AUTHORIZATION TO LEND AND REPRODUCE THE THESIS

As the sole author of this thesis, I authorize Brown University to lend it to other institutions or individuals for the purpose of scholarly research.

Date ______________ ____

Benjamin Simon, Author

I further authorize Brown University to reproduce this thesis by photocopying or other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research.

Date ______________ ____

Benjamin Simon, Author

Randomized Adaptive Vehicle Decomposition for Large-Scale Power Restoration

by

Benjamin Simon

Sc.B., Brown University, 2011

A thesis submitted in partial fulfillment of the
requirements for the Degree of Master of Science
in the Department of Computer Science at Brown University

Providence, Rhode Island

May 2012

This thesis by Benjamin Simon is accepted in its present form by
the Department of Computer Science as satisfying the thesis requirement
for the degree of Master of Science.

Date ______________ ______________________________

Professor Pascal Van Hentenryck, Advisor

Approved by the Graduate Council

Date ______________ ______________________________

Peter M. Weber, Dean of the Graduate School

Abstract

This paper considers the joint repair and restoration of the electrical power system after significant disruptions caused by natural disasters. This problem is computationally challenging because, when the goal is to minimize the size of the blackout, it combines a routing and a power restoration component, both of which are difficult on their own. The joint repair/restoration problem has been successfully approached with a 3-stage decomposition, whose last step is a multiple-vehicle, pickup-and-delivery routing problem with precedence and capacity constraints whose goal is to minimize the sum of the delivery times (PDRPPCCDT). Experimental results have shown that the PDRPPCCDT is a bottleneck and this paper proposes a Randomized Vehicle Decomposition (RVD) to scale to very large power outages. The RVD approach has been shown to produce significant computational benefits and provide high-quality results for infrastructures with more than 1200 needed repairs.

Vita

Benjamin Simon was born in Evanston, Illinois. He graduated from Evanston Township High School in 2007 and entered Brown University in the fall of 2007, where he studied applied mathematics and computer science. He graduated with an Sc.B. in December of 2010 and began a 5th Year Master's in January 2011 which he will complete in December 2011. During his time at Brown, he was a teaching assistant in the Computer Science Department for six semesters and a head teaching assistant for two semesters. He also completed three internships at Google Inc. in the summers of 2009, 2010, and 2011 and will be returning as a full-time software engineer after he graduates.

Acknowledgements

I would like to thank Professor Pascal Van Hentenryck, my advisor, for his support and direction of this research project. I am also very thankful of his Ph.D. student, Carleton Coffrin who provided me with invaluable assistance and who was always available and willing to discuss my results and ideas. Additional thanks to my peers at Brown and I'd like to especially thank Nell Elliott who has played a crucial role in my education for the past eight years.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Every year, seasonal hurricanes threaten coastal areas. The severity of hurricane damage varies from year to year, but hurricanes often cause power outages that have considerable impacts on both quality of life (e.g., crippled medical services) and economic welfare. Therefore considerable human and monetary resources are always spent to prepare for, and recover from, threatening disasters. At this time, policy makers work together with power system engineers to make the critical decisions relating to how money and resources are allocated for preparation and recovery of the power system. Unfortunately, due to the complex nature of electrical power networks, these preparation and recovery plans are limited by the expertise and intuition of power engineers. Moreover, current preparation methods often do not use valuable disaster-specific information.

This research reconsiders the last-mile disaster recovery for power restoration, i.e., how to schedule and route a fleet of repair crews to restore the power network as fast as possible after a disaster. This problem was considered for the first time in [21] which proposed a decomposition approach to handle the significant computational complexity of this application. Indeed, last-mile power restoration combines a combinatorial vehicle routing problem with a traditional power restoration process. A direct approach, which jointly optimizes the vehicle schedules and the power restoration process cannot meet the real-time constraints imposed in disaster recovery. The decomposition approach was shown to improve the practice in the field, significantly reducing the size of the blackout over time. It is deployed in Los Alamos National Laboratory tools and activated to advise the federal government, each time a hurricane of category 3 or above threatens to hit the United States.

The last step of the decomposition approach is a multiple-vehicle, pickup-and-delivery, vehicle routing problem with capacity and precedence constraints whose goal is to minimize the sum of the delivery times. The precedence constraints are introduced to obtain a good restoration plan from a power system perspective, while the objective function is a proxy for minimizing the blackout size. Experimental results indicated that this last routing step was the bottleneck of this approach.

The goal of this paper is to overcome this limitation. It proposes a randomized adaptive vehicle decomposition approach to scale to large-scale disasters, e.g., electrical networks containing more than 1200 damaged components. Randomized adaptive decompositions were proposed for vehicle routing in [8, 9] and exploited spatial and temporal locality. This paper uses a randomized adaptive vehicle decomposition (RAVD) to account for the precedence constraints which are a fundamental difficulty in this context. Experimental results show that the RAVD algorithm produces significant computational benefits over large neighborhood approaches and provides high-quality results for infrastructures on scale with a state. These damage scenarios correspond to vehicle routing problems with as many as 2500 visits.

Chapter 2

Constraint Programming Model

This section presents a constraint-programming model for the multiple-vehicle, pickup-and-delivery routing problem with precedence and capacity constraints whose goal is to minimize the sum of the delivery times (PDRPPCCDT). Figure 1 depicts the model, which is almost a direct translation of the problem specifications. The model is defined in terms of locations, i.e., the pickups, the deliveries, and the starting and ending locations of the vehicles. The decision variables associate with every location l the next location in the visit order, the vehicle visiting l, the load of the vehicle when it arrives at l, and the earliest delivery time for l. The successor variables make up a large circuit by connecting the ending and starting locations of vehicles together. The objective function minimizes the summation of the delivery times of the dropoff locations. Constraint (2) eliminates subtours. Constraints (3)-(7) initialize the initial load and delivery times of vehicles and their first visit has the correct vehicle, load, and delivery time. Constraints (8)-(10) specify the constraints for successors, which have the same vehicle, a modified load, and a larger delivery time than their predecessors. Constraint (11) makes sure that every pickup and delivery pair is served by the same vehicle.

Model 1 A Constraint-Programming Model for the PDRPPCCDT.

Let:

$W^- = \{1 \ldots d\}$

$W^+ = \{d+1 \ldots 2d\}$

$J = W^- \cup W^+$

$H^+ = 2d+1 \ldots 2d+m$

$H^- = 2d+m+1 \ldots 2d+2m$

$L = W^- \cup W^+ \cup H^+ \cup H^-$

$Pair : W^+ \rightarrow W^-$ – The pickup associated with a dropoff

Variables:

$\sigma[L] \in L$ – successor of a location

$vehicle[L] \in V$ – vehicle assignment of a location

$load[L] \in \{0, \ldots, c\}$ – Vehicle load at a location

$EDT[L] \in \{0, \ldots, \infty\}$ – delivery time of a location

Minimize:

$$\sum_{i \in W^+} EDT[i] \qquad (1)$$

Subject To:

$$circuit(\sigma) \qquad (2)$$

for $l \in H^+$

$$vehicle[l] = vehicle[\sigma[l]] \qquad (3)$$

$$load[l] = 0 \qquad (4)$$

$$load[\sigma[l]] = 0 \qquad (5)$$

$$EDT[l] = 0 \qquad (6)$$

$$EDT[\sigma[l]] = T(l, \sigma[l]) + M(\sigma[l]) \qquad (7)$$

for $l \in J$

$$vehicle[l] = vehicle[\sigma[l]] \qquad (8)$$

$$load[\sigma[l]] = load[l] + d(l) \qquad (9)$$

$$EDT[\sigma[l]] \geq s(\sigma[l]) + T(l, \sigma[l]) + EDT[l] \qquad (10)$$

for $l \in W^+$

$$vehicle[l] = vehicle[Pair(l)] \qquad (11)$$

Chapter 3

Large Neighborhood Search

As mentioned earlier, the PDRPPCCDT is solved using LNS and constraint programming in [21]. More precisely, the large neighborhood search, LNS(R), selects a random set of dropoff locations and relaxes the value of the successor and predecessor of each location in the set and its corresponding pickup location. It then uses CP to search for improving solutions in this neighborhood before repeating the process. For this research, we also experimented with three additional neighborhoods.

1. **Spatial Neighborhood (S):** This neighborhood chooses a location l randomly and then selects other locations with a probability inversely proportional to the normalized distance to l.

2. **Temporal Neighborhood (T):** This neighborhood chooses a location l randomly and then selects other locations with a probability inversely proportional to the normalized time difference in delivery times with l.

3. **Vehicle Neighborhood (V):** This neighborhood selects a number of vehicles (about a quarter of the total vehicles) and selects dropoff locations randomly from these vehicles.

In the rest of the paper, we use LNS(R), LNS(S), LNS(T), and LNS(V) to denote the LNS algorithms over the various neighborhoods, LNS(R) being the algorithm in [21].

Chapter 4

Randomized Adaptive Decompositions

To scale to large PDRPPCCDT instances, we use the randomized adaptive decomposition scheme proposed in [8]. Given a routing problem $\mathcal{P}$, its key idea is to use the current solution σ of $\mathcal{P}$ to find a decoupling $(\mathcal{P}_o, \mathcal{P}_s)$ with projected solution σ_o and σ_s. The subproblem $\mathcal{P}_o$ is then reoptimized and its solution is merged with σ_s to obtain a new solution to $\mathcal{P}$. More precisely, the Adaptive Decomposition Scheme (ADS) does the following:

1. Starting from plan σ_0, it produces a sequence of plans $\sigma_1, \ldots, \sigma_j$ such that $f(\sigma_0) \geq f(\sigma_1) \geq \ldots \geq f(\sigma_j)$.

2. At step i, the scheme uses σ_{i-1} to obtain a decoupling $(\mathcal{P}_o, \mathcal{P}_s)$ of $\mathcal{P}$ with projected solutions σ_o and σ_s. It reoptimizes $\mathcal{P}_o$ to obtain σ_o^* and the new plan $\sigma_i = \textsc{merge}(\sigma_o^*, \sigma_{i-1})$

One of the most challenging aspects of ADS is how to perform the merging of the decoupled solutions, i.e, $\sigma_i = \textsc{merge}(\sigma_o^*, \sigma_{i-1})$. In [8], this challenge is addressed by choosing $\mathcal{P}_o$ such that the customers of entire vehicles are removed. The merging operation is then trivial, since the vehicles in ($\mathcal{P}_o$ and $\mathcal{P}_s$) are disjoint. More sophisticated, temporal and spatial, decouplings were also explored in [9].

Precedence constraints in the PDRPPCCDT complicate the more sophisticated decompositions and may also make spatial decompositions much less effective. As a result, this paper considers a Randomized Adaptive Vehicle Decomposition (RAVD). At each step, a quarter of the vehicles are chosen randomly and then LNS(R) is run on the subproblem consisting of all of those vehicles and all of their jobs. In the subproblem, precedence constraints are imposed between the jobs in the decomposition. Additionally, like in [18], the decomposition imposes temporal constraints on the jobs to ensure that optimizing the subproblem does not degrade the overall objective function. In particular, each job is given a lower and upper bound on its earliest delivery time computed from the precedence constraints between it and the jobs not being considered in the decomposition.

Chapter 5

Experimental Results

5.1 Benchmarks

The benchmarks were produced by the Los Alamos National Laboratory based on the electrical infrastructure of the United States. The disaster scenarios were generated using state-of-the-art hurricane simulation tools used by the National Hurricane Center [1, 19]. The damages in the scenarios vary significantly, from 50 to 1000 components to repair, inducing between 100 and 2000 routing visits. The scenarios represent power restoration problems at the scale of US state (e.g., Florida) and are thus of significant computational complexity. Instances with 100 damaged components focus on the transmission network alone, while those with 500 and 1000 damaged items incorporate aspects of both the transmission and distribution networks. For this paper, we used a total of 15 representative disaster contingencies. For each instance, the current "best practice" in the field serves as a baseline, where the "best practice" implements an agent-based greedy routing algorithm that satisfies the restoration order. The experiments were run on quad-core Dell PowerEdge 1855 blade systems with Xeon 2.8 processors. The execution times vary from 1 to 20 hours depending on the instance size. The graphs show the average over 10 executions for each benchmark and configuration.

Instance	Repairs	Hrs	Greedy	LNS(R)	RAVD	% Greedy	% LNS
BM2-h15	67	1	134837	110860	107320	20.41	3.19
BM2-h09	86	1	184894	150780	148300	19.79	1.64
BM2-h03	100	1	286071	227390	220395	22.96	3.08
Average						21.05	2.64
BM3-h15	53	1	67012	57768	48248	28.00	16.48
BM3-h00	61	1	92512	81234	70079	24.25	13.73
Average						26.12	15.11
BM4-h16	54	2	84230	65741	54730	35.02	16.75
BM4-h15	97	2	261220	220140	163820	37.29	25.58
BM4-h09	106	2	307171	246990	182490	40.59	26.11
BM4-h03	121	2	437972	354570	264690	39.56	25.35
Average						38.12	23.45
BM5-h03	206	8	1307329	1033600	674320	48.42	34.76
BM5-h11	255	8	2248539	1779900	1133900	49.57	36.29
BM5-h17	432	8	6341765	5998400	5150500	18.78	14.14
BM5-h00	439	8	6782781	6189600	4609600	32.04	25.53
BM5-h05	504	8	9112107	8220600	6072100	33.36	26.14
Average						36.43	27.37
BM5-h05	504	20	9112107	7958400	5433300	40.37	31.73

Table 5.1: Quality of the Results: Summary on All Benchmarks.

5.2 Quality of the Results

Table 5.1 presents a summary of the quality of the results. For each benchmark, the table gives the number of damaged components, the number of hours the optimization algorithm was run and the average objective value for the greedy solution, LNS from [21] , and the randomized adaptive decompostion. The table also gives the percentage improvement of RAVD with respect to the greedy and the LNS solutions. There is a line for each benchmark and an average for each benchmark class (e.g., BM2) which corresponds to a specific network. The last line reports the results of benchmark BM5-h05 when run for 20 hours. The results indicate that RAVD brings substantial benefits over LNS(R) and that these benefits increase with

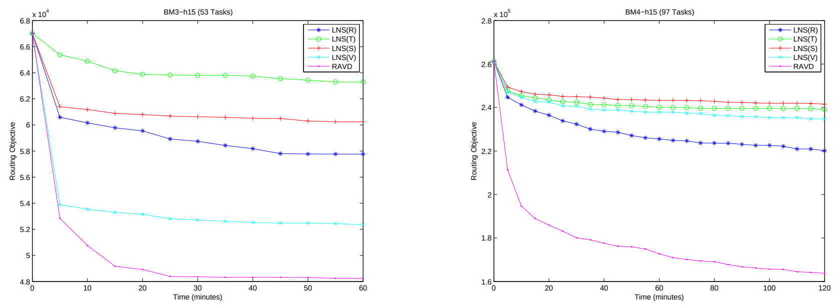

Figure 5.1: Quality of the Results over Time: 53 and 97 Jobs.

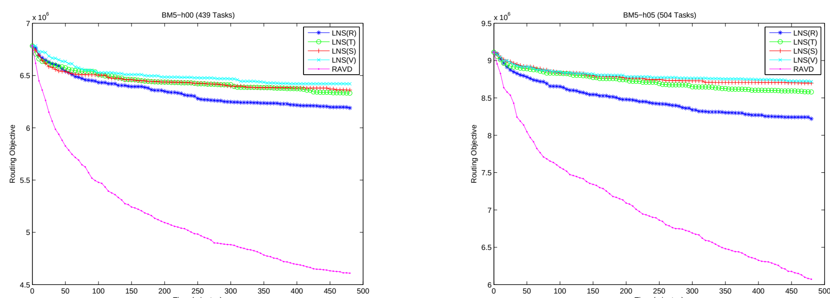

Figure 5.2: Quality of the Results over Time: 439 and 504 Jobs.

the size of the damages. On the BM5 benchmark class, the improvement is about 27% in average, which is substantial. The table only compares RAVD and LNS(R) for reasons that will become clear shortly. Overall, RAVD brings tremendous benefits in solution quality over large neighborhood search and significant improvements over "best practices".

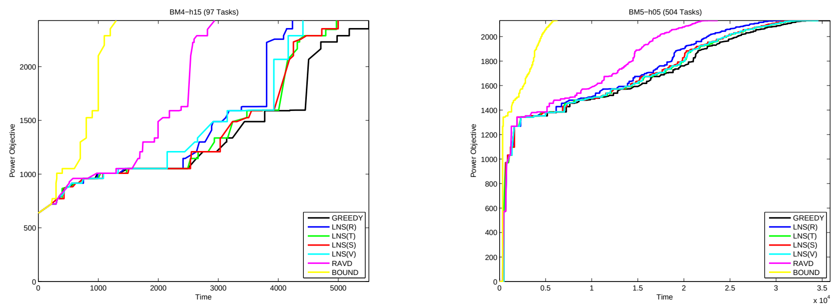

Figure 5.3: Size of the Blackouts: 97 and 504 Tasks.

5.3 LNS Versus RAVD

It is interesting to look at specific benchmarks to understand these results in more detail. Figure 5.1 depicts the average solution quality over time for the BM3-h15 and BM4-h15 benchmarks which are relatively small. On BM3-h15, it is interesting to observe that LNS(V) outperforms the other LNS procedures, although it is dominated by RAVD. On BM4-h15, LNS(R) is the best LNS algorithm but it is significantly dominated by RAVD. Figure 5.2 presents the same results for two large benchmarks with 439 and 504 jobs. Once again, LNS(R) dominates the other LNS algorithms and RAVD provides significant benefits over all LNS algorithms. Figure 5.3 shows the restored power over time resulting from single solutions from different routing algorithms for networks with 97 and 504 damage components. The figure also gives a very crude lower bound obtained by ignoring travel distances, i.e., viewing the problem as a pure restoration without taking account the travel times of repair crews. Both benchmarks show significant reductions in blackout sizes. On the 97-damage network, RAVD almost cuts in half the gap between the crude lower bound and the LNS(R) algorithm, giving some indirect evidence of its quality.

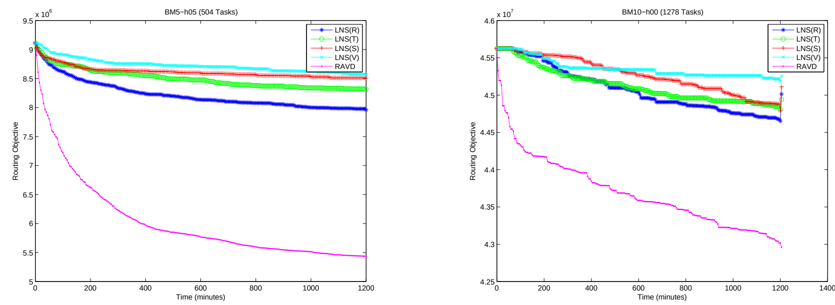

Figure 5.4: Convergence of the Results over Time: 504 and 1278 Tasks.

5.4 Convergence of the Results

RAVD produces improved routings steadily over time, with the more significant improvements coming early. However, the experimental results indicate that these instances are computationally challenging and RAVD may continue to improve the solutions for many hours. Figure 5.4 shows how the 504-damage benchmark and an even larger 1278-damage problem behave when RAVD is given 20 hours of CPU time. The 504-damage benchmark has typically reached its best solution at that point: Several of the individual runs have in fact reached a plateau and the remaining ones are close to reaching that solution. The largest benchmark with 2556 visits is still improving after 8 hours but has produced significant improvements over the LNS algorithms. Overall, the graphs all show that RAVD produces significant improvements over the LNS algorithms early in the run and then continue with a steeper rate of solution improvement before converging.

5.5 The Impact of the Precedence Constraints

It is not immediately clear why RAVD produces such significant improvements over the LNS approaches. Past work has shown that focused neighborhoods such

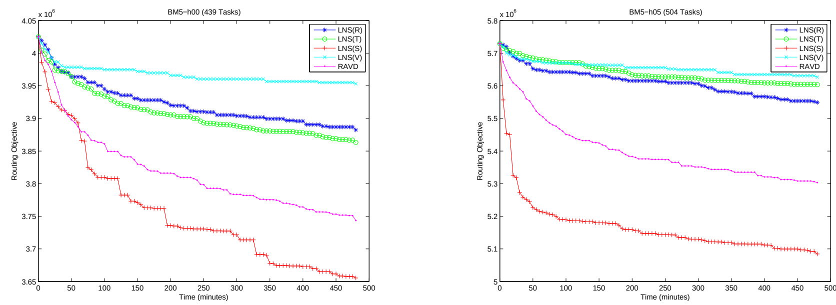

Figure 5.5: The Impact of Precedence Constraints: 439 and 504 Tasks.

as the ones studied here are normally able to overcome scaling issues. To better understand the behavior of the algorithms, we applied the algorithms to the instances but with the precedence constraints removed. Figure 5.5 reports the results on two large instances with 439 and 504 tasks, i.e., routing problems with 878 and 1008 visits. The results are quite interesting. They indicate that, without precedence constraints, LNS(S) is the best algorithm, followed by RAVD and then the other LNS algorithms. LNS(S) clearly dominates the other LNS algorithms significantly, while RAVD is really in between LNS(S) and the other LNS approaches. The figure also shows that RAVD dominates LNS(V) which, in this context, is the worst algorithm. The success of LNS(S) seems to indicate that LNS can normally scale to very large instances. Our conjecture is that there are significant benefits to a two-level optimization approach that

- Focuses on a well-isolated subproblem;
- Optimizes this subproblem with LNS.

Note that, on the PDRPPCCDT problem, RAVD exploits a vehicle decomposition for the subproblems since precedence constraints make it difficult to obtain a natural spatial or temporal decomposition. It then uses LNS(R), the best LNS algorithm, for optimizing the subproblems.

Chapter 6

Conclusion

This paper reconsidered the joint repair and restoration of the electrical power system after significant disruptions caused by natural disasters. This problem is computationally challenging when the goal is to minimize the size of the blackout because it combines a routing and a power restoration component, both of which are difficult on their own. The joint repair/restoration problem has been successfully approached with a 3-stage decomposition in [21], whose last step is a multiple-vehicle, pickup-and-delivery routing problem with precedence and capacity constraints whose goal is to minimize the sum of the delivery times (PDRPPCCDT). Experimental results have shown that this routing problem was the bottleneck of the approach.

This paper remedied this limitation and proposed a Randomized Adaptive Vehicle Decomposition (RAVD) that scales to very large power outages. The RAVD algorithm was shown to produce significant computational benefits over various LNS algorithms and provides high-quality results for infrastructures with more than 1200 damaged components.

The experimental results have also isolated the difficulties raised by precedence constraints for spatial neighborhoods. Moreover, randomized adaptive decompositions seem to leverage LNS strengths to another level. In particular, the ability of randomized adaptive decompositions to optimize subproblems with LNS seems to produce significant benefits in solution quality and speed. Future work will attempt to confirm this conjecture on other problems and neighborhoods.

Bibliography

[1] Fema hazus overview. Available online at http://www.fema.gov/plan/prevent/hazus/, 2010.

[2] M. Adibi. *Power System Restoration(Methodologies & Implementation Strategies)*. 2000.

[3] M.M. Adibi and L.H. Fink. Power system restoration planning. *Power Systems, IEEE Transactions on*, 9(1):22 –28, feb. 1994.

[4] M.M. Adibi, L.R.J. Kafka, and D.P. Milanicz. Expert system requirements for power system restoration. *Power Systems, IEEE Transactions on*, 9(3):1592 –1600, aug. 1994.

[5] J.J. Ancona. A framework for power system restoration following a major power failure. *Power Systems, IEEE Transactions on*, 10(3):1480 –1485, aug. 1995.

[6] R. Bent and P. Van Hentenryck. A Two-Stage Hybrid Local Search for the Vehicle Routing Problem with Time Windows. *Transportation Science*, 8(4):515–530, 2004.

[7] R. Bent and P. Van Hentenryck. A Two-Stage Hybrid Algorithm for Pickup and Delivery Vehicle Routing Problems with Time Windows. *Computers and Operations Research (Special Issue on Applications in Combinatorial Optimization)*, pages 875–893, 2006.

[8] R. Bent and P. Van Hentenryck. Randomized Adaptive Spatial Decoupling For Large-Scale Vehicle Routing with Time Windows. In *Proceedings of the 22th National Conference on Artificial Intelligence (AAAI'07)*. AAAI Press, July 2007.

[9] R. Bent and P. Van Hentenryck. Spatial, Temporal, and Hybrid Decompositions For Large-Scale Vehicle Routing with Time Windows. In *16th International Conference on Principles and Practice of Constraint Programming 2010 (CP 2010)*, St Andrews, Scotland, 2010. Springer Verlag.

[10] Ann Melissa Campbell, Dieter Vandenbussche, and William Hermann. Routing for relief efforts. *Transportation Science*, 42(2):127–145, 2008.

[11] A. Delgadillo, J.M. Arroyo, and N. Alguacil. Analysis of electric grid interdiction with line switching. *Power Systems, IEEE Transactions on*, 25(2):633 –641, may. 2010.

[12] E.B. Fisher, R.P. O'Neill, and M.C. Ferris. Optimal transmission switching. *Power Systems, IEEE Transactions on*, 23(3):1346 –1355, aug. 2008.

[13] J.A. Huang, L. Audette, and S. Harrison. A systematic method for power system restoration planning. *Power Systems, IEEE Transactions on*, 10(2):869 –875, may. 1995.

[14] J.A. Huang, F.D. Galiana, and G.T. Vuong. Power system restoration incorporating interactive graphics and optimization. pages 216 –222, may. 1991.

[15] A.L. Morelato and A.J. Monticelli. Heuristic search approach to distribution system restoration. *Power Delivery, IEEE Transactions on*, 4(4):2235 –2241, oct. 1989.

[16] H. Mori and Y. Ogita. A parallel tabu search based approach to optimal network reconfigurations for service restoration in distribution systems. volume 2, pages 814 – 819 vol.2, 2002.

[17] T. Nagata, H. Sasaki, and R. Yokoyama. Power system restoration by joint usage of expert system and mathematical programming approach. *Power Systems, IEEE Transactions on*, 10(3):1473 –1479, aug. 1995.

[18] Dario Pacino and Pascal Van Hentenryck. Large neighborhood search and adaptive randomized decompositions for flexible jobshop scheduling. In Toby Walsh, editor, *IJCAI*, pages 1997–2002. IJCAI/AAAI, 2011.

[19] Dorothy A. Reed. Electric utility distribution analysis for extreme winds. *Journal of Wind Engineering and Industrial Aerodynamics*, 96(1):123–140, 2008.

[20] T. Sakaguchi and K. Matsumoto. Development of a knowledge based system for power system restoration. *Power Apparatus and Systems, IEEE Transactions on*, PAS-102(2):320 –329, feb. 1983.

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

[22] M.H. Yolcu, Z. Zabar, L. Birenbaum, and S.A. Granek. Adaptation of the simplex algorithm to modeling of cold load pickup of a large secondary network distribution system. *Power Apparatus and Systems, IEEE Transactions on*, PAS-102(7):2064 –2068, jul. 1983.