SAND2016-8516C

Modeling Infrastructure Dependencies to Inform and Improve Electric Power Grid Resilience

Eric D. Vugrin
Sandia National Laboratories
P.O. Box 5800 MS1138
Albuquerque, New Mexico 87185-1138 USA
E-mail: edvugrin@sandia.gov

Introduction

Historically, U.S. government policy toward critical infrastructure security has focused on physical protection. However, following the terrorist attacks of September 11, 2001; the devastation from Hurricane Katrina in 2005; and a series of other disasters in the early 2000s, the infrastructure security community in the United States and globally recognized that it was simply not possible to prevent all threats to all assets at all times. Consequently, critical infrastructure resilience emerged as a complementary goal to prevention-focused activities. Whereas critical infrastructure security policies primarily focused on prevention of terrorism, accidents, and other disruptions, critical infrastructure resilience activities emphasize the infrastructure's ability to continue providing goods and services even in the event of disruptions.

For decades, reliability has been the primary standard for evaluating electrical power grid operations, but reliability metrics (e.g., SAIDI and SAIFI) were designed for outages that might occur during normal conditions. They cannot and are not used to evaluate grid operations for outages that occur as a result of hurricanes, earthquakes, cyber attacks, geomagnetic disturbances, and other low probability, high consequence events. Hence, mirroring the larger critical infrastructure community, the electrical power sector is demonstrating an increased interest in using resilience as a complementary objective to reliability goals [1].

Though there is much current debate and research on how to formally define and measure grid resilience, little debate exists about the essential role that electric power has in contributing to the resilience of dependent critical infrastructure systems and communities. For example, when Superstorm Sandy made landfall on the eastern coast of the United States in 2012, the resulting loss of power caused (at least in part):

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

- The failure of a quarter of the cell phone towers in the 10 states affected by the storm;
- Extensive gasoline shortages when electricity-powered pumps could not pump gas to potential customers;
- Interruption of mass transit systems in New York and New Jersey; and
- The evacuation of 6 New York hospitals.

These anecdotal observations have been further emphasized in federal resilience policies. Presidential Policy Directive 21 (PPD-21) Critical Infrastructure Security and Resilience, which establishes national U.S. policy on infrastructure security and resilience, designates the energy sector as "uniquely critical due to the enabling functions [it provides] across all critical infrastructure sectors." Consequently, electrical power grid resilience efforts should consider the infrastructure systems and community functions that depend upon electrical power for their operations.

Over the past decade Sandia National Laboratories (Sandia) has actively developed an infrastructure resilience modeling an analysis program. A core element of this program is the modeling of infrastructure dependencies on electrical power and the consequences that can result from power outages. Grid dependencies have been modeled for transportation networks [2], military installations [3], telecommunications [4], health care [5,6], and manufacturing supply chains [7].
This paper describes two of these modeling efforts:

- Hospital emergency and evacuation planning for potential power outages; and
- Modeling the impact on chemical manufacturing supply chains that result from power outages caused by hurricanes.

The following sections describe the analysis goals that motivated the efforts and the modeling approaches employed. The paper concludes with a discussion on how these kinds of efforts can be used to inform grid resilience planning efforts.

Emergency Planning for Hospitals and Power Outages

In 2012 Superstorm Sandy caused significant damage across the 10 states it affected, and one of the most challenging impacts was the evacuation of 6 hospitals in New York. These hospitals had backup generators for potential power outages, but flooding within the hospitals caused the backup generation system to fail. Consequently, many of the hospitals were required to evacuate without power in the middle of the storm. These evacuations occurred only a year after Hurricane Irene caused similar hospital evacuations in New York.

Subsequent studies looking into decision making processes for these evacuations noted the many complex factors that make the evacuation decision so difficult [8]. To address that challenge, Sandia partnered with the Veteran's Health Administration and the Veterans Emergency Management Evaluation Center to create a hospital evacuation decision support model. The evacuation model is made up of two primary components:

1. A patient care hierarchy: the hierarchy represents the complex structure of dependencies of patient care upon hospital operations and resources.; and
2. A constrained optimization model: this mathematical formulation compares the patient care requirements to the expected resource availability to estimate the hospital's ability to provide adequate care to its patients throughout the projected duration of infrastructure disruption.

Within the hierarchy, the analyst can identify which key equipment, operations, and services are affected by a loss of power. With the optimization model, the analyst can determine how hospital operations and staff adapt to cope with the loss of power and whether an evacuation (full or partial) is necessary to provide sustainable patient care.

Patient Care Hierarchy

The hierarchy itself represents operational dependencies within a hospital and consists of six elements: patients, services, functions, resources, utilities, and consumables (Figure 1). Patients occupy the apex of the hierarchy. Multiple patient types can be represented (eg, intensive care unit (ICU) patient, surgical patient, etc.). Adequate patient care consists of a specified minimum set of service requirements which vary according to the patient type. Services support patient care directly and consist of hospital activities provided by various hospital departments required to provide medical care to patients. The next level down in the hierarchy, functions, directly supports services. Functions include medical procedures for individuals and non-medical activities that contribute to patient care. Resources, utilities, and consumables are the tangible commodities, equipment, and people required to perform functions and services. Resources and consumables differ in that resources are reusable and consumables are not. Utilities are a special class of resource that includes water, electricity, and communications. The hospital evacuation model is designed to incorporate specifically the consequences of the loss of utilities on a hospital's viability.

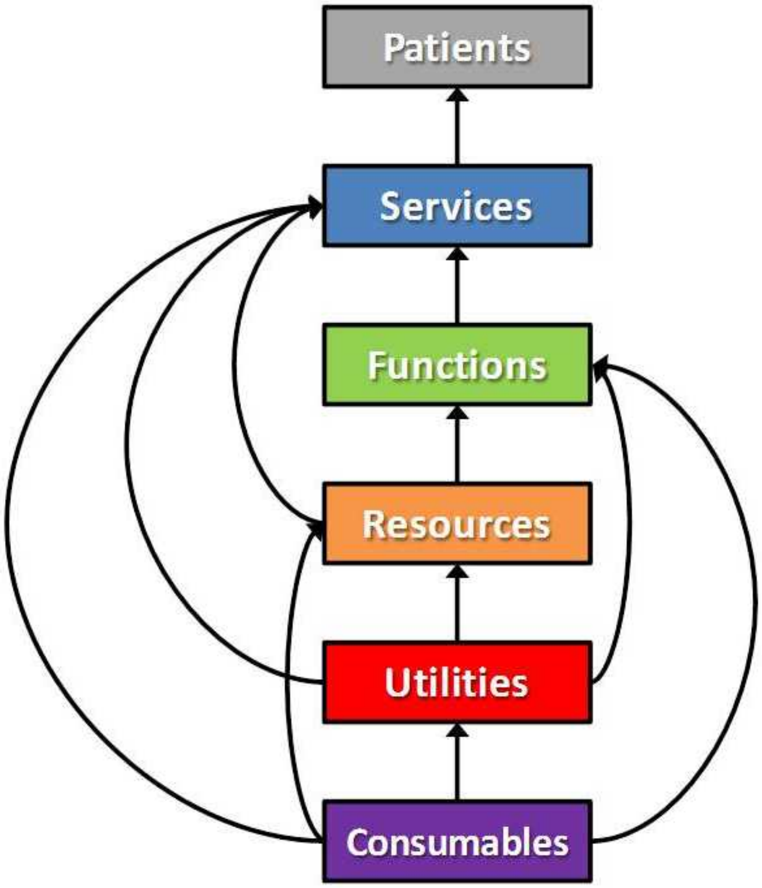

Figure 1: Patient Care Hierarchy

Loss of utilities can compromise a hospital's ability to provide patient services. If hospitals do not have temporary alternatives to sustain activities, the disruption of these critical utilities will affect important services needed for patient care severely. Temporary alternatives, such as backup power generators generally depend upon consumables (ie, generators require fuel). Because hospitals typically maintain limited quantities of consumables, their ability to sustain adequate patient care over time is determined by their stores of required consumables, the rate of consumption, and at what point in the consumable restocking schedule a disruptive event occurs.

Substitution possibilities exist to help the hospital cope with loss of infrastructure services. For example, some medical devices may be able to run using batteries when utility power is lost. The limited capacity of batteries means that they can serve only as a temporary substitute. Backup generators provide another substitution for utility power, but they are dependent upon fuel supplies. If the duration of the power outage exceeds the capacity of fuel stored at the hospital, additional fuel will need to be transported to the hospital. Although substitutions provide hospitals with a means for handling disruptions, their use generally presents an alternative set of requirements that must be met.

The Optimization Model

The model's resource-constrained optimization formulation is used to address the following questions:

1. Can the hospital provide adequate care for its patients in the event of a power outage?
2. If adequate care cannot be provided for all patients, how many patients need to be evacuated, and which patients should be prioritized for evacuation?
3. How long will the evacuation take?
4. What resources are required to ensure adequate care is provided for all patients, both those requiring evacuation and those sheltering-in-place?
5. What are the primary resource constraints (eg, staff, equipment, and so forth) that have a detrimental effect on patient care and/or evacuation?

The model is designed to maximize the number of patients in the adequate-care state. To compensate for a power loss, the model searches among defined substitution options to provide required functions and services, making the best use of available resources, utilities, and consumables to maximize patient care. In the event that available substitution options that meet minimum patient care requirements do not exist, the model will recommend a partial or full evacuation.

Example Results

Vugrin et al. [5,6] have demonstrated how this modeling approach could be used for emergency planning and resilience analysis from a hospital perspective, but additional opportunities exist for the use of this kind of model for grid resilience actitivities. Consider a hypothetical hospital and its intensive care unit (ICU) with a patient care hierarchy as shown in Figure 2. (Model parametrization specifics can be found in Appendix A of [5]). The ICU contains 16 patients, and 6 of them require ventilation support to breathe and are therefore deemed to be "high risk." The remaining 10 patients do not require ventilation and are therefore deemed "low risk."

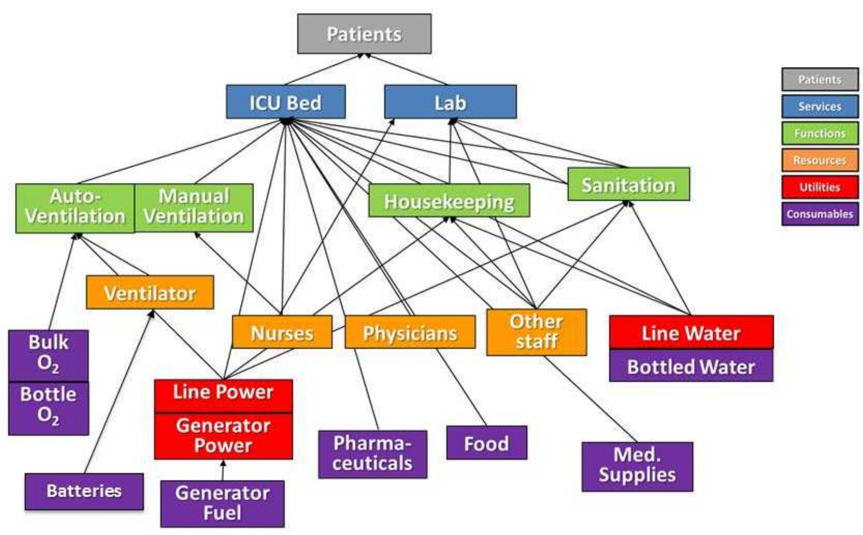

Figure 2: Dependency Mapping for Illustrative Intensive Care Unit

Figure 3 shows the results of a parameter study that included a series of power outage scenarios that varied according to the utility power outage duration and backup power generation capacity. Power outage duration is varied between and 2 and 34 hours; backup generation capacity is specified to be relative to nominal needs, i.e., 0 indicates no backup generation is operational and 1.0 indicates backup generation entirely replaces utility power during the outage.

The model recommends evacuations for high risk ICU patients even for relatively short outages (4 or more hours), regardless of backup generation capacity (Figure 3a); however, evacuation of the low risk ICU patients is only recommended for the most extreme outages that involve longer outages and less backup generating capacity (Figure 3b). When these results are combined (Figure 3c), one can develop an overall strategy for evacuation of the ICU patients:

- Evacuate the high risk ICU patients for outages expected to exceed 4 hours;
- When the hospital has backup generation capacity that can meet at least 50% of nominal load requirements, evacuate high and low risk patients in the event that the outage is expected to last more than 30 hours;
- If the hospital has backup generation capacity that can meet between 10 and 50% of nominal load requirements, evacuate high and low risk patients in the event that the outage is expected to last more than 18 hours;
- If the hospital has backup generation capacity that cannot meet at least 10% of nominal load requirements, evacuate high and low risk patients in the event that the outage is expected to last more than 8 hours.

This information is not only beneficial to hospital emergency planners, but it can also be used to inform power grid resilience planning activities. The potential uses are described in the Conclusions section.

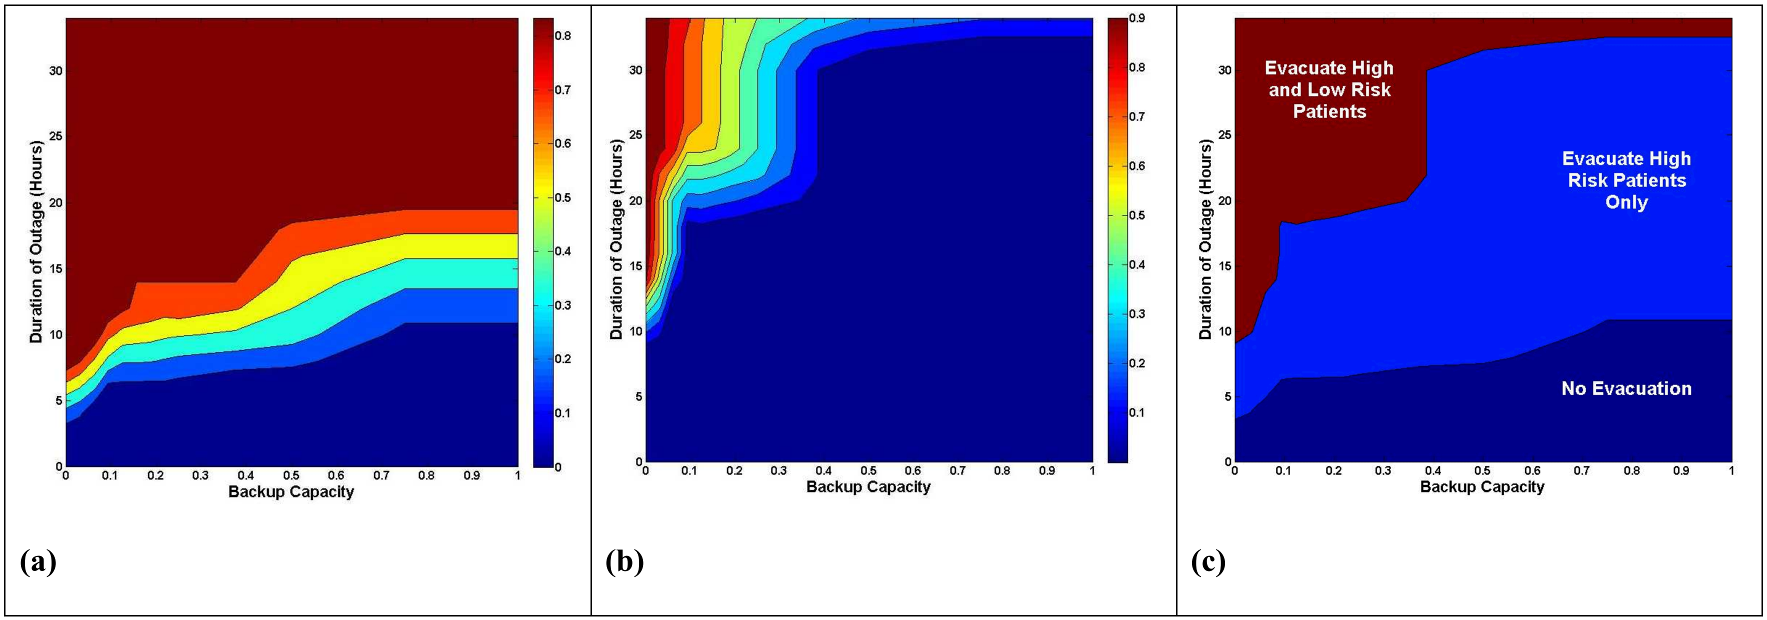

Figure 3: Recommended Fraction of Patients to Evacuate for Power Outage Scenarios that Vary by Outage Duration and Backup Generator Capacity: (a) high risk, (b) low risk and (c) total patient populations

Chemical Production Impacts from Hurricane-Related Power Outages

As a founding partner in the National Infrastructure Simulation and Analysis Center (NISAC), Sandia provides "strategic, multidisciplinary analyses of interdependencies and the consequences of infrastructure disruptions across all 16 critical infrastructure sectors" [9] to the U.S. Department of Homeland Security (DHS). Because the chemical sector supplies more than 70,000 chemicals to 350,000 firms for use in the manufacturing, energy, chemical, transportation, public health, and other sectors, the chemical sector has been designated as 1 of 16 critical infrastructure sectors in the United States [10]. A high density of chemical manufacturing facilities, especially ones producing petrochemicals, are located along the Gulf Coast, making petrochemical supply chains vulnerable to disruption by hurricanes. Consequently, DHS funded Sandia to develop a modeling and analysis capability to estimate the potential consequences of a hurricane on domestic petrochemical production.

Figure 4 illustrates the overall process for estimating the consequences of power outages on chemical manufacturing supply chains. The process begins by characterizing the expected hurricane. Hurricane forecasts by the National Oceanic and Atmospheric Administration are used to estimate windspeeds, the hurricane's path, geospatial location, and other features. These data, in addition to electrical power infrastructure data (lines, stations, geospatial locations, etc.,) are passed to a series of models: CICLOPS estimates the probability that an outage will occur in a specific location and EPRAM estimates the time to restore power by location*. These data are then passed along to a supply chain model. The supply chain model has two primary components: the Chemical Data Model and the the N-ABLE™ agent-based, microeconomics simulation tool.

The CDM is a database of domestic and foreign chemical plants, chemical productions, commodity flows, and chemical infrastructure (for example, pipelines, rail networks, and water transport networks). In 2009, the CDM contained data for almost 4000 domestic and foreign consumers and producers of 63 commodity petrochemicals. Each of the firms in the CDM either makes a primary feedstock petrochemical (benzene, toluene, ethylene, propylene, xylene, o-xylene, p-xylene), converts these chemicals into other petrochemicals, or produces other chemical and non-chemical products based on these chemicals. Using the stoichiometric (chemistry-based) and
other production "recipes" for each chemical, the CDM can identify the basic relationships between these chemicals.

Sandia uses the petrochemical CDM in conjunction with the N-ABLE™ agent-based, microeconomics simulation tool to simulate disruptions to the petrochemical sector from various types of disturbances. N-ABLE™ is a collection of tools to perform supply chain analysis, the analysis of ways individual firms within multi-tiered, multi-product economic systems purchase input goods, produce products, sell them in markets, and ship them via different modes of transportation. The actions of individual plants are the primary units of analyses in N-ABLE™.

* CICLOPS and EPRAM have been developed at Los Alamos National Laboratory (LANL) for the NISAC program [11]. When chemical supply chain analyses are performed, staff from LANL exercise these models and provide outputs to staff at Sandia.

Each agent-based enterprise firm is composed of buyers, production workers, supervisors, sellers, and strategic planners who conduct their real-world analog tasks within the enterprise and among enterprises. Using a generic data-driven object structure, firms representing the range of economic activity in a supply chain (such as manufacturing, transportation, and consumption) can be modeled by specifying such things as production functions, buying and selling behaviors, inventory capacities, and long-term strategic planning. Production decisions are made by a production manager, independently of input costs. The production manager adjusts production based on the inventory of finished goods and the presence of market signals (orders). Every day, the plant orders enough to raise its inventory position to a predetermined level, taking into account expected usage patterns and historic averages of delivery times. Entire supply chains are constructed from collections of firms, based on this enterprise design, with each participating firm interacting with others through markets and physical infrastructure. When applied to the petrochemical CDM, N- ABLE™ simulations can provide dynamic information about how disruptions affect the petrochemical sector.

In the context of hurricane disruption analyses, the analyst use the output from CICLOPS and EPRAM to specify the expected duration of power outages. Using the the CDM, N-ABLE™ specifies that the facilities located within the outage region as "non-operational" for the duration. Facilities located outside of the duration are modeled as attempting to continue operations; however, impacts to chemical production are not merely calculated by aggregating the production capacities of the facilities in the outage areas. Instead, N-ABLE™'s agent-based formulation enables it to describe how customers of the closed plants will find new suppliers, the higher transportation costs associated with those suppliers; the use of chemical substitutes; and the implementation of different production technologies and recipes to adapt to a disruption.

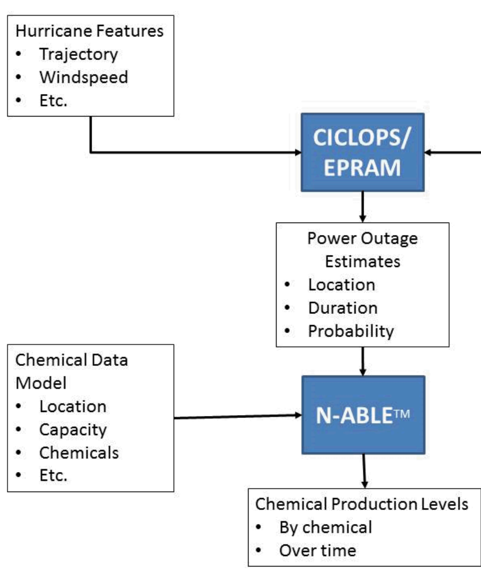

Figure 4: Process for Estimating Effects of Hurricane-Related Power Outages on Chemical Production

Figure 5 shows estimated power outage regions for a hypothetical hurricane scenario that was designed to resemble when Hurricane Ike made landfall near Houston, TX in 2008. Figure 2a shows the extent of expected outages, and Figure 5b shows the location of petrochemical assets overlaid with the outage contours. Almost 1400 firms, representing almost 600,000 daily short tons of product supply are potentially affected by the power outages. On average, petrochemical facilities within the affected region are projected to be without power for 23 days. It is common practice for Gulf Coast petrochemical production facilities in the projected path of a hurricane to shut down operations 48 hours prior to hurricane landfall. Hence, N-ABLE^{TM} was parameterized for a 365 day period under the assumption that facilities within the 50-75% and >75% probability of power outage regions are offline for 25 days (days 201-225). Figures 6a and 6b show chemical production levels for two sets of petrochemical supply chains, vinyl acetate monomer (VAM) and ethylene, respectively. Because of the location of VAM production facilities and its relatively simple supply chain structure, VAM production is almost completely stopped for the entire outage period. Production for ethylene supply chain does not realize its maximum losses until approximately midway through the outage.

These data and analyses can be useful for chemical sector emergency planning , but they can also be used to inform power grid resilience planning activities. The potential uses are described in the Conclusions section.

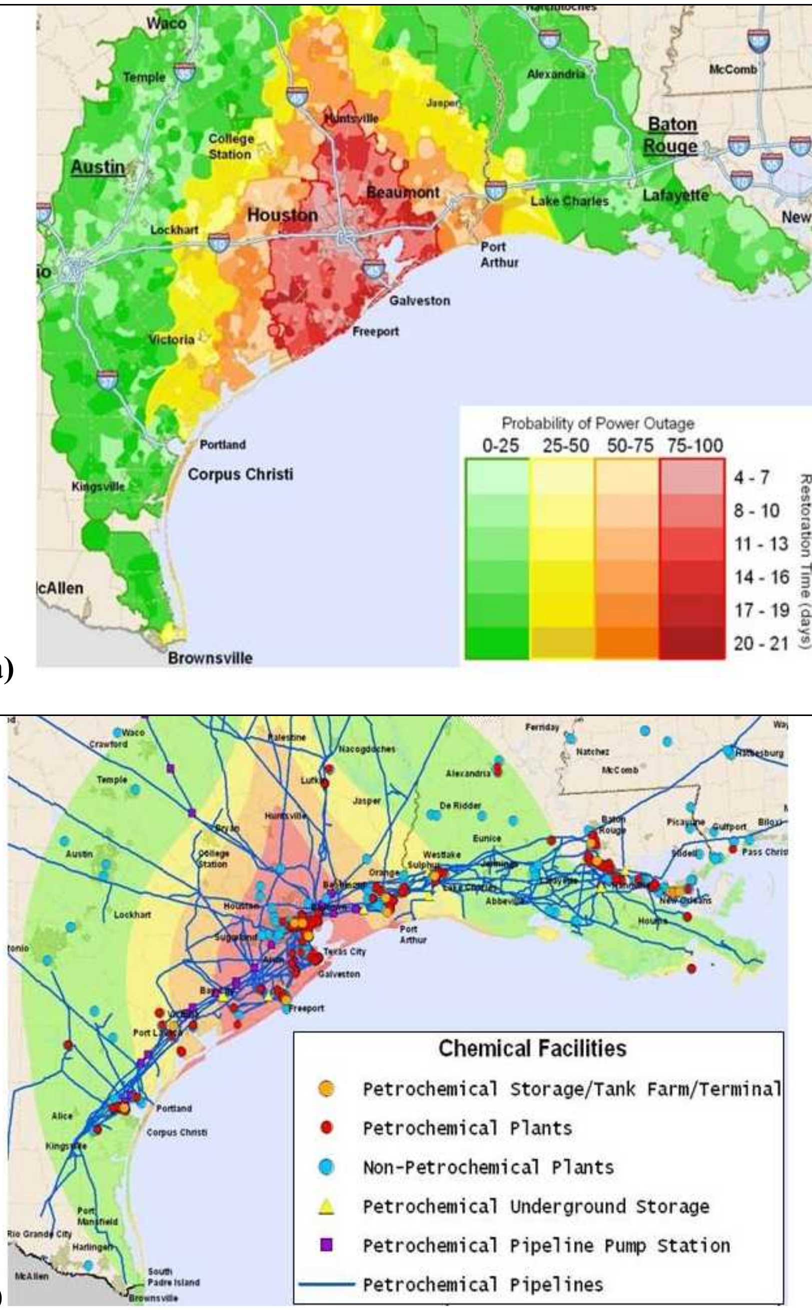

(a)

(b)

Figure 5: Project Grid Outage Regions: (a) probability and duration and (b) with chemical facilities.

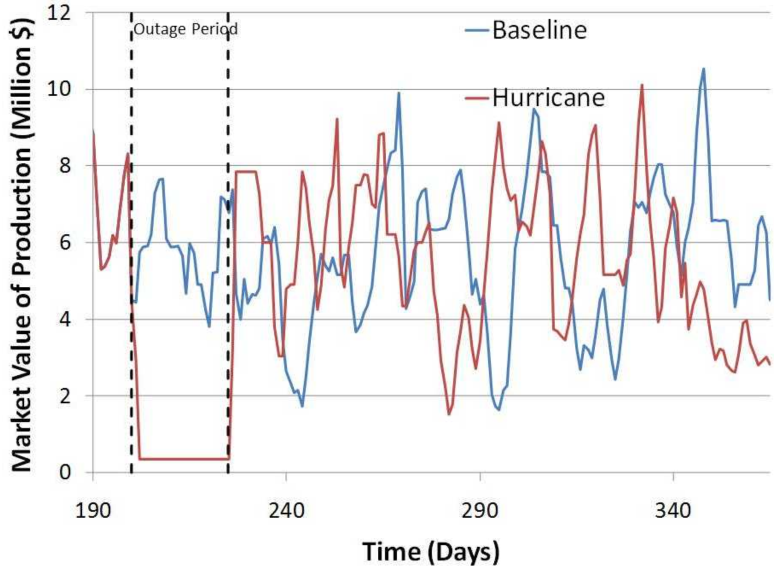

(b)

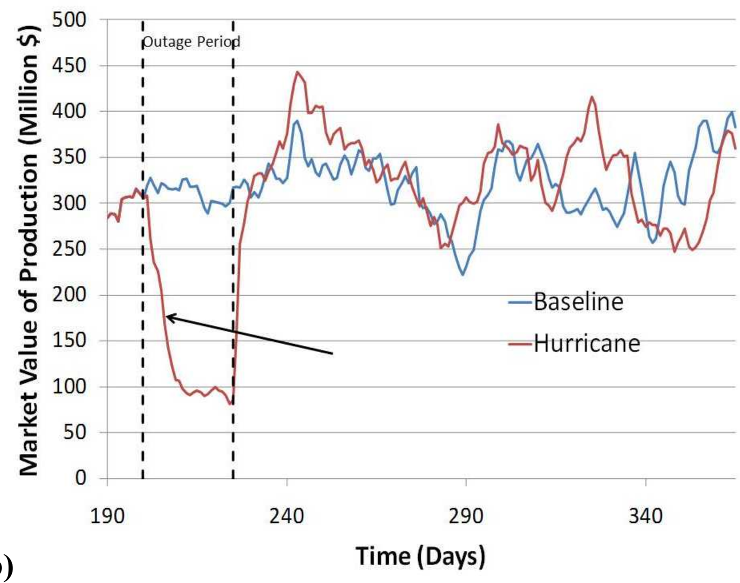

(b)

Figure 6: Project Impacts to Chemical Production (measured in market value of production): (a) vinyl acetate monomer and (b) ethylene supply chain.

Conclusions: Informing Grid Resilience with Infrastructure Dependency Modeling

As previously noted, the development of grid resilience metrics and analysis tools is a current area of research [1]. Many of these activities recommend the use of grid outage consequences to be the basis grid resilience metrics. These metrics would be the drivers for informing grid resilience planning, investment, response, and design activities [e.g., 1,3]. Many different consequences of categories have been suggested, including total time to recovery, power demand not met during the disruption (e.g., measured in total MWh), number of people without power (e.g., in total person-hours), cost of response and recovery activities, grid revenue losses, economic impacts, etc. One category of proposed metrics includes impact to dependent critical infrastructure (e.g., number of emergency services facilities without power, business disruption, etc.) These metrics could be estimated, at least in part, by infrastructure dependency modeling efforts like the ones described above. For example, if one was interested in improving grid resilience to lessen impacts to emergency response facilities such as hospitals, the aforementioned hospital model could be used to quantify the effect of proposed mitigation options by measuring the impact they had on decreasing the number of required evacuations. If business interruptions and economic impacts were a chosen metric for grid resilience, the chemical supply chain modeling capabilities could similarly be used to measure the estimated impact that grid outage mitigation options had on decreasing impacts to the chemical sector.

Not only can infrastructure dependency models measure the potential effectivemess of proposed solutions, they can also help identify solution strategies. For example, in the chemical supply chain example described in the previous section, one can conclude that the VAM supply chain is far more fragile than the ethylene supply chain. Hence, grid hardening may be the preferred option to prevent power outages and keep VAM production facilities online. Additionally, much of the losses for the ethylene could be avoided by merely shortening the outage durations. Expediting power restoration activities to keep power outages to no more than 7 or 10 days would significantly benefit the ethylene supply chain. See Vugrin et al. [7] for additional information on how Sandia has used chemical supply chain modeling capabilities for resilience analysis.

Similarly, infrastructure dependency modeling can provide resilience goals. For example, in the hospital evacuation model, restoration of power within 4 hours could be specified as a resilience goal in order to prevent the evacuation of the high risk ICU patients. A secondary target could be restoration of power within 8 or 18 hours for hospitals with <10% and 10-50% backup generation capacity, respectively. Thus, even though the scenario and hospital used in this example is notional and thus these results are merely illustrative, one can see how this type of modeling can time-to-restoration targets for grid resilience planning activities.

Significant work remains to establish formal methods and metrics for grid resilience. However, infrastructure dependency modeling can provide valuable information to inform grid resilience related research and activities.

References

1. Watson Jean-Paul, et al. Conceptual Framework for Developing Resilience Metrics for the Electricity, Oil, and Gas Sectors in the United States, Sandia National Laboratories Technical Report SAND2014-18019. 2014.
2. Brown Nathanael JK, et al. A Modeling Framework for Investment Planning in Interdependent Infrastructures in Multi-Hazard Environments, Sandia National Laboratories Technical Report SAND2013-7755. 2013.
3. Vugrin, Eric D., et al. "A mission resilience assessment methodology for infrastructure disruptions." 2012 IEEE International Carnahan Conference on Security Technology (ICCST). 2012.
4. Conrad, Stephen H., "An Overview of Energy and Telecommunications Interdependencies Modeling at NISAC," in *Computational Models of Risks to Infrastructure*, D. Skanata and D.M. Byrds Eds., IOC Press, 2007.
5. Vugrin, Eric, et al. "Modeling Hospitals' Adaptive Capacity during a Loss of Infrastructure Services." *Journal of healthcare engineering* 6.1 (2015): 85-120.
6. Vugrin, Eric D., et al. "Modeling evacuation of a hospital without electric power." *Prehospital and disaster medicine* 30.03 (2015): 279-287.
7. Vugrin, Eric D., Drake E. Warren, and Mark A. Ehlen. "A resilience assessment framework for infrastructure and economic systems: Quantitative and qualitative resilience analysis of petrochemical supply chains to a hurricane." *Process Safety Progress* 30.3 (2011): 280-290.

8. Ricci, Karen A., et al. "Evacuate or shelter-in-place? The role of corporate memory and political environment in hospital-evacuation decision making." Prehospital and disaster medicine 30.03 (2015): 233-238.
9. Sandia National Laboratories, "National Infrastructure Simulation and Analysis Center," accessed August 18, 2016 at http://www.sandia.gov/nisac/ .
10. Obama, Barack O., Presidential Policy Directive 21 (PPD-21) Critical Infrastructure Security and Resilience, accessed August 18, 2016 at https://www.whitehouse.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil .
11. Stamber Kevin L., et al. *Modeling and Simulation for Homeland Security*. 20th Proceedings of the International Congress on Modelling and Simulation, Adelaide, Australia, 1–6 December 2013.