage robust optimization for n-k contingency-constrained unit commitment. *IEEE Transactions on Power Systems*, 28(3):2366–2375, Aug 2013. ISSN 0885-8950. doi: 10.1109/TPWRS.2013.2244619.

X. Wang and V. Vittal. System islanding using minimal cutsets with minimum net flow. In *IEEE PES Power Systems Conference and Exposition, 2004.*, pages 379–384 vol.1, Oct 2004. doi: 10.1109/PSCE.2004.1397589.

Yezhou Wang, Chen Chen, Jianhui Wang, and Ross Baldick. Research on resilience of power systems under natural disasters—a review. *IEEE Transactions on Power Systems*, 31(2):1604–1613, 2016a.

Yezhou Wang, Chen Chen, Jianhui Wang, and Ross Baldick. Research on resilience of power systems under natural disasters? a review. *IEEE Transactions on Power Systems*, 31(2):1604–1613, 2016b.

Yifei Wang, Mohammad Shahidehpour, Loi Lei Lai, LPLP Huang, HLHL Yuan, and Fang Yuan Xu. Resilience-constrained hourly unit commitment in electricity grids. *IEEE Transactions on Power Systems*, 2018.

Zhaoyu Wang and Jianhui Wang. Self-healing resilient distribution systems based on sectionalization into microgrids. *IEEE Transactions on Power Systems*, 30(6): 3139–3149, 2015.

Jean-Paul Watson, Regan Murray, and William E Hart. Formulation and optimization of robust sensor placement problems for drinking water contamination warning systems. *Journal of Infrastructure Systems*, 15(4):330–339, 2009.

Jean-Paul Watson, Ross Guttromson, Cesar Silva-Monroy, Robert Jeffers, Katherine Jones, James Ellison, Charles Rath, Jared Gearhart, Dean Jones, Tom Corbet, et al. Conceptual framework for developing resilience metrics for the electricity, oil, and gas sectors in the united states. *Sandia National Laboratories, Albuquerque, NM (United States), Tech. Rep*, 2014.

Xuan Wu and Antonio J Conejo. An efficient tri-level optimization model for electric grid defense planning. *IEEE Transactions on Power Systems*, 32(4):2984–2994, 2017.

Zheng Yi Wu, Rong He Wang, Thomas M Walski, Shao Yu Yang, Daniel Bowdler, and Christopher C Baggett. Efficient pressure dependent demand model for large water distribution system analysis. In *Water Distribution Systems Analysis Symposium 2006*, pages 1–15, 2008.

Do Guen Yoo, Donghwi Jung, Doosun Kang, Joong Hoon Kim, and Kevin Lansey. Seismic hazard assessment model for urban water supply networks. *Journal of Water Resources Planning and Management*, 142(2):04015055, 2015.

YX Yu and CY Jin. Empirical peak ground velocity attenuation relations based on digital broadband records. In *The 14th world conference on earthquake engineering*. Beijing, China, 2008.

Majid Zare, Ali Abbaspour, Mahmud Fotuhi-Firuzabad, and Moein Moeini-Aghaie. Increasing the resilience of distribution systems against hurricane by optimal switch placement. In *Electrical Power Distribution Networks Conference (EPDC), 2017 Conference on*, pages 7–11. IEEE, 2017.

Q. Zhang, I.E. Grossmann, C.F. Heuberger, A. Sundaramoorthy, and J. M. Pinto. Air separation with cryogenic energy storage: Optimal scheduling considering electric energy and reserve markets. *AIChE Journal*, 61(5):1547–1558, 2015a.

Q. Zhang, I.E. Grossmann, A. Sundaramoorthy, and J.M. Pinto. Data-driven construction of convex region surrogate models. *Optimization and Engineering*, 2015b.

Qi Zhang, M.F. Morari, I. E. Grossmann, A. Sundaramoorthy, and J. M. Pinto. A projection-based method for production planning of multiproduct facilities. *AIChE Journal*, 55(10):2614–2630, 2009.

Qi Zhang, M.F. Morari, I. E. Grossmann, A. Sundaramoorthy, and J. M. Pinto. An adjustable robust optimization approach to scheduling of continuous industrial processes providing interruptible load. *Computers & Chemical Engineering*, 86(4): 106–119, 2016.

M. Zima and G. Andersson. Wide area monitoring and control as a tool for mitigation of cascading failures. In *2004 International Conference on Probabilistic Methods Applied to Power Systems*, pages 663–669, Sept 2004.

R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas. Matpower: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Transactions on Power Systems*, 26(1):12–19, Feb 2011. ISSN 0885-8950. doi: 10.1109/TPWRS.2010.2051168.

Ray D Zimmerman and Carlos E Murillo-Sanchez. Matpower 5.1-user’s manual. *Power Systems Engineering Research Center (PSEERC)*, 2015.

©2018 IEEE Part of this dissertation is reprinted with permission from “Tightening McCormick Relaxations Toward Global Solution of the ACOPF Problem” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., *IEEE Transactions on Power Systems*, 2018.

In reference to IEEE copyrighted material which is used with permission in this thesis, the IEEE does not endorse any of Purdue University’s products or services.

Part of this dissertation is reprinted from “13th International Symposium on Process Systems Engineering (PSE 2018)”, Volume 44, Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., “Strengthened SOCP Relaxations for ACOPF with McCormick Envelopes and Bounds Tightening”, pages 1555–1560, 2018, with permission from Elsevier.

Part of this dissertation is reprinted from “Evaluating Demand Response Opportunities for Power Systems Resilience Using MILP and MINLP Formulations” by Bynum, M., Castillo, A., Watson, J.P., and Laird, C.D., to appear in *AIChE Journal*, 2018.

Part of this dissertation is reprinted with permission from “A Software Framework for Assessing the Resilience of Drinking Water Systems to Disasters with an Example Earthquake Case Study” by Klise, K.A., Bynum, M., Moriarty, D., and Murray, R., *Environmental Modelling & Software*, 2017.