

DISTRIBUTED AND OPTIMAL RESILIENT PLANNING OF LARGE-SCALE INTERDEPENDENT CRITICAL INFRASTRUCTURES

Linan Huang
Juntao Chen
Quanyan Zhu

Department of Electrical and Computer Engineering
New York University
2 MetroTech Center, Brooklyn
New York, 11201, USA

ABSTRACT

The complex interconnections between various critical infrastructure sectors make the system of systems (SoS) vulnerable to failures and highlight the importance of robustness and resilience. To this end, we first establish holistic probabilistic networks to model the interdependencies between infrastructure components. To capture the underlying failure and recovery dynamics, we further propose a Markov decision processes (MDP) model in which the response policy determines a long-term performance. To address the challenge of a large dimensionality, we exploit the sparsity of the network interconnections and solve an approximate linear program by the variable elimination, which leads to a distributed control policy under mild assumptions. Finally, we use a case study of the interdependent power and subway systems to corroborate the results and show that the optimal resilience resource planning and allocation can reduce the failure probability and mitigate the impact of failures caused by natural or artificial disasters.

1 INTRODUCTION

Presidential Policy Directive 21 (PPD-21) identifies 16 critical infrastructure sectors including energy, communication and transportation systems as so vital that their malfunctions and incapacitation can lead to an enormous economic loss and public safety threat (Obama 2013). Driven by the recent advances in information and communication technologies (ICTs) and the Internet of Things (IoTs), these sectors become highly interconnected, enabling faster information exchange and a higher level of situational awareness for real-time operations. For example, (Zimmerman et al. 2016) and (Zimmerman et al. 2018) investigate the interdependency between food, energy and water. (Gao et al. 2014) shows that on the one hand, supervisory control and data acquisition (SCADA) systems control the power generation, transmission, and distribution according to the monitoring information from the communication sector. On the other hand, the energy sector provides the power to guarantee the normal operation of the SCADA systems.

Protection of interdependent critical infrastructures (ICIs) against natural or human-made disasters is essential for providing support to a reliable and sustainable economy. The increasing interdependencies, however, make systems vulnerable to these catastrophes because a failure of one system will propagate to others and cause cascading failures without appropriate, rapid responses. Lessons of the September 11 attack, hurricanes Sandy and Irma have highlighted that protection and prevention against such disruptions are not always possible. Hence a paradigm shift to emphasize the preparedness and response is indispensable to enhance the resilience of ICIs as shown in Fig. 1.

Many works have investigated the modeling, simulation, and design of interdependent critical infrastructures (Ouyang 2014; Zimmerman et al. 2017). Network flows (Lee II et al. 2007) and the interacting dynamic model (Rosato et al. 2008) are used to assess and manage the risks. Built upon the modeling



Figure 1: The systematic resilience consists of three phases. First is the robustness against disruptions to prevent failures. Second is the impact mitigation to guarantee the essential services under failures. Third is the prompt response and recovery. A paradigm shift from protections to the mitigation and recovery is indispensable for catastrophes.

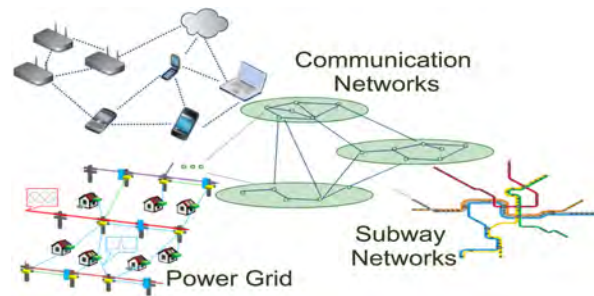


Figure 2: Multi-layer networks representing interdependent infrastructures such as power grid, subway and communication networks. They can be viewed as a large-scale aggregated network. The directed link between two nodes indicates either physical, cyber or logical influences from the source to the sink.

of interdependent infrastructures, (Yuan et al. 2015) focuses on designing a resilient control for power grid. (Wang et al. 2013) has proposed an L1Simplex architecture to address both physical and software failures in the cyber-physical systems. If an intelligent attacker rather than a random accident causes the network component failure, (Zhu and Basar 2015; Huang et al. 2017b) apply dynamic game frameworks to incorporate the strategic participants who adapt to the defense policy.

Despite these pioneering works, challenges still remain to build a holistic model that captures not only distinct characteristics of individual infrastructures but also complex dependencies among sectors due to geographical, logical, cyber-physical, or human connections (Rinaldi et al. 2001). Without a holistic framework, we cannot integrate the data or models of individual sectors to make a globally optimal planning and defense decisions. For example, system operators should not only make plans for normal operations of their own system but also be concerned about the possible failures of other dependent sectors and prepare for the potential influences. The framework should also be general enough to include various sectors and straightforward enough to enable real-time planning for large-scale SoS, which motivates the MDP model in Section 2 to abstract the complex networks, unify the multi-layer interdependences, and design resilient control mechanisms.

Moreover, the probabilistic and persistent transition of system status requires system operators to prepare for the uncertainties in a dynamic perspective and avoid the following two extremes. First is the blind optimism that satisfies with the current normal operation and makes no sufficient contingency plan for probabilistic failures. Second is the pursuit of absolute security that would result in a huge inspection and maintenance cost. Thus, it is crucial to make proactive and cost-effective controls based on the real-time status of SoS to reduce the long-term risk economically and globally.

In this paper, we first establish a holistic network model to capture the relationships between nodes within an individual infrastructure and across the infrastructures. We use a probabilistic model to assess the impact of the couplings between different components in a system, in which the failure probability of one component depends not only on its state but also on state of its connected ones. One fundamental characteristic of the infrastructures is their dynamics in real-time operations (Zimmerman et al. 2018). For example, the load on a bus in a power grid changes due to the time-varying demands, and the bus would encounter a load shedding if the demand exceeds the supply. The concept of resilience involves the phase of recovery of infrastructure after disastrous events. This process is naturally a dynamic one (Chen et al. 2017). Hence, building on the probabilistic failure model, we propose a Markov decision processes (MDP) model to take into account the underlying dynamic processes of ICIs and develop resilient mechanisms for heterogeneous interdependent infrastructures.

The MDP model has naturally a large state-space due to the large-scale feature of infrastructure networks and exponential increases of the system states with respect to the number of components (Huang et al. 2017b; Hayel and Zhu 2015). The curse of dimensionality of the MDP makes the control design computationally intractable even in a moderate-size system with 15 nodes. To this end, we reduce the dimension by first using basis functions to reformulate the problem as an approximate linear program (ALP) and then leveraging the sparsity of the interconnections to simplify the ALP with factored graphs. The approximated problem yields a low computational complexity. Under mild assumptions on the dependence structures and the separability of reward and control cost functions at each node, the resilient control policies can be designed in a distributed fashion. We can show that the system performance under the distributed control scheme is equivalent to the one achieved under the centralized one. We corroborate the established MDP model and analytical results with case studies of an interdependent power and subway network.

The remainder of the paper is organized as follows. Section 2 establishes the MDP model of the interdependent infrastructures. Problem analysis is given in Section 3. The structural results including the optimal and distributed resilient control policies are obtained in Section 3.3. The case study of a power and transit network is presented in Section 4, and Section 5 concludes the paper.

2 SYSTEM MODEL

In this section, we introduce the network model to represent the complex interdependent SoS and the MDP model to produce resilient plannings of ICIs. The resilient plannings are both proactive, i.e., maintain high-risk nodes in advance of failures to reduce the future risks, and cost-effective, i.e., delay the repair of low-influence nodes due to the restrained or costly resources.

2.1 Network Model

ICIs can be modeled by a multi-layer network as shown in Fig. 2, where nodes represent heterogeneous components in multiple infrastructures such as buses in power grid, subway stations in transportation, and base stations in the communication network. The links connecting nodes in the network represent the dependencies either homogeneous, e.g., the failure of one subway station can disrupt the service of the next station, or heterogeneous, e.g., the power loss of the grid leads to the shutdown of subway stations. Moreover, the inter-dependency implies the ping-pong effect. For example, the shutdown of subway stations can, in turn, slow down the power recovery because maintainers may take extra time to commute. The vicious spiral can lead to the cascading failure. Without loss of generality, we use an aggregated network $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ to capture the heterogeneous components in the multi-layer network, where \mathcal{N} denotes a set of nodes and \mathcal{E} is a set of links.

2.2 MDP Model

Each infrastructure node $i \in \mathcal{N}$ in the network \mathcal{G} has a state X_i representing its working status. State X_i admits values from a binary set $\{0, 1\}$. Specifically, $X_i = 1$ indicates that node i works normally and $X_i = 0$ means otherwise. In a network of $|\mathcal{N}| = n$ nodes, the system state can be denoted by $X = [X_i]_{i \in \mathcal{N}} \in \mathcal{X} := \{0, 1\}^n$. Due to the interdependencies between different nodes, the failure probability of one component depends not only on its own state but also on its connected ones. Similarly, each infrastructure node $i \in \mathcal{N}$ can be controlled by taking action $A_i \in \{0, 1\}$. Specifically, $A_i = 1$ represents repair (if node is faulty) or maintenance (if normal) of node i , and $A_i = 0$ indicates that no action is applied. The system action is a vector $A = [A_i]_{i \in \mathcal{N}} \in \mathcal{A} := \{0, 1\}^n$ including all the actions of each node.

In order to design resilient mechanisms for the real-time operation of ICIs, we use an MDP model to capture the dynamics of the infrastructure states as well as the underlying failure and recovery processes of each node. Our objective is to determine a stationary policy $\pi : \mathcal{X} \mapsto \mathcal{A}$ that yields the optimal control of nodes to achieve the largest long-term benefit of ICIs.

To achieve this goal, we further define the reward function $R: \mathcal{X} \times \mathcal{A} \mapsto \mathbb{R}$ and the objective function $V_\pi(x): \mathcal{X} \mapsto \mathbb{R}$. The reward function measures the performance of the ICIs at one step by taking the number of working nodes into account. The objective function accumulates the obtained reward at each step over an infinite horizon: $V_\pi(x) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t R(X^t, \pi(X^t)) | X^0 = x]$, where $\gamma < 1$ is a positive discounted factor, x is an initial state, and X^t is a random variable denoting the system state at stage t . The state transition probability follows the Markov property, i.e., the system state at the next time step only depends on the current system state and the adopted action. For example, $P(x'|x, a) = \mathbb{P}\{X^{t+1} = x' | X^t = x, A^t = a\}$, $\forall t$, denotes the transition probability from state $x \in \mathcal{X}$ to $x' \in \mathcal{X}$ under control $a \in \mathcal{A}$. Note that transition probability P captures the interdependencies between connected infrastructure nodes. Further, the value function $V(x) = \max_\pi V_\pi(x) := V_{\pi^*}(x)$ is the maximum achievable economic benefit of the ICIs starting from state x under the optimal policy π^* . A larger value of $V(x)$, $\forall x$, indicates a better performance of ICIs, and the network is resilient to failures by using the control policy π^* .

3 PROBLEM ANALYSIS

In this section, we introduce a sequence of solution techniques of the MDP problem and tackle the curse of dimension by exploiting the sparse nature of the large-scale infrastructure networks. Approximation with the factor graph and the variable elimination technique leads to equivalent distributed policies.

3.1 Linear Programming

One way to characterize the optimal policy of the MDP is via dynamic programming with the Bellman operator $(TV)(x) := \max_\pi R(x, \pi(x)) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, \pi(x))V(x')$, $\forall x$. Then, the value function $V(x)$ is the fixed point of T , i.e., $V(x) = (TV)(x)$, $\forall x$. Linear programming (LP) provides a convenient and efficient approach to solve the Bellman equation with value functions as LP variables, i.e.,

$$\begin{aligned}
 \text{(LP)} : \min_{V(x)} : & \sum_{x \in \mathcal{X}} \alpha(x)V(x) \\
 \text{s.t.} & V(x) \geq R(x, a) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, a)V(x'), \\
 & \forall x \in \mathcal{X}, \forall a \in \mathcal{A}.
 \end{aligned} \tag{1}$$

LP has the benefit to be easily extended to include both endogenous and exogenous constraints. For example, if we add an endogenous action constraint $\sum a_i = 1$ which limits to single repair due to the limited resources at each stage, the LP formulation remains the same with only a reduced dimension of action space \mathcal{A}' as in (Huang et al. 2017a). Similarly, we can directly add exogenous action constraints such as $\sum c_i(a_i) < C$ which gives a bound on the action cost due to limited budgets at each step.

The disadvantage of (1) is that the curse of dimension appears in both state space $|\mathcal{X}| = 2^n := N$ and action space $|\mathcal{A}| = 2^n := N$ when solving the LP directly. Therefore, we use ALP and FMDP method to reduce the computation complexity (Guestrin et al. 2003).

First, we use the linear value function (Farias and Pucci 2002) to approximate the value functions $V(x) = \sum_{i=1}^k w_i h_i(x)$, where h_i and w_i is the i^{th} basis function and its weight, respectively. Hence the number of LP variables is reduced from $N = 2^n$ to k and we obtain the ALP with new variables $\{w_1, w_2, \dots, w_k\}$.

$$\begin{aligned}
 \text{(ALP)} : \min_w & \sum_{x \in \mathcal{X}} \alpha(x) \sum_{i \in \mathcal{N}} w_i h_i(x) \\
 \text{s.t.} & \sum_{i \in \mathcal{N}} w_i h_i(x) \geq R(x, a) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, a) \sum_{i \in \mathcal{N}} w_i h_i(x'), \\
 & \forall x \in \mathcal{X}, \forall a \in \mathcal{A}.
 \end{aligned} \tag{2}$$

Our second step is to reduce the dimension of the transition probability matrix $P(x'|x, a) \in \mathbb{R}^{N \times N \times N}$. The system transition probability $P(x'|x, a)$ can be factored into a multiplication of local transition probabilities

$\prod_{i=1}^n P(x_i|x_{\bar{\Omega}_i})$ where $\bar{\Omega}_i$ is the node set that can affect the local transition probability of node i . With the assumption that the network \mathcal{G} is sparse, we have $|\bar{\Omega}_i| < n$. Notice that the co-domain of each basis function $h_i(x)$ is determined by some subset of \mathcal{X} . The effective domain of $g_i(x, a) := \sum_{x'} P(x'|x, a)h_i(x')$ is also restricted to a set much smaller than \mathcal{X} and a special case is discussed under assumption 1. Define $\bar{g}_i(x, a) := \gamma g_i(x, a) - h_i(x)$, the constraints $\sum_i w_i h_i(x) \geq R(x, a) + \gamma \sum_i w_i g_i(x, a), \forall x, a$ is rewritten in a more compact form $0 \geq \max_{x, a} R(x, a) + \sum_i w_i \bar{g}_i(x, a), \forall a \in \mathcal{A}$.

The reward function $R(x, a)$ can also be factored into a sum of functions with the domain restrained to some subsets of \mathcal{X} . Then, we can use the variable elimination method to eliminate each element of x and a step by step (Huang et al. 2017a).

3.2 Approximate Optimal Policy

The optimal centralized policy under the approximate form of LP is obtained by maximizing the objective function given the initial state x , i.e., $a^* \in \arg \max_{a \in \mathcal{A}} [R(x, a) + \gamma \sum_{x' \in \mathcal{X}} P(x'|x, a)V(x')]$. Similarly, an approximation is adopted to reduce the computation to $k + 1$ multiplications and k summations for any given pair of (x, a) , i.e.,

$$a^* \in \arg \max_{a \in \mathcal{A}} \left[R(x, a) + \gamma \sum_{i=1}^k w_i g_i(x, a) \right]. \quad (3)$$

To solve (3), for every state x , we need to search for action a , which means that we have to compute (3) for $|\mathcal{X}| \times |\mathcal{A}| = 2^{2n}$ times. However, by exploiting the structure of the problem, we can make assumptions in Section 3.3 to achieve a distributed computation of the optimal policy.

3.3 Main Structural Results

To study the interdependent infrastructure network consisting of n nodes, we choose the basis function $h_i(x)$ to be the indicator function of node i 's state, i.e., $h_i(x_i, x_{-i}) = x_i, \forall x_{-i} \in \{0, 1\}^{n-1}$. Hence the number of basis functions is of the same size of the nodes, i.e., $k = n$. Weight parameter w_i shows the importance of node i in the measure of economic benefits.

The advantage of choosing such basis functions is that the value function $V(x)$ at state $x = [x_i]_{i \in \mathcal{N}}$ is approximated by the summation of the weight w_i of all working node $x_i = 1$ at state x . With a larger number of critical nodes functioning, the system achieves a higher reward in the long run. This choice of basis functions is particularly suitable for sparse networks because they primarily capture local effects. We can add new basis functions such as $h_i(x_i, x_j, x_{\mathcal{N} \setminus \{i, j\}}) = x_i \cdot x_j, \forall j \in \bar{\Omega}_i$, which is necessary when non-local effects play a major role. We make the following reasonable assumptions to achieve a distributed repair policy.

Assumption 1 The local transition probability of node i is not affected by the control value of other nodes, i.e., $P(x_i|x_{\bar{\Omega}_i}, a) = P(x_i|x_{\bar{\Omega}_i}, a_i)$.

Under the above assumption, we obtain the following equations:

$$\begin{aligned} g_i(x, a) &= \sum_{x' \in \mathcal{X}} P(x'|x, a)h_i(x') = \sum_{i \in \mathcal{N}} \prod_{j \in \mathcal{N}} P(x'_j|x_{\bar{\Omega}_j}, a)h_i(x'_i) \\ &= \sum_{x'_i} P(x'_i|x_i, x_{\bar{\Omega}_i}, a)h_i(x'_i) = P(X'_i = 1|x_{\bar{\Omega}_i}, a_i). \end{aligned}$$

Assumption 2 We can separate and factor the reward function $R(x, a)$ into the net reward $\sum_{i=1}^n r_i(x_{\bar{\Omega}_i})$ and the action cost $-\sum_{i=1}^n c_i(a_i)$, where $r_i(x_{\bar{\Omega}_i})$ and $c_i(a_i)$ are the local reward and the cost function of node i , respectively.

Theorem 1 A distributed policy $a_i \in \arg \max_{a_i} \gamma w_i g_i(x_{\bar{\Omega}_i}, a_i) - c_i(a_i)$ achieves the same reward obtained by the centralized optimal policy using (3) based on assumption 1 and 2.

Theorem 1 holds as we can represent the RHS of Eq. (3) in a factored form, i.e.,

$$\max_{a_1, \dots, a_n} [R(x, a) + \gamma \sum_i w_i g_i(x, a)] = \sum_{i=1}^k r_i(x_{\bar{\Omega}_i}) + \sum_{i=1}^n \max_{a_i} \gamma w_i g_i(x_{\bar{\Omega}_i}, a_i) - c_i(a_i).$$

The distributed policy only takes into account the states of node $j \in \bar{\Omega}_i$ and reduces the computation complexity from $O(|\mathcal{X}| \times |\mathcal{A}|)$ to $n \cdot 2^{\Omega_i}$. In many applications, the policy designer can only control a limited number of nodes. Thm. 1 still holds as the sparse control condition can be viewed as a special case.

Our next step is to find a pattern to quickly identify each node's optimal action, i.e., repair or not. Thm. 2 achieves the above objective under an arbitrarily given set of controllable nodes without explicitly computing $a_i \in \arg \max_{a_i} \gamma w_i g_i(x_{\bar{\Omega}_i}, a_i) - c_i(a_i)$. With assumption 3, we can find the repair policy thresholds as summarized in Thm. 2, which means that we switch to the opposite action if the condition is not satisfied. Let $\bar{\Omega}_i = \{i, \Omega_i\}$ and $x_{\bar{\Omega}_i} = [x_i, x_{\Omega_i}]$. Besides, we define a shorthand notation $P_{x_{\Omega_i}=j} := P(X'_i = 1 | X_{\Omega_i} = j, A_i = 0)$.

Assumption 3 We further assume that a repair is always effective $P(X'_i = 1 | x, a_i = 1) = 1, \forall i, \forall x$ and the node stays faulty without a repair $P(X'_i = 1 | x_i, x_{\Omega_i}, a_i = 0) = x_i \cdot P_{x_{\Omega_i}=1}, \forall i, \forall x$. In addition, the cascading failure means that $P_{x_{\Omega_i}=1} > P_{x_{\Omega_i} \neq 1}$, where $\mathbf{1}$ represents a vector with all n elements as 1.

Theorem 2 With assumptions 1, 2 and 3, we only repair node i when it is faulty and its action cost is relatively low, i.e., when c_i satisfies $c_i(1) - c_i(0) < w_i$. If $0 < c_i(1) - c_i(0) < w_i(1 - P_{x_{\Omega_i}=1})$, we repair node i even when nodes in $\bar{\Omega}_i$ are all working normally. We choose to repair node i , i.e., $A_i = 1$, when i is working $X_i = 1$ yet other nodes in Ω_i have states $x_{\Omega_i} \neq \mathbf{1}$ under the condition $w_i(1 - P_{x_{\Omega_i}=1}) < c_i(1) - c_i(0) < w_i(1 - P_{x_{\Omega_i} \neq \mathbf{1}})$.

To show Thm. 2, we first look at the case when only one given node is controllable. Since action a_i is binary, we can obtain the optimal policy by comparing which action leads to a larger value. We generalize the one-node case to the case of multiple nodes because the Thm. 1 guarantees that each node's policy can be decided independently of others' actions.

4 CASE STUDY

In this section, we use a case study motivated by Hurricane Sandy, which struck New Jersey on Oct. 29, 2012, resulting in an estimated \$50 billion economic loss and at least 147 direct deaths. It wreaked havoc on critical infrastructures such as power grid, transportation and communication systems. In this case study, we focus on the impact of the power and subway failures in the lower Manhattan and downtown Brooklyn as shown in Fig. 4.

4.1 Power and Subway

The most direct impact of Sandy has been on the power grid. Much of Manhattan south of 39th Street has suffered a massive power cut because of the floods and the high wind. A blast of one Consolidated Edison (ConEd) substation on 14th Street at 9 p.m. on Oct. 29 has intensified the power outage. Besides, ConEd took parts of its grid offline due to the rising water. Load shedding arose when the demand strained the capacity during the storm. The massive damages made ConEd unable to restore most of the power outages in lower Manhattan until 2 November. The whole power system has taken weeks to recover completely, and nearly 8.5 million people have suffered from power outages.

As shown in Fig. 3, the normal operation of subway system relies heavily on a stable power supply. Thus, the subway system has been significantly affected as well. The Metropolitan Transportation Authority (MTA) shuts down the entire subway system from 7 p.m. on Oct. 28 to move trains away from lowland and vulnerable area. Floodwaters brought by the storm surge of Sandy began to enter the subway tunnels and stations after 9 p.m. on 29 October. All seven East River subway tunnels connecting Manhattan,

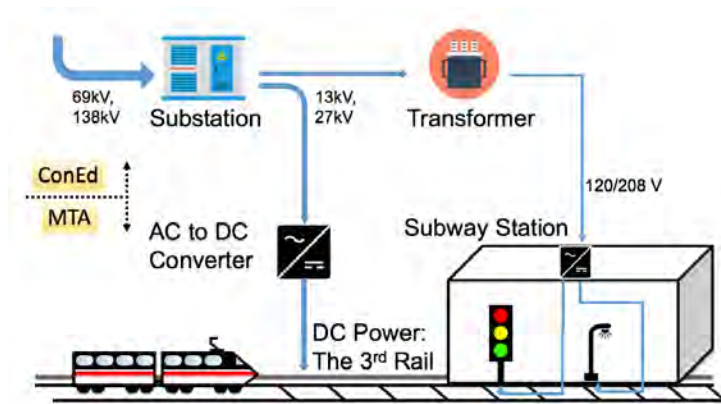


Figure 3: After the high-voltage transmission of 69/138kV, area substations transform the power voltage to 13/27kV, convert it to direct power current and feed it for the 600-volt third rail of the track to propel the train. Besides, transformers in the distribution grid further lower the voltage to 120/208V to power signals, lights, and track switches in subway stations.

Brooklyn, and Queens have suffered from floods a day later. Besides, Hurricane Sandy wiped out tracks on the A train in the Rockaways and severely damaged new South Ferry station.

Since after flooding, water can short-circuit electrical devices including signals, switches, and the electrified third rail, the subway has to remain closed before dewatering. The pump system of the normal drainage does not function during Sandy because of the power cut. Thus, it takes a long time to pump out water with limited backup power generators. After seawater has been pumped out, MTA needs extra effort to clean the salt deposits and debris, inspect, test and repair the electrical components to guarantee the safety of train operations. Unfortunately, not only the resilience of electric power has a significant impact on the recovery of the subway system, the repairing process of electric power systems is strongly influenced by the subway as well because the subway shutdown creates difficulty for ConEd crews to reach faulty components.

4.2 Network Construction

Lessons learned from Sandy has indicated that the lack of sufficient backup power decreases the recovery efficiency due to the failure of conventional pumping systems. Hence, backup generators (e.g., diesel generators) or local energy storage devices should be prepared in case of power failures. For instance, Verizon’s Garden City central office powered by seven fuel cells has saved it from the loss of power (Kwasinski 2013). Thus, Microgrids can be adopted to supply power to local users when the centralized power distribution network breaks down as what has happened during Sandy. In this case study, we construct a 20-node interdependent network illustrated in Fig. 4, where each microgrid node 1-10 represents a local backup generator located at the corresponding subway station 11-20. Transmission lines in red connect the microgrid so that loads of each node can change due to its working status. The blue lines construct three simplified subway lines. Let $X_i = 1$ be the state of normal operation while $X_i = 0$ indicates the failure of node i .

Backups can fail and we set up the failure model of the microgrids based on our previous work (Huang et al. 2017a). The failure of microgrids at subway stations can lead to the meltdown of the public transit system. For example, if subway station 17 is out of service due to the power outage of node 7, the entire line 4/5 has to suspend as trains cannot get through. On the other hand, the subway system affects the microgrids through reward $r_i(x_{\bar{Q}_i})$ of node i . We adopt the proposed framework to illustrate how we achieve the resilience of the microgrids and the subway lines under our optimal policy. For each node i , $A_i = 1$

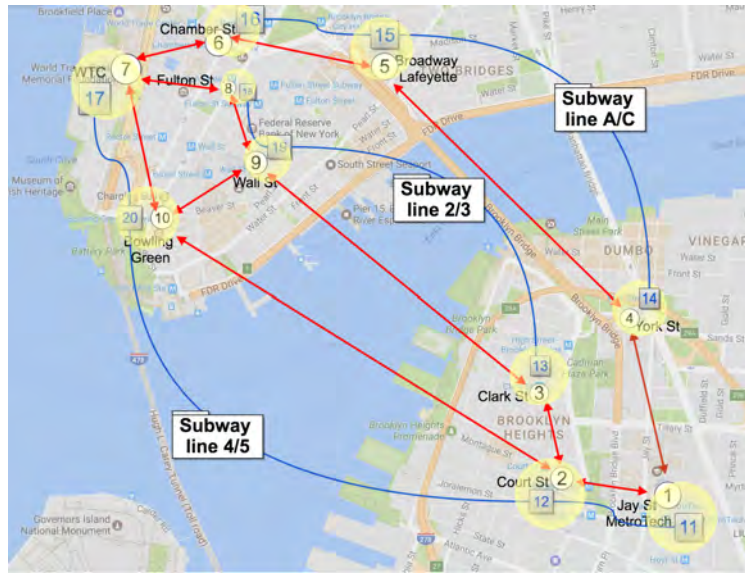


Figure 4: Network representation of the suggested interdependent microgrids and subway system at the lower Manhattan and downtown Brooklyn area. Red and blue links represent power and subway lines, respectively. Power nodes are indexed from 1 to 10 while the coupled subway stations are indexed from 11 to 20. The size of node i is proportional to its reward $r_i(x_{\Omega_i})$, indicated by the people flow density data from the MTA website. Power and subway nodes of the same place inside one yellow circle share the same local reward.

means to repair or maintain the node based on the state of the node. On the other hand, $A_i = 0$ means that no action is taken.

The coupling of the power and subway system exists as an illustration of the interdependent infrastructures. The failure of microgrids at subway stations can lead to the meltdown of the public transit system. The subway system affects the microgrids through reward $r_i(x_{\Omega_i})$ to node i .

4.3 Computation Result

Table 1 shows that the weight of each node obtained by ALP follows the same ordering as the local reward except for node 1 and 5. This result arises from the fact that the topology of the network has less influence on the weight than the local reward. The centralized and the distributed search methods yield the same policy as indicated by Thm. 1 and the optimal policy satisfies the pattern stated in Thm. 2. The optimal policy is not to just myopically repair defective nodes but also maintain some healthy nodes to achieve a long-term system-wide performance. For instance, we observe that the policy of the most critical node (WTC) is not repaired only when node 6 and 7 are both working as shown in Fig. 5. We choose to maintain node 7 when $X_6 = 0, X_7 = 1$ to prevent node 7's outage in later stages. Other neighboring nodes 6 and 8 are not influential on the policy of node 7 because they have relatively small weights as shown in Table 1. As for the computation complexity, it takes the centralized search method about 13843.7 seconds to solve the case after weights w are computed while our distributed search process requires only 2.1 seconds.

4.4 Simulation Result

We choose $X = 1$ as the initial state and show the system resilience under our optimal policy in Fig. 7. Averaged under 50 iterations, the data of the y-axis shows the average number of working nodes and also measures the resilience of the system. Without controls, although the prior failure probability is set to a small number 0.01, cascading failure still occurs, and the survival number of nodes quickly decreases to

Table 1: Results of ALP with all nodes controllable.

Node Index and Name			Local Reward $r_i(x_{\bar{\Omega}_i})$	Weight w_i
1	Jay St.	11	15.3215	18.5749
2	Court St.	12	7.1197	8.7800
3	Clark St.	13	2.4738	3.8601
4	High St.	14	2.9452	3.9617
5	Broadway	15	15.7205	17.5535
6	Chamber St.	16	14.1986	14.2076
7	WTC	17	20.4828	25.2162
8	Fulton St.	18	1.7065	2.6374
9	Wall St.	19	9.3785	11.2510
10	Bowling Gr	20	10.6528	12.8872

zero. On the other hand, our optimal policy repairs the proper nodes so that the system quickly recovers and stays at the healthy state $X = 1$. The variance of the subway system is larger than that of the microgrid because one subway station out of service makes the entire line break down.

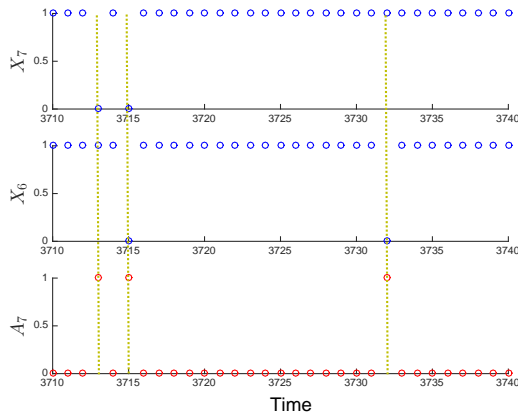


Figure 5: The blue circles at the top two sub-figures show the state transition of node 7 and 6 respectively. The subplot at the bottom shows the corresponding action of node 7.

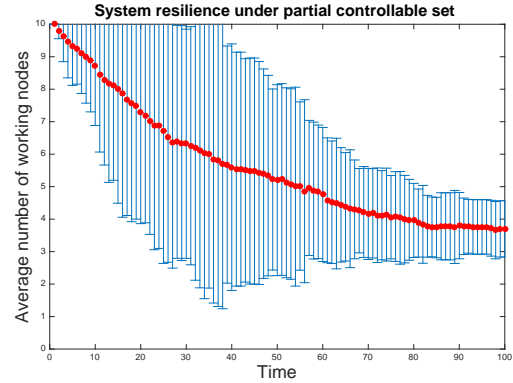


Figure 6: Variations of the average number of working nodes when only node 7 (WTC) and node 8 (Fulton St.) are controllable. The red dots indicate an average number of the working node with error bars shown in blue.

We compare three other policies. Policy 1 is a straightforward strategy, i.e., to repair every faulty node and leave alone working nodes. Policy 2 represents a random strategy that chooses to restore each outage node with probability $4/5$ and repair each non-outage node with probability $1/5$. Taking no action is applied as policy 3. Our simulation results show that policy 1, 2, 3 achieve the value of 979.81, 947.17 and 796.87 with initial state $X = 1$, while our optimal policy achieves a higher value of 985.44. The optimal value obtained by ALP is 989.2963, which is close to the ‘actual value’ obtained from the simulation. Note that it is computationally expensive to directly solve the exact LP. We can approximate the exact value by simulating over a long horizon and averaging over a sufficient large number of samples.

Finally, we study the case when only part of nodes is controllable. In Fig. 6, we maintain the resilience of some nodes with a sparse manageable set. For example, optimal control of nodes 6 and 8 keeps on average four nodes healthily working. More controllable nodes lead to a larger value function as shown in Fig. 6 and the growth is approximately linear with the size of controllable nodes.

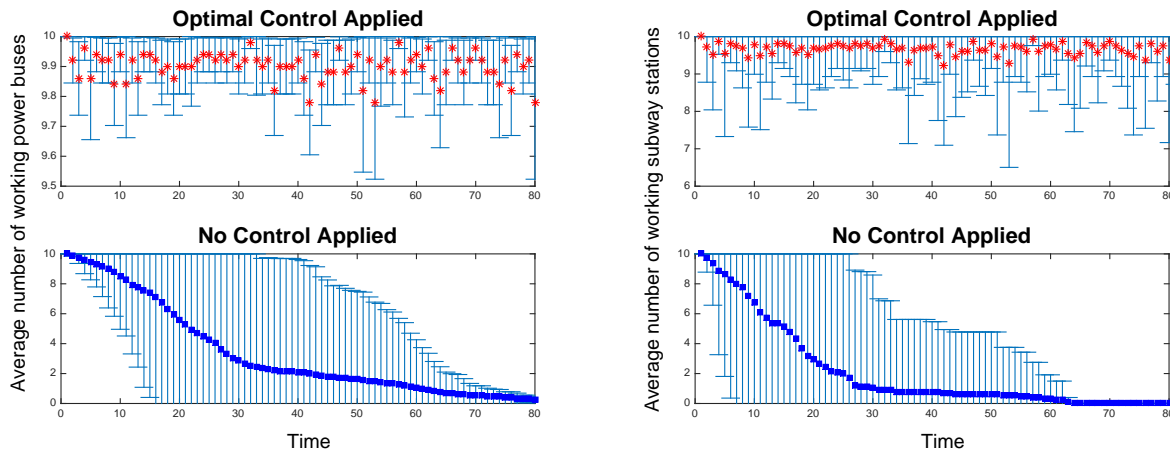


Figure 7: The resilience of the microgrid and subway system with and without the optimal policy. The dot indicates the mean value while the error bar shows the variance of the sample trajectory. Our optimal policy largely reduces the variance and thus makes the system more stable.

4.5 Large-scale Demo

We illustrate the cascading failure of a two-layer network (<https://drive.google.com/drive/u/0/folders/0B6-Q8-SnvO6lYmlxX2FuN3ZuVIU>) with 100 nodes representing mixed components. The failure probability of each node is proportional to the number of its connected faulty nodes and is illustrated via the color bar. A lighter color corresponds to a higher failure probability, and the white color means the node has failed. If a node is repaired or maintained, then it will not fail in a few stages afterward and is shown in black. We introduce both edge-based and node-based resilience metrics, i.e., connectivity and the reward of the working nodes, respectively. They are normalized to 100 with their maximum values. As we can see from the video demo, the system starts with a two-node failure and propagates to the neighboring nodes quickly without effective controls.

On the contrary, the dynamic recovery demo shows how the network can recover quickly and wholly from the catastrophic event where almost all nodes fail. Also, the resilient planning manages to maintain a stable, healthy state of the entire network in the long run after the recovery. We assume that the limited resources can only support one node recovery at each stage to obtain a clear visualization. Our optimal control policy achieves a tradeoff between a high-level resilience and the cost it takes.

5 CONCLUSION

In this paper, we have formulated a factored MDP model for large-scale ICIs and significantly reduced the complexity of computing optimal policies. A distributed optimal control policy is designed to enhance the resilience and security of ICIs. The proposed framework has been applied to a case study motivated by Hurricane Sandy. We have shown that our policy manages to harden the security by reducing the failure probability and achieve resilience by acting on proper nodes. Moreover, the cost achieved by the found policies is the lowest. The framework developed in this work would be useful to mitigate large-scale infrastructure networks through optimal dynamic resilience planning. Even with a small subset of hardened nodes through planning, we have observed a sizable loss reduction for extreme events. We have used microgrids and subway systems in lower Manhattan and downtown Brooklyn as a case study to illustrate the significant resilience improvement in event of Hurricane Sandy. Our framework can be applied to develop resilience policies for other interdependent infrastructures including water distribution systems, food systems, data centers, and communication infrastructures. The future work would aim at understanding

implementation and resource constraints on the policies and developing a large-scale solver that would provide usable infrastructure solutions.

ACKNOWLEDGMENTS

This research is partially supported by a DHS grant through Critical Infrastructure Resilience Institute (CIRI) and grants SES-1541164, and ECCS-1550000 from National Science Foundation (NSF).

REFERENCES

- Chen, J., C. Touati, and Q. Zhu. 2017. “A Dynamic Game Analysis and Design of Infrastructure Network Protection and Recovery”. *ACM SIGMETRICS Performance Evaluation Review* 45(2):125–128.
- Farias, D., and D. Pucci. 2002. *The Linear Programming Approach to Approximate Dynamic Programming: Theory and Application*. Ph. D. thesis, Stanford University.
- Gao, J., D. Li, and S. Havlin. 2014. “From a Single Network to a Network of Networks”. *National Science Review* 1(3):346–356.
- Guestrin, C., D. Koller, R. Parr, and S. Venkataraman. 2003, October. “Efficient Solution Algorithms for Factored MDPs”. *Journal of Artificial Intelligence Research* 19(1):399–468.
- Hayel, Y., and Q. Zhu. 2015. “Resilient and Secure Network Design for Cyber Attack-induced Cascading Link Failures in Critical Infrastructures”. In *Proceedings of the 2015 49th Annual Conference on Information Sciences and Systems (CISS)*, 1–3. IEEE.
- Huang, L., J. Chen, and Q. Zhu. 2017a, April. “A Factored MDP Approach to Optimal Mechanism Design for Resilient Large-Scale Interdependent Critical Infrastructures”. In *Proceedings of the 2017 Workshop on Modeling and Simulation of Cyber-Physical Energy Systems (MSCPES)*, 1–6. IEEE.
- Huang, L., J. Chen, and Q. Zhu. 2017b. “A Large-Scale Markov Game Approach to Dynamic Protection of Interdependent Infrastructure Networks”. In *Proceedings of the International Conference on Decision and Game Theory for Security*, edited by S. e. a. Rass, 357–376. Cham: Springer.
- Kwasinski, A. 2013. “Lessons from Field Damage Assessments about Communication Networks Power Supply and Infrastructure Performance during Natural Disasters with a Focus on Hurricane Sandy”. In *FCC Workshop on Network Resiliency 2013*, Volume 2013.
- Lee II, E. E., J. E. Mitchell, and W. A. Wallace. 2007. “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.
- Obama, B. 2013. “Presidential Policy Directive 21: Critical Infrastructure Security and Resilience”. *Washington, DC*.
- Ouyang, M. 2014. “Review on Modeling and Simulation of Interdependent Critical Infrastructure Systems”. *Reliability engineering & System safety* 121:43–60.
- Rinaldi, S. M., J. P. Peerenboom, and T. K. Kelly. 2001. “Identifying, Understanding, and Analyzing Critical Infrastructure Interdependencies”. *IEEE Control Systems* 21(6):11–25.
- Rosato, V., L. Issacharoff, F. Tiriticco, S. Meloni, S. Porcellinis, and R. Setola. 2008. “Modelling Interdependent Infrastructures Using Interacting Dynamical Models”. *International Journal of Critical Infrastructures* 4(1-2):63–79.
- Wang, X., N. Hovakimyan, and L. Sha. 2013, April. “L1Simplex: Fault-tolerant Control of Cyber-physical Systems”. In *Proceedings of the 2013 ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS)*, 41–50. IEEE.
- Yuan, Y., F. Sun, and Q. Zhu. 2015. “Resilient Control in the Presence of DoS Attack: Switched System Approach”. *International Journal of Control, Automation and Systems* 13(6):1423–1435.
- Zhu, Q., and T. Basar. 2015. “Game-theoretic Methods for Robustness, Security, and Resilience of Cyberphysical Control Systems: Games-in-games Principle for Optimal Cross-layer Resilient Control Systems”. *IEEE Control Systems* 35(1):46–65.

- Zimmerman, R., Q. Zhu, F. de Leon, and Z. Guo. 2017. “Conceptual Modeling Framework to Integrate Resilient and Interdependent Infrastructure in Extreme Weather”. *Journal of Infrastructure Systems* 23(4):04017034.
- Zimmerman, R., Q. Zhu, and C. Dimitri. 2016. “Promoting Resilience for Food, Energy, and Water Interdependencies”. *Journal of Environmental Studies and Sciences* 6(1):50–61.
- Zimmerman, R., Q. Zhu, and C. Dimitri. 2018. “A Network Framework for Dynamic Models of Urban Food, Energy and Water systems (FEWS)”. *Environmental Progress & Sustainable Energy* 37(1):122–131.

AUTHOR BIOGRAPHIES

LINAN HUANG received the B.Eng. degree in Electrical Engineering from Beijing Institute of Technology, China, in 2016. He is currently a second-year Ph.D. candidate in the Laboratory for Agile and Resilient Complex Systems, Tandon School of Engineering, New York University, NY, USA. His research interests include optimal dynamic decision making of the multi-agent system, cyber-physical security, and resilient infrastructure networks. His e-mail address is lh2328@nyu.edu.

JUNTAO CHEN received the B.Eng. degree in Electrical Engineering and Automation with honor from Central South University, Changsha, China, in 2014. He is currently pursuing the Ph.D. degree in the Laboratory for Agile and Resilient Complex Systems, Tandon School of Engineering, New York University, NY, USA. His research interests include game theory, cyber-physical systems, interdependent complex networks, Internet of things, and mechanism design. His e-mail address is jc6412@nyu.edu.

QUANYAN ZHU is an assistant professor in the Department of Electrical and Computer Engineering at New York University. He received the B. Eng. in Honors Electrical Engineering with distinction from McGill University in 2006, the M.A.Sc. from University of Toronto in 2008, and the Ph.D. from the University of Illinois at Urbana Champaign (UIUC) in 2013. From 2013-2014, he was a postdoctoral research associate at the Department of Electrical Engineering, Princeton University. He is a recipient of many awards including NSERC Canada Graduate Scholarship (CGS), Mavis Future Faculty Fellowships, and NSERC Postdoctoral Fellowship (PDF). He spearheaded and chaired INFOCOM Workshop on Communications and Control on Smart Energy Systems (CCSES), Midwest Workshop on Control and Game Theory (WCGT), and 7th Game and Decision Theory for Cyber Security (GameSec). His current research interests include resilient and secure interdependent critical infrastructures, energy systems, cyber-physical systems, and smart cities. His e-mail address is qz494@nyu.edu.