

Computational Disaster Management

Pascal Van Hentenryck

Optimization Research Group, NICTA, Australia

Computing and Information Systems, University of Melbourne, Australia

1 Introduction

The frequency and intensity of natural disasters have been increasing significantly over the past decades and this trend is predicted to continue. Natural disasters have a dramatic impact on human lives and on the socio-economic welfare of entire regions. They were identified in 2011 by the World Bank [The World Bank, 2011] as one of the major risks of the East Asia and Pacific region, which represents 85 percents of all people affected since 2007. Moreover, this exposure will likely double by 2050 due to rapid urbanization and climate change. To understand the magnitude of such disasters, consider Irene, a category 3 hurricane that hit the East Coast of the United States in August 2011. It killed 56 people, inflicted damages now estimated in the range of 15 billion dollars, and created blackouts that lasted for several days. Hurricane Sandy and the Tohoku tsunami in Japan were even more dramatic, affecting human welfare in entire regions and damaging entire segments of the economy. For instance, Japanese manufacturers lost a significant market shares after the tsunami.

It is possible however to mitigate the impact of these disasters through appropriate public and corporate policies, investment in infrastructure resilience, real-time decision-support systems, and education. In general, the focus of decision-support systems for disaster management is on *situational awareness*, i.e., communicating to decision makers the situation in the field as accurately as possible. Situational awareness is obviously a critical component of any decision-support system for disaster management. However, Hurricane Katrina, a traumatic event for the United States, indicated the need to go beyond the communication of timely information; it is critical to enhance the cognitive abilities of decision makers through advanced use of optimization and simulation technology. The Katrina report [United-States Government, 2006], which should be required reading for policy-makers and emergency officials around the world, pointed out that “*the existing planning and operational structure for delivering critical resources and humanitarian aid clearly proved to be inequale to he task.*” In addition, the report recommended that “*the Federal government must develop the capacity to conduct large-scale logistical operations*”. Similar observations have been made elsewhere: A European ambassador at a donor conference for Tsunami relief even said that “we do not need another donors’ conference; we need

a logistics conference”. It seems clear that a novel, holistic approach is necessary for disaster management.

The United States responded strongly to the Katrina report, starting major initiatives in various departments, including the Department of Homeland Security (DHS). Substantial progress has been achieved since 2006 and this paper reviews some efforts that are now deployed to mitigate the impact of disasters. However, much more is needed and this paper sketches some of the challenges and a long-term vision for *computational disaster management*.

This paper articulates the role of optimization for disaster management and its inherent complexity . It reviews a number of case studies in this space, some of which in deployment, to highlight the benefits of optimization . It concludes by articulating a potential long-term vision for computational disaster management, articulating some of the broader computational challenges. Obviously, space limits make it difficult to enter into deep technical issues but the hope is that this paper will provide the incentives to learn more about this important and fascinating field and, hopefully, to join this effort.

2 Optimization for Disaster Management

Many optimization problems in disaster management always start with spatial and temporal information about communities and infrastructures. It is critical, for instance, to know where people live and where hospitals and distribution centers are located, as well as to have accurate models of various networks, such as the transportation and power systems. The second key input is a prediction about what the disaster might do and how it will affect communities, assets, and infrastructures. For hurricanes such as Irene and Sandy, the National Hurricane Center (NHC) of the National Weather Service in the United States is highly skilled at generating ensembles of possible hurricane tracks.¹ These tracks are then run through fragility simulators to determine their impact, identifying resource needs and damages to assets and infrastructures. These simulation steps are illustrated in Figure 1. The resulting scenario sets, illustrated in Figure 2, is the input to various types of optimization problems.

Optimization problems exploiting these inputs arise at different levels: Strategic, tactical, response, and recovery. Op-

¹Some countries, such as Australia, have much weaker prediction tools for cyclones however.



Figure 1: The Simulation Steps Before Optimization.




Figure 2: Optimization for Disaster Management: Inputs.

timization at the strategic level include applications such as how to stockpile relief supplies or repair parts for infrastructures, how to schedule planned burns for minimizing the risk of bush fires, where to build levees for flood management, and how best to evacuate a region. Strategic planning typically takes place before the hurricane, bushfire, or flood seasons. The tactical level starts when a disaster first materializes, e.g., from a few days before a hurricane hits, and ends at the time of impact. It considers problems such as asset repositioning, sandbagging, and evacuation and sheltering in place. Responding to a disaster involves search & rescue operations, relief distribution, evacuations, and damage assessment to name only a few. Finally, the recovery phase, which is often critical for social and economic welfare, is concerned with all aspects of restoring critical infrastructures such as the transportation network, the power system, and telecommunication networks.

3 Computational Complexity

In disaster management, decision makers are faced with optimization problems of daunting complexity, which explains why they may be overwhelmed by the magnitude of the task. The field of humanitarian logistics has investigated some of these problems since the 1990s and recent disasters have brought increased attention to the logistics aspects [Wassenhove, 2006; Beamon, 2008; United-States Government, 2006; Fritz Institute., 2008]. It is well recognized that innovative research is required to meet the underlying challenges (e.g., [Wassenhove, 2006; Beamon, 2008]). The complexity of computational disaster management can be attributed to many factors. Here are a few that are particularly striking for optimization experts.

1. **Large Scale** - The optimization problems are large-scale and concern entire cities or states [Van Hentenryck *et*

al., 2010; Coffrin *et al.*, 2011b], or even international multimodal supply chains.

2. **Non-Standard Objective Functions** - In computational disaster management, the goal is not to maximize profit or minimize costs: Rather it is to maximize some notion of social welfare or an equibilty objective [Barbarosoglu *et al.*, 2002; Campbell *et al.*, 2008; Balcik *et al.*, 2008]. These problems are much less studied and often harder than their traditional counterparts.
3. **Stochastic Aspects** - Computational disaster managements operate in inherently uncertain environment due to the disaster itself, the way people react, and the limitations in information gathering. Preparations and recovery plans must be robust with respect to many scenarios [Duran *et al.*, 2009; Gunnec and Salman, 2007].
4. **Multiple Objectives** - High-stake disaster situations often have to balance conflicting objectives such as operational costs, speed of service, and unserved demand [Barbarosoglu *et al.*, 2002; Duran *et al.*, 2009; Balcik *et al.*, 2008; Gunnec and Salman, 2007].
5. **Complex Infrastructures** - Disasters typically damage complex infrastructures such as the transportation network and the power system. Optimization applications make discrete decisions over these complex infrastructures, which are often modeled by complex systems of constraints.
6. **Multiple Stakeholders** - Disasters often involve multiple infrastructures, agencies, and communities. It becomes critical to incentivize different stakeholders to maximize social welfare.

As a result, disaster management applications present unique computational challenges for optimization. Commercial, off-the-shelf packages simply do not scale for these applications [Campbell *et al.*, 2008].

4 Case Studies

The last couple of years have witnessed significant progress in this area. It is beyond the scope of this paper to review all of them. Instead, this section reviews some applications our research group has been directly involved with. Some of these have been deployed at Los Alamos National Laboratories as part of the National Infrastructure Simulation and Analysis Center (NISAC) to aid federal organizations in the United States and used on hurricanes Irene and Sandy. It is important to keep in mind however that these applications were a first attempt at these computationally challenging problems: Proving optimality is beyond the reach of existing technology. Instead, our methodology consists in finding high-quality solutions within the time constraints imposed by these applications and, whenever possible, in providing lower bounds to evaluate solution quality. It is likely that improvements to these results will be obtained in the future.

4.1 Relief Distribution

Relief distribution is critical in disaster management: It is important to deliver water, food, medicine, and other commodi-




Figure 3: The Scalability of Fleet Routing.

ties as quickly as possible after a disaster strikes. For a particular commodity (or collection of commodities), the problem consists in deciding

1. where to store the supply and in how much quantity;
2. how to deliver the commodity as far as possible.

The first question is strategic or tactical in nature; the second one addresses the response. The overall problem is stochastic and the goal is to minimize a lexicographic objective function consisting of minimizing the unsatisfied demand, the latest delivery time, and the storage costs, using either stochastic or robust optimization. The second objective, which is unconventional, was a DHS requirements and differs from traditional objective functions in vehicle routing which often aim at minimizing travel distance or cost.

The relief distribution is too difficult to handle globally: It is a 2-stage stochastic inventory, location, and routing. It is also large-scale as any solution technique must be able to scale to large instances (e.g., the state of Florida). For some commodities such as water, there may be multiple trips per location. Our solution [Van Hentenryck *et al.*, 2010; Coffrin *et al.*, 2011b], which was the first last-mile humanitarian logistics covering all aspects of the problem (stochasticity, inventory, location, and routing), is a decomposition into four subproblems:

1. a stochastic warehouse allocation to decide where to store the commodity and how much;
2. a customer allocation to decide which warehouses supply which demands;
3. a repository routing to decide how to serve the demand associated with a warehouse;
4. a fleet routing to decide how to route all available vehicles globally.

The first subproblem is solved before the hurricane seasons or a few days before a hurricane hits. The remaining subproblems are the response. The subproblems were solved by different optimization technology, i.e., MIP solvers for the planning phase and hybrid optimization involving constraint programming and large neighborhood search for the response.

The approach was evaluated on benchmarks based on the United States infrastructure using hurricane scenarios from




Figure 4: The Benefits of Optimization.




Figure 5: Corporate Giving with a Central Location.

simulation tools from the National Hurricane Center. The largest benchmarks deal with entire states and involve more than 1,000,000 decision variables. The proposed approach was compared with the practice in the field. Figure 3 highlights the computational difficulties of disaster management. It shows the runtime behavior of two commercial solvers on small instances of the fleet routing as the budget increases and compares it with the approach in [Van Hentenryck *et al.*, 2010; Coffrin *et al.*, 2011b] (The y-axis depicts the runtime in log-scale). The approach finds optimal solutions on small instances and improves the practice in the field significantly on large ones as Figure 4 indicates. The latest delivery time is often reduced by 50%.

Optimization is also a powerful tool for evidence-based policy making, providing the scientific basis for complex public or corporate policy decisions. The Tohoku disaster clearly showed that international or domestic aid may complicate relief distribution and must be handled carefully. Figures 5 and 6 were part of a study we undertook to understand the role of corporate donations after a disaster. It shows drastic contrast in benefits when the corporate warehouse is in central or distant locations. Until the demand is met, the latest delivery times increase, which is normal, but the increase is significantly larger than in the case of a central location. When the budget is large, there is no benefit to corporate giving.



Figure 6: Corporate Giving with a Distant Location.

ing from a distant location, although it has significant benefits from a central location [Van Hentenryck *et al.*, 2013 Submitted for Publication].

4.2 Power Network Restoration

Computational disaster management often operates on infrastructures that have been affected by the disaster. It is a critical task to restore these infrastructures as quickly as possible. Blackouts in the power system for instance are not only dangerous for human lives; they are also extremely costly. A 24-hour blackout in San Diego in 2011 was estimated to have cost about 300 millions. Infrastructure restoration once again has strategic, tactical, response, and recovery aspects.

A strategic optimization was studied in [Coffrin *et al.*, 2011a] and considers how to stockpile electrical parts to recover the power systems as quickly as possible. This optimization used robust or stochastic optimization over a large number of scenarios. Each of these scenarios determines how best to repair the power system with the supply at hand.

The recovery of the power system [Van Hentenryck *et al.*, 2011] is particularly interesting in that it highlights the inherent complexity of computational disaster management. Its goal is to schedule a fleet of repair crews to minimize the size of the blackout. The complexity arises from the tension between two aspects of the problems:

1. the large-scale pickup and delivery vehicle routing problem that picks the electrical parts and use them to repair the damaged component
2. the power restoration which determines in which order to repair the components to minimize the blackout size.

Figure 7 depicts the two components and their interaction on a small example. The right figure shows the size of the blackout over time with the component repairs superimposed on the blackout. For instance, at time 20, component 3 is repaired, leading to a significant increase in flow. Component 6 is repaired at time 30 with no increase in power flow, while component 4 is fixed at time 40, restoring the entire power flow. The left figure depicts the routing of the repair crews. The top crew picks up parts 1 and 2 to repair components 3 and 4, then goes on repairing components 3 and 4.




Figure 7: The Two Aspects of Power System Restoration.

The two aspects of power restoration are NP-hard, yet the application must solve both at the same time. Moreover, the modeling of the blackout, even if one restricts attention to the steady states of the power system, requires to solve the traditional power flow equations which are nonlinear and extremely challenging to solve outside normal operating conditions [Overbye *et al.*, 2004]. Figure 8 depicts one step of this restoring process, the so-called activation problem, which ignores the routing and sequencing aspects and linearizes the power flow equations to obtain the so-called linearized DC model. This activation is extremely challenging computationally [Fisher *et al.*, 2008] and amounts to finding which component to activate (variable y_i). The underlying power system is modeled in terms of the voltage phase angles at the buses and the real power flows at the nodes and lines. Of particular interest are the flow conservation constraints (Constraints 7) and the power equations (Constraints 10–11) which must use a variable z_i to determine if the line is operational.

The power restoration problem, if handled in its globality, would have to integrate such activation problems in a pickup and delivery vehicle routing model to evaluate the size of the blackout for every (partial) routing configuration. Such an approach is outside the scope of state-of-the-art optimization technology. For this reason, a 3-step decomposition approach based on the concept of constraint injection was proposed in [Van Hentenryck *et al.*, 2011]. The approach decomposes the power restoration and routing aspects but its key idea to inject precedence constraints into the pickup and deliver problem to obtain a routing that minimizes the blackout size.

The approach was applied to the US transmission systems and hurricane scenarios obtained from state-of-the-art simulators, including on transmission systems with several thousand components and damages on about a third of the network. Figure 9 depicts some of the results and a comparison with the practice in the field. In Figure 9, the restoration plan on the right is the practice in field, the plan in the middle is the optimized restoration, and the dashed line on the left is a lower bound assuming infinitely many repair crews. On this scenario, optimization reduces the blackout size by about 50%. Results on large-scale transmission systems can be found in [Simon *et al.*, 2012].

There are many aspects that this brief overview is omitting. Power restoration raises some fundamental issues on power systems and it is necessary to work on better approximations of the power flow equations [Coffrin *et al.*, 2012b; 2012a;

Inputs:

$\mathcal{PN} = \langle N, L \rangle$ the power network
 D the set of damaged items
 R the set of repaired items
 $MaxFlow$ the maximum flow (MW)

Variables:

$y_i \in \{0, 1\}$ - item i is activated
 $z_i \in \{0, 1\}$ - item i is operational
 $P_i^l \in (-\hat{P}_i^l, \hat{P}_i^l)$ - power flow on line i (MW)
 $P_i^v \in (0, \hat{P}_i^v)$ - power flow on node i (MW)
 $\theta_i \in (-\frac{\pi}{6}, \frac{\pi}{6})$ - phase angle on bus i (rad)

Minimize

$$MaxFlow - \sum_{b \in N^b} \sum_{i \in N_b^l} P_i^v \quad (1)$$

Subject to:

$$y_i = 1 \quad \forall i \in (N \cup L) \setminus D \quad (2)$$

$$y_i = 0 \quad \forall i \in D \setminus R \quad (3)$$

$$z_i = y_i \quad \forall i \in N^b \quad (4)$$

$$z_i = y_i \wedge y_j \quad \forall j \in N^b, \forall i \in N_j^g \cup N_j^l \quad (5)$$

$$z_i = y_i \wedge y_{L_i^+} \wedge y_{L_i^-} \quad \forall i \in L \quad (6)$$

$$\sum_{j \in N_i^l} P_j^v = \sum_{j \in N_i^g} P_j^v + \sum_{j \in L_{L_i}} P_j^l - \sum_{j \in L_{O_i}} P_j^l \quad \forall i \in N^b \quad (7)$$

$$0 \leq P_i^v \leq \hat{P}_i^v * z_i \quad \forall j \in N^b, \forall i \in N_j^g \cup N_j^l \quad (8)$$

$$-\hat{P}_i^l * z_i \leq P_i^l \leq \hat{P}_i^l * z_i \quad \forall i \in L \quad (9)$$

$$P_i^l \geq B_i * (\theta_{L_i^+} - \theta_{L_i^-}) + M * (\neg z_i) \quad \forall i \in L \quad (10)$$

$$P_i^l \leq B_i * (\theta_{L_i^+} - \theta_{L_i^-}) - M * (\neg z_i) \quad \forall i \in L \quad (11)$$

Figure 8: A MIP Model for Minimizing Unserved Load.

Coffrin and Van Hentenryck, 2012]. Complex infrastructures are typically interdependent (e.g., [Lee *et al.*, 2007]) and it is necessary to restore them together [Coffrin *et al.*, 2012c] or to incentivize organizations to cooperate [Abeliuk *et al.*, 2013]. In general, the state of the power system is not known exactly and it is necessary to perform both damage assessment and repair simultaneously [Van Hentenryck *et al.*, 2012]. Much research is needed to handle these issues at large scales.

4.3 Large-Scale Flood Evacuation

Evacuation planning has received increasing attention in recent years. The optimization team at NICTA is designing algorithms for massive evacuations, such as those arising in major flooding events. Of particular interest is the Hawkesbury Nepean region (West Sydney) “where a one-in-1000 flood would cause up to \$8 billion in total damages” [Tim Barlass, 2012]. Evacuation planning must consider the dynamics of the flood and the state of the transportation network over time. Our recent research [Pillac *et al.*, 2013] shows that it is possible to design a conflict-based column-generation algorithm to evacuate about 70,000 people in real time (a MIP program would use about 160 million variables and is not able to solve the linear relaxation at the root node). Figure 10 depicts the modeling of evacuation. The top figure depicts an evacuation scenario, the node to evacuate (node 0), the evacuation centers (nodes A and B), and the transportation network, including information about when links are no longer available. The bottom figure describes the evacuation graph, including the population, the link capacities, and their lengths.




Figure 9: Benefits of Optimization for Power Restoration.

Evacuation planning raises many issues. Like in power systems, it is important to model the transportation network with reasonable accuracy. Moreover, evacuations raise fundamental issues in modeling human behavior. For instance, in some countries, most people disregard evacuation orders. Such knowledge should be incorporated in optimization algorithms and/or new policies may need to be enacted to encourage the population to cooperate.

5 A Vision for Disaster Management

Given the costs in human, social, and economic welfare, it is desirable to envision a future in which every city, every state, and every nation will have an advanced decision-support system for disaster management to mitigate the consequences of natural disasters. Implementing such a vision raises fundamental computational challenges at the intersection of computer science, engineering, operations research, and social sciences. Such a system would need to optimize over interdependent infrastructures and emergency-response processes in uncertain environments, simulate natural disasters with appropriate accuracy and response time, collect, aggregate, and visualize critical data in real time to provide accurate situational awareness and reduce uncertainty. It also requires understanding how decision makers, responders, and the population at large react in emergency situations in order to formulate effective plans.

The disaster management team at NICTA aims at building the science and technology to help make this vision a reality. It views computational disaster management holistically, integrating a variety of functional layers (see Figure 11) in a comprehensive decision-support platform. These layers indicate the magnitude of the task ahead and the wealth of scientific problems raised by computational disaster management.

- The *geospatial modeling* layer aims at building a complete map of a country or region, including its infrastructures, its terrain and vegetation, and its cities (including buildings and their occupancies) to name only a few. For instance, flood studies require fine-grained elevation and bathymetry maps.




Figure 10: Modeling Evacuation Planning Problem

- The *sensing and monitoring* layer collects information in real-time and monitors the situation on the field. It may use sensors, analyses of social media, crowd sourcing, and techniques such as hyperspectral imaging in order to provide timely situation awareness.
- The *data* layer manages information, including aggregation and fusion of multiple, heterogeneous, and possibly contradictory sources of uncertain information.
- The *behavioral* layer synthesizes and populates models on how individuals and groups react during emergencies in order to produce effective plans for situations involving evacuation and sheltering, isolation, search and rescue, and aid deliveries.
- The *simulation and forecasting* layer generates potential scenarios, achieving appropriate tradeoffs between accuracy and speed. The scenarios concern, not only the disaster itself, but also the interaction between the response, the disaster, and the critical infrastructures.
- The *optimization* layer provides decision support for strategic and tactical planning, response, and recovery and was the main topic of this paper.
- The *visualization* layer will offer unprecedented situational awareness and decision support, applying modern 3D-Visualization and information rendering techniques to disaster management. It will let decision makers explore and visualize various disaster and response scenarios, warn them about critical decisions and events, and provide them with a training platform.




Figure 11: Functionality Layers for Disaster Management.

This brief description highlights that computational disaster management is truly multi-disciplinary and goes well beyond the fundamental optimization challenges articulated in this paper. Optimization relies on the contributions of many fields to recommend proper decisions. Computational data science, data mining, and machine learning are critical to derive the behavioral models for optimization. Game theory and mechanism design are critical in incentivizing interdependent agencies to maximize social welfare, which is the topic of some recent work [Abeliuk *et al.*, 2013]. Tight integrations of optimization and simulation are likely to emerge from computational disaster management as optimization models need to incorporate accurate modeling of complex infrastructures. Crowdsourcing offers unprecedented potential for coordinating sensing, preparedness and recovery [Rutherford *et al.*, 2013] but its numerous successful applications may have overshadowed the potential for elaborate sabotages. Once again, game-theoretic frameworks can illuminate the fundamental tension between efficiency and vulnerability of crowdsourcing for disaster management [Naroditskiy *et al.*, 2013].

In summary, disaster management is a field filled with multi-disciplinary computational challenges whose solutions will have a fundamental impact on human, social, and economic welfare. It is an area where science can make a substantial difference to alleviate human suffering.

Acknowledgments

I would like to express my gratitude to Francesca Rossi for her kind invitation to present this work at IJCAI'13 and to my collaborators in this space. Special thanks to Russell Bent at Los Alamos for our long-term collaboration, Carleton Coffrin for being such an amazing collaborator, and Manuel Cebrian for broadening my view of disaster management. Thanks also to Nabeel Gillani, Ben Simon, and Nell Elliott at Brown and the team at NICTA, including Jeff Cotter, Victor Pillac, and Caroline Even. NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

References

- [Abeliuk *et al.*, 2013] A. Abeliuk, G. Berbeglia, and P. Van Hentenryck. Interdependent Scheduling Games. 2013. Submitted for Publication.
- [Balcik *et al.*, 2008] B. Balcik, B. Beamon, and K. Smilowitz. Last mile distribution in humanitarian relief. *Journal of Intelligent Transportation Systems*, 12(2):51–63, 2008.
- [Barbarosoglu *et al.*, 2002] Gulay Barbarosoglu, Linet Ozdamar, and Ahmet Cevik. An interactive approach for hierarchical analysis of helicopter logistics in disaster relief operations. *European Journal of Operational Research*, 140(1):118 – 133, 2002.
- [Beamon, 2008] B. Beamon. Humanitarian relief chains: Issues and challenges. *34th International Conference on Computers & Industrial Engineering*, pages 77–82, 2008.
- [Campbell *et al.*, 2008] Ann Melissa Campbell, Dieter Vandenbussche, and William Hermann. Routing for relief efforts. *Transportation Science*, 42(2):127–145, 2008.
- [Coffrin and Van Hentenryck, 2012] C. Coffrin and P. Van Hentenryck. A Linear-Programming Approximation of AC Power Flows. *ArXiv e-prints*, June 2012.
- [Coffrin *et al.*, 2011a] C. Coffrin, P. Van Hentenryck, and R. Bent. Strategic stockpiling of power system supplies for disaster recovery. *Proceedings of the 2011 IEEE Power & Energy Society General Meetings*, 2011.
- [Coffrin *et al.*, 2011b] Carleton Coffrin, Pascal Van Hentenryck, and Russell Bent. Spatial and objective decompositions for very large scaps. In *CPAIOR*, volume 6697 of *Lecture Notes in Computer Science*, Springer, 2011.
- [Coffrin *et al.*, 2012a] C. Coffrin, P. Van Hentenryck, and R. Bent. Approximating Line Losses and Apparent Power in AC Power Flow Linearizations. *Proceedings of the 2012 IEEE Power & Energy Society General Meetings*, 2012.
- [Coffrin *et al.*, 2012b] C. Coffrin, P. Van Hentenryck, and R. Bent. Smart Load and Generation Scheduling for Power System Restoration. *Proceedings of the 2012 IEEE Power & Energy Society General Meetings*, 2012.
- [Coffrin *et al.*, 2012c] Carleton Coffrin, Pascal Van Hentenryck, and Russell Bent. Last-mile restoration for multiple interdependent infrastructures. In Jörg Hoffmann and Bart Selman, editors, *AAAI*. AAAI Press, 2012.
- [Duran *et al.*, 2009] Serhan Duran, Marco Gutierrez, and Pinar Keskinocak. Pre-positioning of emergency items worldwide for care international. *Interfaces*, 41(3), 2011.
- [Fisher *et al.*, 2008] E.B. Fisher, R.P. O’Neill, and M.C. Ferris. Optimal transmission switching. *IEEE Transactions on Power Systems*, 23(3):1346–1355, 2008.
- [Fritz Institute., 2008] Fritz Institute. Fritz Institute Website. <http://www.fritzinstitute.org>, 2008.
- [Gunnec and Salman, 2007] D. Gunnec and F. Salman. A two-stage multi-criteria stochastic programming model for location of emergency response and distribution centers. In *INOC*, 2007.
- [Lee *et al.*, 2007] E.E. Lee, J.E. Mitchell, and W.A. Wallace. Restoration of services in interdependent infrastructure systems: A network flows approach. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 37(6):1303–1317, nov. 2007.
- [Naroditskiy *et al.*, 2013] Victor Naroditskiy, Nicholas R Jennings, Pascal Van Hentenryck, and Manuel Cebrian. Crowdsourcing dilemma. *arXiv preprint arXiv:1304.3548*, 2013.
- [Overbye *et al.*, 2004] T.J. Overbye, Xu Cheng, and Yan Sun. A comparison of the ac and dc power flow models for Imp calculations. In *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*, page 9 pp., 2004.
- [Pillac *et al.*, 2013] V. Pillac, P. Van Hentenryck, and C. Even. Conflict-Based Heuristic Column Generation for Evacuation Planning. 2013. Submitted for Publication.
- [Rutherford *et al.*, 2013] Alex Rutherford, Manuel Cebrian, Sohan Dsouza, Esteban Moro, Alex Pentland, and Iyad Rahwan. Limits of social mobilization. *Proceedings of the National Academy of Sciences*, 110(16):6281–6286, 2013.
- [Simon *et al.*, 2012] Ben Simon, Carleton Coffrin, and Pascal Van Hentenryck. Randomized adaptive vehicle decomposition for large-scale power restoration. In *CPAIOR*, volume 7298 of *Lecture Notes in Computer Science*, 2012.
- [The World Bank, 2011] The World Bank. Navigating Turbulence, Sustaining Growth. EAP Economic Update, 2011.
- [Tim Barlass, 2012] Tim Barlass. Urgent flood prevention measures needed for Warragamba Dam. The Sydney Morning Herald, November 2012.
- [United-States Government, 2006] United-States Government. The Federal Response to Hurricane Katrina: Lessons Learned, 2006.
- [Van Hentenryck *et al.*, 2010] Pascal Van Hentenryck, Russell Bent, and Carleton Coffrin. Strategic Planning for Disaster Recovery with Stochastic Last Mile Distribution. In *CPAIOR*, volume 6140 of *Lecture Notes in Computer Science*, pages 318–333. Springer, 2010.
- [Van Hentenryck *et al.*, 2011] 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, Stockholm, Sweden*, 2011.
- [Van Hentenryck *et al.*, 2012] Pascal Van Hentenryck, Nabeel Gillani, and Carleton Coffrin. Joint assessment and restoration of power systems. In *Frontiers in Artificial Intelligence and Applications (ECAI’12)*, pages 792–797. IOS Press, 2012.
- [Van Hentenryck *et al.*, 2013 Submitted for Publication] P. Van Hentenryck, R. Bent, and C. Coffrin. Strategic planning for disaster recovery with stochastic last mile distribution. 2013 (Submitted for Publication).
- [Wassenhove, 2006] L. Van Wassenhove. Humanitarian aid logistics: supply chain management in high gear. *Journal of the Operational Research Society*, 57(1), 2006.