 DOI: 10.1111/risa.12923

Interdependent Network Recovery Games

Andrew M. Smith,[1,*] Andrés D. González,[2,3] Leonardo Dueñas-Osorio,[2] and Raissa M. D'Souza[1,4,5]

Recovery of interdependent infrastructure networks in the presence of catastrophic failure is crucial to the economy and welfare of society. Recently, centralized methods have been developed to address optimal resource allocation in postdisaster recovery scenarios of interdependent infrastructure systems that minimize total cost. In real-world systems, however, multiple independent, possibly noncooperative, utility network controllers are responsible for making recovery decisions, resulting in suboptimal decentralized processes. With the goal of minimizing recovery cost, a best-case decentralized model allows controllers to develop a full recovery plan and negotiate until all parties are satisfied (an equilibrium is reached). Such a model is computationally intensive for planning and negotiating, and time is a crucial resource in postdisaster recovery scenarios. Furthermore, in this work, we prove this best-case decentralized negotiation process could continue indefinitely under certain conditions. Accounting for network controllers' urgency in repairing their system, we propose an *ad hoc* sequential game-theoretic model of interdependent infrastructure network recovery represented as a discrete time noncooperative game between network controllers that is guaranteed to converge to an equilibrium. We further reduce the computation time needed to find a solution by applying a best-response heuristic and prove bounds on ϵ-Nash equilibrium, where ϵ depends on problem inputs. We compare best-case and *ad hoc* models on an empirical interdependent infrastructure network in the presence of simulated earthquakes to demonstrate the extent of the tradeoff between optimality and computational efficiency. Our method provides a foundation for modeling sociotechnical systems in a way that mirrors restoration processes in practice.

KEY WORDS: Game theory; infrastructure recovery; optimization

1. INTRODUCTION AND MOTIVATION

A collection of interdependent networks is crucial to the function of modern society with such networks pervasive in our financial, infrastructure, and cyber systems. Due to society's fundamental reliance on these systems, it is imperative to assess the risk in security from natural or adversarial threats,[1,2] as well as their resilience given the recent interest in restoring community-level normalcy.[3–6] Furthermore, when these systems do fail, plans must be made to restore commodities as quickly as possible to reduce the total cost/damage inflicted on the community and the economy. Yet, analyzing

[1]Department of Computer Science, University of California, Davis, CA, USA.
[2]Department of Civil Engineering & Environmental Engineering, Rice University, Houston, TX, USA.
[3]School of Engineering, Universidad de los Andes, Bogotá, Colombia.
[4]Department of Mechanical and Aerospace Engineering, University of California, Davis, CA, USA.
[5]Santa Fe Institute, Santa Fe, NM, USA.
*Address correspondence to Andrew M. Smith, Department of Computer Science, University of California, Davis, CA 95616, USA; amsmi@ucdavis.edu.

 0272-4332/17/0100-0001$22.00/1 © 2017 Society for Risk Analysis

critical systems in isolation is not sufficient to mitigate security or recoverability risks that exist from dependencies on other systems. Interdependencies have been studied in infrastructure systems from the perspective of modeling,[7,8] security,[9,10] and resilience.[11,12] With respect to recovery of interdependent systems, previous work has focused on the perspective of a centralized network controller or utility operator,[13–15] without explicitly accounting for the decision processes that take place across the distinct network controllers.

Realistically, infrastructure systems, specifically those that exchange multiple interdependent commodities, are managed by several potentially noncooperative entities. Game theory has been used to describe such strategic interactions between agents for control,[16,17] security,[18,19] and organizational partnerships in disaster recovery.[20,21] Approaches to decentralized decision making for resource allocation in the physical recovery process have been suggested as solutions to alleviate shortcomings in restoration processes seen in the field.[22] Applying game-theoretic tools to resource allocation in infrastructure recovery has received relatively little attention, however. Strategic recovery for critical infrastructure is a unique problem in that, unlike optimal control, resilience, or security, decisionmakers are operating in an out-of-nominal environment with limited access to resources and little communication. In developing a model, we must consider key attributes about multiple decisionmakers involved in a recovery scenario:

(1) each network controller can only use resources to recover components in its utility network;
(2) communication between network controllers is limited;
(3) network controllers have no *certain* information about the recovery plans of other network controllers (i.e., they cannot count on a dependency in another network to be restored at a given time);
(4) a large cost is incurred by each controller for leaving demand *in its network* unmet.

The first attribute assumes that recovery resources are commodity specific. In general, there are certain resources that can potentially be shared (money, workers, etc.), but this would require communication and coordination between controllers. The second attribute is a result of the unknown availability of communication resources postdisaster, as well as each controllers' inherent uncertainty about the management hierarchy of other utilities. The most limiting factor to network controllers, outlined in attribute 3, is the lack of information about the structure, utility, and recovery plans of other controllers' networks, specifically in conjunction with limited communication. Attributes 1 and 4 together cause each network controller to be self-interested, and favor strategies that repair their network *quickly* and efficiently. While the optimal reconstruction plan for the whole collection requires considering the entire time horizon of the recovery process, network controllers in postdisaster scenarios instead have to make quick decisions with the resources they have, given the limited information and communication available.

Due to the properties of network controllers (or "players") in a postdisaster reconstruction process, we consider a strategic model that prioritizes minimizing time to find a solution (computation time), that models interactions in a short time period (to account for the *ad hoc* nature of the recovery process), and that contains a realistic model of the limited information available to each player. For our contribution, we propose a decentralization of a recently developed optimal recovery optimization problem for interdependent infrastructure, which provides a utility function for each player involved in each game-theoretic model. We then formally analyze the convergence and tractability properties of a "best-case" information sharing model proposed by Sharkey *et al.*,[23] proving that certain scenarios may cause this process to never converge. Addressing the issues of computation time raised by the information sharing model, we propose and analyze a novel *ad hoc* interdependent network recovery game (INRG) as an imperfect information interdependent network recovery game (sequential INRG), as well as a more time-efficient heuristic best-response solution. Our models are validated using simulated earthquake disasters on a real-world interdependent infrastructure network.

The rest of this article is organized as follows. Section 2 reviews related literature on strategic models of decentralized infrastructure design and recovery. Section 3 introduces the baseline centralized recovery optimization problem and proposed decentralization techniques. Section 3.3 examines solution properties of strategic decentralization with information sharing. Section 3.4 describes and analyzes our novel imperfect information recovery game (sequential INRG), with Section 3.5 providing a heuristic best-response method with compute time relative to the number of players, and proving an upper bound on cost efficiency compared to the exact equilibrium solution. We apply the described decentralized models to a sample recovery scenario in Section 3.6. Finally, Section 4 contrasts the solution quality of strategic decentralization using long-term planning and information sharing with *ad hoc* best-response dynamics on real-world infrastructure data. We find that, while *ad hoc* methods are more indicative of realistic recovery scenarios and are computationally tractable, the best-case information sharing method results in recovery strategies closer to optimal, especially in the presence of extreme failure events. Furthermore, player ordering in the *ad hoc* model has a large effect on cost disparity between players, but a relatively low impact on total recovery cost. In Section 5, we conclude with the implications of our findings and motivate future directions.

2. PREVIOUS WORK

Recently, the effects of interdependence in infrastructure systems on resilience have been studied as reviewed here. Sharkey *et al.* identify, classify, and analyze various types of restoration interdependencies that arose during Hurricane Sandy.[24] Since such interdependencies across multiple infrastructures can create complex scheduling problems in restoration processes, performance and risk metrics have been developed in order to better assess recovery strategies in interdependent infrastructure systems.[11,25] Accounting for the fact that control is usually decentralized in such systems, Reilly *et al.* show how investment decisions made by strategic network controllers are likely to result in underperformance when compared to the performance of a centralized controller, which could have a negative impact in a failure scenario.[16] Game theory has been used on networks to study security in the presence of random failures or probabilistic threats using simultaneous-move games,[10,19] as well as intelligent adversaries with cascading failures using Stackelberg games.[18] Rao *et al.* analyze a cyber-physical security game specific to coupled power and cyber infrastructure, where network properties such as degradation are public information to the attacker.[26]

While understanding how to strategically secure and maintain systems in the presence of failures and attackers is certainly important, recovery after failure is also critical. Centralized optimization of recovery processes in interdependent infrastructure networks have been developed for use in operations research to minimize the flow cost during restoration efforts after a disaster scenario.[27] Matisziw *et al.* formulate a multiobjective restoration optimization problem, noting that while minimizing costs is important in restoration efforts, maximizing flow on the network (and thus satisfying demand) is also significant.[28] Physical restoration on independent infrastructure networks, namely, the class of integrated network design and scheduling (INDS) problems, has been studied, and optimization approaches[29] and computational heuristics[30] have been developed to maximize satisfied demand. Extending these methods to include interdependencies, Cavdaroglu *et al.* develop a mixed-integer linear program (MILP) to minimize the recovery cost for arc reconstruction decisions (as well as several efficient heuristics), accounting for resource availability and dependencies that exist between nodes.[13] González *et al.* generalize the optimization method to include node reconstruction, geographical spaces, varying timescales, and complex interdependencies (including one-to-many and many-to-one dependencies).[14,15] We decentralize this formulation to develop utility functions for each independent player in our model (see Section 3).

Game-theoretic (decentralized) models of recovery in interdependent systems have received relatively little attention. Coles and Zhuang investigate a cooperative model between disaster relief organizations.[20] Since players are explicitly tasked with contributing to relief efforts, cooperation is a reasonable assumption. However, at the level of (potentially profit-driven) network controllers, explicit cooperation cannot be assumed. At the network-controller level, Sharkey *et al.* extend the centralized optimization model in Cavdaroglu *et al.*[13] to be decentralized by partitioning the objective function among multiple players.[23] Each player then solves its objective function (using various assumptions about the initial state of the other players' variables), and shares its recovery strategy with the other players. Given this new information the players revisit their decisions and may revise their solutions and then share their revised recovery strategy, which can prompt other agents to then revise their solution, etc. Convergence and efficiency properties are analyzed using a computational study. In our work, we formally analyze the convergence properties of this model, and use it as a "best-case" recovery cost decentralization baseline (if convergence is reached). We develop an alternative game model that prioritizes minimizing computation time and guarantees convergence. To our knowledge, there has been no other work in applying game theory to recovery at the network-controller level, or with regard to reducing computation time while proving error approximations.

3. DECENTRALIZED MODELS OF INTERDEPENDENT NETWORK RECOVERY

Centralized recovery models serve to discover best-case resource allocation solutions. Recently, several models of recovery in interdependent infrastructure systems have been developed.[13–15] However, in real-world systems, resource allocation decisions and ownership are decentralized; therefore, we must develop tractable formulations that are also decentralized. For our decentralization methods and empirical studies, we will extend the time-dependent interdependent network design problem (td-INDP) described in González *et al.*[15] It is important to note that the techniques discussed in this article would also apply to other centralized recovery optimization problems (including that from Cavdaroglu *et al.*[13]).

3.1. Centralized td-INDP Model

In the td-INDP model, the objective function minimizes the recovery cost for all discrete timesteps $t \in \{0 \ldots |T|\}$ in a time horizon T for an interdependent infrastructure network $\mathcal{G} = (\mathcal{N}, \mathcal{A})$, where $\mathcal{N}$ is the set of nodes and $\mathcal{A}$ is the set of directed arcs in $\mathcal{G}$. Table I lists and describes all necessary notation to define the td-INDP. For our model, a set of parameters and decision variables exist for each timestep $t \in T$; if t is not referenced as a subscript in the notation, we assume a constant value across the entire time horizon. We represent the nodes and arcs that have failed or have been destroyed by $\mathcal{G}' = (\mathcal{N}', \mathcal{A}')$. $\mathcal{G}$ may be separated into several subnetworks or layers to represent different utility systems and their commodities (i.e., power and water subnetworks), such that $\mathcal{G} = \mathcal{G}_1 \cup \cdots \cup \mathcal{G}_{|\mathcal{K}|}$, where $\mathcal{K}$ is the set of layers, and $\mathcal{G}_k = (\mathcal{N}_k, \mathcal{A}_k)$, $\forall k \in \mathcal{K}$. Interdependencies between layers are modeled as arcs between select nodes in different layers, and are represented by the variables $\gamma_{ijk\hat{k}}$, where $\gamma_{ijk\hat{k}} > 0$ if a node $j \in \mathcal{N}_{\hat{k}}$ is dependent on a node $i \in \mathcal{N}_k$ and $k \neq \hat{k}$. Each node in the infrastructure network has a specified supply (demand, if negative), b_{ik}, $i \in \mathcal{N}_k$. The infrastructure network can also be divided into geographical spaces, $\mathcal{S}$, where $\alpha_{iks} = 1$ if node $i \in \mathcal{N}_k$ is in space $s \in S$, and $\beta_{ijks} = 1$ if arc $(i, j) \in \mathcal{A}_k$ is in space $s \in S$.

The costs required as input for the td-INDP optimization problem can be divided into operational costs and recovery costs. Operational costs include unit flow costs, c_{ijk}, for arc $(i, j) \in \mathcal{A}_k$, and supply penalties, M^+_{ik} for being oversupplied, and M^-_{ik} for being undersupplied, for each node $i \in \mathcal{N}_k$. Corresponding surplus and deficit supply variables are represented by δ^+_{ikt} and δ^-_{ikt}. Recovery costs are included to represent the cost of repairing nodes, q_{ik}, $\forall i \in \mathcal{N}_k$, and the cost of repairing arcs, f_{ijk}, $\forall (i, j) \in \mathcal{A}_k$. Costs for preparation of a geographical space, g_s, $\forall s \in \mathcal{S}$ can also be included to address a common problem of co-location and shared conveyance of utility systems. The preparation of a geographical space, s, denoted by z_{st}, only has to happen once for all nodes and arcs in the space that are set to be recovered simultaneously at time t.

The availability of a recovery resource is represented by v_r, where $r \in \mathcal{R}$ and $\mathcal{R}$ is the set of available resources. We assume that $|\mathcal{R}| = |\mathcal{K}|$ (i.e., each layer has its own independent resources), and $v_r = 1$, $\forall r \in \mathcal{R}$ for each timestep. Further, we assume the amount of resources required to repair each node, $\nu_{ikr} = \nu_{ik}$, $\forall i \in \mathcal{N}_k$, and each arc, $h_{ijkr} = h_{i,j,k}$, $\forall (i, j) \in \mathcal{A}_k$ is equal to 1. Given all of the above inputs, the td-INDP optimization problem finds the optimal flow to assign to each arc, and the nodes and arcs to recover that minimizes the total cost. Flow on an arc $(i, j) \in \mathcal{A}_k$ is designated by x_{ijkt}. A node $i \in \mathcal{N}'_k$ is set to be recovered at time t if the binary variable $\widetilde{w}_{ikt} = 1$. Similarly, an arc $(i, j) \in \mathcal{A}'_k$ is set to be recovered if $\widetilde{y}_{ijkt} = 1$. Our reduced MILP formulation of the td-INDP model is shown in Equations (1a)–(1s). Equation (1a) defines the objective function, which minimizes the total cost of the recovery process over all timesteps in T, with cost types outlined in boxes. From left to right and top to bottom, these cost types are geographical space reconstruction cost, arc reconstruction cost, node reconstruction cost, under- and oversupply penalties, and flow cost. Constraints (1b) limit flow at each node, allowing for excess and deficit flow. Constraints (1c)–(1e) disallow flow on nodes, arcs, or arcs connected to nodes that are not functioning due to direct failure or a failed interdependency. Constraint (1f) allows the use of $|\mathcal{K}|$ resources per layer. Note that this does allow the sharing of resources among layers as well. Constraint (1g) disallows nodes from functioning if

Table I. Table of Notation for td-INDP and Game-Theoretic Models

Sets	
$\mathcal{N}$	Set of nodes in the entire network
$\mathcal{A}$	Set of arcs in entire network
$\mathcal{S}$	Set of geographical spaces
$\mathcal{K}$	Set of infrastructure networks (layers)
$\mathcal{R}$	Set of unique resources for reconstruction efforts
$\mathcal{N}_k$	Set of nodes in layer $k \in \mathcal{K}$
$\mathcal{A}_k$	Set of arcs in layer $k \in \mathcal{K}$
$\mathcal{N}'_k$	Set of failed nodes in layer $k \in \mathcal{K}$
$\mathcal{A}'_k$	Set of failed arcs in layer $k \in \mathcal{K}$
T	Set of discrete timesteps in the network recovery time horizon

td-INDP decision variables	
δ^+_{ikt}	Excess supply at node i in layer k at time t
δ^-_{ikt}	Unmet demand at node i in layer k at time t
x_{ijkt}	Flow through arc (i, j) in layer k at time t
w_{ikt}	Binary variable that indicates if node i in layer k is functional at time t
y_{ijkt}	Binary variable that indicates if arc (i, j) in layer k is functional at time t
$\tilde{w}_{ikt}$	Binary variable that indicates if node i in layer k is repaired at time t
$\tilde{y}_{ijkt}$	Binary variable that indicates if arc (i, j) in layer k is repaired at time t
z_{st}	Binary variable that indicates if space s has to be prepared at time t

td-INDP parameters	
v_r	Availability of resource r at time t
h_{ijkrt}	Usage of resource r in recovering arc (i, j) in layer k at time t
v_{ikrt}	Usage of resource r in recovering node i in layer k at time t
M^+_{ikt}	Cost of excess supply at node i in layer k at time t
M^-_{ikt}	Cost of unmet demand at node i in layer k at time t
β_{ijkst}	Indicates if repairing arc (i, j) in layer k at time t requires preparing space s
α_{ikst}	Indicates if repairing node i in layer k at time t requires preparing space s
$\gamma_{ijk\tilde{k}t}$	Indicates node j in layer $\tilde{k}$ is dependent on node i in layer k at time t
g_{st}	Cost of preparing space s at time t
f_{ijkt}	Cost of recovering arc (i, j) in layer k at time t
q_{ikt}	Cost of recovering node i in layer k at time t
c_{ijkt}	Unit flow cost through arc (i, j) in layer k at time t
u_{ijkt}	Flow capacity of arc (i, j) in layer k at time t
b_{ikt}	Demand (negative)/supply (positive) of node i in layer k at time t

Game-theoretic notation	
$\mathcal{P}$	Set of players (network controllers) in strategic model
$\mathcal{Q}$	Set of actions available to players (usually $\mathcal{Q} = (N' \cup A')$)
$\mathcal{Q}_p$	Set of actions available to player $p \in \mathcal{P}$ (usually $\mathcal{Q}_p = (N'_p \cup A'_p)$)
$\mathcal{H}$	Set of choice nodes in a sequential game, where each element is mapped to a player $p \in \mathcal{P}$ and a set of possible actions from $\mathcal{Q}_p$
$\mathcal{Z}$	Set of terminal nodes in a sequential game that assigns a payoff value for each $p \in \mathcal{P}$
$\rho_\tau(n)$	Utility value of repairing node $n \in \mathcal{N}$ at time $\tau \in T$, assuming all dependencies of n are satisfied
$\alpha_\tau(n)$	Utility value of repairing node $n \in \mathcal{N}$ at time $\tau \in T$, given the actual functionality of the dependencies of n

$$\text{td-INDP} = \text{minimize} \sum_{t\in T|t>0}\left(\boxed{\sum_{s\in\mathcal{S}} g_{st}\widetilde{z}_{st}} + \sum_{k\in\mathcal{K}}\left(\boxed{\sum_{(i,j)\in\mathcal{A}'_k} f_{ijkt}\widetilde{y}_{ijkt}} + \boxed{\sum_{i\in\mathcal{N}'_k} q_{ikt}\widetilde{w}_{ikt}}\right)\right)$$

$$\sum_{t\in T}\left(\boxed{\sum_{i\in\mathcal{N}_k}\left(M^+_{ikt}\delta^+_{ikt} + M^-_{ikt}\delta^-_{ikt}\right)} + \boxed{\sum_{(i,j)\in\mathcal{A}_k} c_{ijkt}x_{ijkt}}\right) \quad (1a)$$

subject to

$$\sum_{j:(i,j)\in\mathcal{A}_k} x_{ijt} - \sum_{j:(j,i)\in\mathcal{A}_k} x_{jit} = b_{it} - \delta^-_{it} + \delta^+_{it}, \quad \forall k\in\mathcal{K}, \forall i\in\mathcal{N}_k, \forall t\in T, \quad (1b)$$

$$x_{ijkt} \le u_{ijkt}w_{ikt}, \quad \forall k\in\mathcal{K}, \forall(i,j)\in\mathcal{A}_k, \forall t\in T, \quad (1c)$$

$$x_{ijkt} \le u_{ijkt}w_{jkt}, \quad \forall k\in\mathcal{K}, \forall(i,j)\in\mathcal{A}_k, \forall t\in T, \quad (1d)$$

$$x_{ikjt} \le u_{ijkt}y_{ijkt}, \quad \forall k\in\mathcal{K}, \forall(i,j)\in\mathcal{A}'_k, \forall t\in T, \quad (1e)$$

$$\sum_{k\in\mathcal{K}}\left(\sum_{i\in\mathcal{N}'_k} v_{ikt}\widetilde{w}_{ikt} + \sum_{(i,j)\in\mathcal{A}'_k} h_{ijkt}\widetilde{y}_{ijkt}\right) \le |\mathcal{K}| \quad \forall t\in T \mid t>0, \quad (1f)$$

$$\sum_{j\in\mathcal{N}_{\hat{k}}} w_{j\hat{k}t}\gamma_{jik\hat{k}t} \ge w_{ikt}, \quad \forall k,\hat{k}\in\mathcal{K}, \forall i\in\mathcal{N}_k, \forall t\in T, \quad (1g)$$

$$w_{ikt} \le \sum_{\tilde{t}=1}^{t} \widetilde{w}_{ik\tilde{t}}, \quad \forall k\in\mathcal{K}, \forall i\in\mathcal{N}'_k, \forall t\in T|t>0, \quad (1h)$$

$$y_{ijkt} \le \sum_{\tilde{t}=1}^{t} \widetilde{y}_{ijk\tilde{t}}, \quad \forall k\in\mathcal{K}, \forall(i,j)\in\mathcal{A}'_k, \forall t\in T|t>0, \quad (1i)$$

$$\widetilde{w}_{ikt}\alpha_{ikst} \le z_{st}, \quad \forall k\in\mathcal{K}, \forall i\in\mathcal{N}_k, \forall s\in\mathcal{S}_k, \forall t\in T, \quad (1j)$$

$$\widetilde{y}_{ijkt}\beta_{ikst} \le z_{st}, \quad \forall k\in\mathcal{K}, \forall(i,j)\in\mathcal{A}_k, \forall s\in\mathcal{S}_k, \forall t\in T, \quad (1k)$$

$$\delta^+_{ikt} \ge 0, \quad \forall k\in\mathcal{K}, \forall i\in\mathcal{N}_k, \forall t\in T, \quad (1l)$$

$$\delta^-_{ikt} \ge 0, \quad \forall k\in\mathcal{K}, \forall i\in\mathcal{N}_k, \forall t\in T, \quad (1m)$$

$$x_{ijkt} \ge 0, \quad \forall k\in\mathcal{K}, \forall(i,j)\in\mathcal{A}_k, \forall t\in T, \quad (1n)$$

$$w_{ikt} \in \{0,1\}, \quad \forall k\in\mathcal{K}, \forall i\in\mathcal{N}_k, \forall t\in T, \quad (1o)$$

$$y_{ijkt} \in \{0,1\}, \quad \forall k\in\mathcal{K}, \forall(i,j)\in\mathcal{A}_k, \forall t\in T, \quad (1p)$$

$$\widetilde{w}_{ikt} \in \{0,1\}, \quad \forall k\in\mathcal{K}, \forall i\in\mathcal{N}'_k, \forall t\in T, \quad (1q)$$

$$\widetilde{y}_{ijkt} \in \{0,1\}, \quad \forall k\in\mathcal{K}, \forall(i,j)\in\mathcal{A}'_k, \forall t\in T, \quad (1r)$$

$$z_{st} \in \{0,1\}, \quad \forall s\in\mathcal{S}, \forall t\in T. \quad (1s)$$

their dependencies are not functioning. Constraints (1h) and (1i) enforce that a node (or arc) has to be repaired in a previous timestep in order to be operational. Constraint (1k) requires a geographical space to be prepared once for any recovery actions taken in that space. Finally, constraints (1l)–(1s) define the nonnegativity for continuous variables, and the binary nature for integer variables.

3.2. Decentralized INDP Model

The centralized model assumes there is one controller for the entire interdependent infrastructure system. We wish to decentralize the optimization problem to allow for multiple controllers or decisionmakers (which we will call players). First, we assume each player uniquely controls one layer of the network. Formally, we define the set of players in the decentralized system as $\mathcal{P}$, where $|\mathcal{P}| = |\mathcal{K}|$. The nodes, $\mathcal{N}$, are partitioned into $|\mathcal{P}|$ sets, such that if node $n \in \mathcal{N}_k$, then $n \in \mathcal{N}_p$, where $p \in \mathcal{P}$, $\forall n$. The resulting decentralized td-INDP (denoted td-DINDP) objective function for a player $p \in \mathcal{P}$ is shown in Equation (2). Constraints (1b)–(1s) can be adapted to this formulation by replacing the set $\mathcal{K}$ (and corresponding indices, k and $\hat{k}$) with $\mathcal{P}$ (and corresponding index, p, to indicate the optimizing player, and p' to index over other players). We also constrain each player to only have one resource (i.e., we change the right-hand side of constraint (1f) to 1 instead of $|\mathcal{K}|$).

$$\text{td-DINDP}(p) = \text{minimize}$$

$$\sum_{t \in T | t>0} \left(\sum_{s \in \mathcal{S}} g_{st} z_{st} + \sum_{(i,j) \in \mathcal{A}'_p} f_{ijpt} \tilde{y}_{ijpt} + \sum_{i \in \mathcal{N}'_p} q_{ipt} \tilde{w}_{ipt} \right)$$
$$+ \sum_{t \in T} \left(\sum_{i \in \mathcal{N}_p} \left(M^+_{ipt} \delta^+_{ipt} + M^-_{ipt} \delta^-_{ipt} \right) + \sum_{(i,j) \in \mathcal{A}_p} c_{ijpt} x_{ijpt} \right) \tag{2}$$

To put Equation (2) in terms of a payoff function for each player, we define the payoff, $\pi_{p,t}(\tilde{\mathcal{G}}, \tilde{\mathcal{S}})$, for a player p at time $t \in T$ for recovering some set of nodes, $\tilde{\mathcal{N}}_p \subseteq \mathcal{N}'_p$, arcs, $\tilde{\mathcal{A}}_p \subseteq \mathcal{A}'_p$ (where $\tilde{\mathcal{G}}_p = (\tilde{\mathcal{N}}_p \cup \tilde{\mathcal{A}}_p)$), and preparing geographical spaces $\tilde{\mathcal{S}} \subseteq \mathcal{S}$ in Equation (3). The minimization term refers to a minimum cost flow problem with respect to constraints (1b)–(1e) and (1g) (i.e., we are only optimizing for flow cost and not recovery actions), where $w_{ipt} = 1$ if node $i \in \{\mathcal{N}_p \setminus \mathcal{N}'_p\} \cup \tilde{\mathcal{N}}$ and $w_{ipt} = 0$ otherwise.

$$\pi_{p,t}(\tilde{\mathcal{G}}, \tilde{\mathcal{S}}) = \sum_{s \in \tilde{\mathcal{S}}} g_{st} + \sum_{(i,j) \in \tilde{\mathcal{A}}_p} f_{ijpt} + \sum_{i \in \tilde{\mathcal{N}}_p} q_{ipt}$$
$$+ \min \left(\sum_{i \in \mathcal{N}_p} \left(M^+_{ipt} \delta^+_{ipt} + M^-_{ipt} \delta^-_{ipt} \right) + \sum_{(i,j) \in \mathcal{A}_p} c_{ijpt} x_{ijpt} \right) \tag{3}$$

It is also possible in the centralized model for all resources to be allocated to one layer for certain timesteps in order to repair critical dependencies present in other layers. In a distributed environment, however, players are unaware of these dependencies in other networks. This unawareness could leave a player waiting for the availability of a dependency, when, in fact, this dependency may never be repaired. This is further exacerbated if all players are waiting on cyclical dependencies to be repaired. In practice, utility network controllers choose the best repair option, given the limited information available. To prevent such a deadlock in our optimization problem, we force players to take an action in every recovery stage, even if it may not be immediately beneficial. This is a practical constraint since, eventually, the network will be (mostly) fully restored. Equation (4) shows the corresponding constraints.

$$\sum_{i \in \mathcal{N}'_p} \tilde{w}_{it} + \sum_{(i,j) \in \mathcal{A}'_p} \tilde{y}_{ijt} = \min(1, |\mathcal{N}'_p| + |\mathcal{A}'_p|), \quad \forall t \in T \mid t > 0. \tag{4}$$

3.3. Information Sharing in INRGs

In Sharkey *et al.*,[23] a method of using information sharing is used to explore decentralization in the recovery of interdependent infrastructure systems. The information sharing process begins with each agent independently forming a plan to recover her system for every discrete timestep during the restoration, which includes acknowledgment of dependencies on resources controlled by other agents. Initially, each player either assumes the other players have repaired all required nodes (optimistic) or that they have not repaired any (pessimistic). After each agent has formed its initial plan, all the players then exchange plans and the planning process is repeated given the new information.

This process of planning, sharing information, and reformulating the restoration plan iterates until the restoration plans of all players converge (i.e., after exchanging information, no agent wishes to change its restoration plan). Given convergence for all the participating agents and assuming the agents are self-interested and not externally incentivized, the information sharing process represents a best-case decentralization technique where agents have unlimited foresight and negotiation time.

While information sharing is a good baseline mechanism, it suffers from two main issues. First, forming a recovery plan for the entire time domain is not practical. Oftentimes, recovery decisions are made *ad hoc*, either by fixing components that are most convenient to resource location, or the most beneficial to the network at the current time.[31] The second issue involves the potentially large amount of time required for exchanging plans and negotiating. During this process, repairs cannot be made until a recovery strategy is agreed upon. Below, we prove sufficient conditions that this negotiation process may never converge.

3.3.1. Nonconvergence in Information Sharing

We now provide sufficient conditions in any INRG using information sharing for nonconvergence. To assist in the proof, we define the *actual value*, $\alpha_t(n)$, of a node n at time t as the utility node n actually gives to a player when repaired at time t, considering whether or not the dependencies are functional.

The *potential value*, $\rho_t(n)$, of a node n at time t is the utility a node n could potentially give a player when repaired at time t, assuming all of its dependencies are functional. We define $\mathcal{N}_x$ to be the nodes in player p_x's network, and $\mathcal{N}'_x \subseteq \mathcal{N}_x$ to be the destroyed nodes in player p_x's network. Note that $\rho_\tau(n) \geq \alpha_\tau(n)$ always holds.

Theorem 1. *If at some recovery time, τ, the following conditions hold:*

- *a player p_x controls nodes i and i' such that $\rho_\tau(i') > \rho_\tau(i) > \rho_\tau(n) > 0, \forall n \in \mathcal{N}'_x \setminus \{i'\}$, where i' is dependent on some node $j \in N'_y$ in some other player p_y's network;*
- *a player p_y controls nodes j and j' such that $\rho_\tau(j') > \rho_\tau(j) > \rho_\tau(n) > 0, \forall n \in \mathcal{N}'_y \setminus \{j'\}$, where j' is dependent on some node $i \in \mathcal{N}'_x$;*
- *$i' \neq i$, $j' \neq j$, $\alpha_\tau(i) = \rho_\tau(i)$, $\alpha_\tau(j) = \rho_\tau(j)$; and*
- *players p_x and p_y are either both initially pessimistic or both optimistic,*

then, information sharing will fail to converge.

Proof. Assume two players p_1 and p_2 control nodes $i, i' \in \mathcal{N}'_1$ and $j, j' \in \mathcal{N}'_2$, at some time τ. Node i' is dependent on node j and node j' is dependent on i and $\rho_\tau(i') > \rho_\tau(i)$ for player p_1 and $\rho_\tau(j') > \rho_\tau(j)$ for player p_2. At time τ, $\rho_\tau(i') > \rho_\tau(i) > \rho_\tau(n) > 0, \forall n \in \mathcal{N}'_1 \setminus \{i'\}$ and $\rho_\tau(j') > \rho_\tau(j) > \rho_\tau(m) > 0, \forall m \in \mathcal{N}'_2 \setminus \{j'\}$.

- *If players are initially optimistic*: For the initial planning stage, player p_1 would assume node j is repaired at time τ, and thus believes $\alpha_\tau(i') = \rho_\tau(i')$. Because of the assumption, player p_1 would choose to repair node i' at time τ. Similarly, player p_2 assumes node i is repaired at time τ, and thus believes $\alpha_\tau(j') = \rho_\tau(j')$, and would choose to repair node j'.
  In the second stage of planning, player p_1 now sees that player p_2 recovered node j' at time τ (instead of the prior belief of j). As a result, $\alpha_\tau(i') = 0$, since node i' is dependent on node j, and node j has not yet been repaired at time τ. Player p_1 will now revise its plan and choose to recover i, since $\alpha_\tau(i) = \rho_\tau(i)$ and $\alpha_\tau(i) > \alpha_\tau(n)$, $\forall n \in \mathcal{N}'_1$. Symmetrically, player p_2 observes that player p_1 recovered node i' at time τ. By a similar proof used for player p_1, we can show that player p_2 will choose to recover j.
  After stage 3 of planning, we now have identical strategies for both players at time τ as in the initial planning stage. This cycle will repeat indefinitely, and the information sharing process will fail to converge.
- *If players are initially pessimistic*: For the initial planning stage, player p_1 would assume node j is not repaired at time τ. So $\alpha_\tau(i') = 0$, and player p_1 will choose to recover some node i, since $\alpha_\tau(i) = \rho_\tau(i)$ and $\alpha_\tau(i) > \alpha_\tau(n)$, $\forall n \in \mathcal{N}'_1$. A symmetrical situation exists for player p_2. This strategy is identical to that of the second stage of planning in the optimistic case. Since the evolution of the strategy from this planning stage forward is only dependent on the current strategy, the same cycle described in the optimistic case is reached. Furthermore, using the same argument, if this strategy point is ever reached in any planning stage, the information sharing process will fail to converge. □

Note that Theorem 1 does not discount other more complex cycles from occurring, but shows that the simple two-edge interdepencency cycle and symmetric preferential ordering are sufficient (but not necessary) conditions for nonconvergence.

3.4. Imperfect Information INRGs

To address the need for a decentralized model that conforms more with the *ad hoc* nature of repair actions in infrastructure recovery, and that can be efficiently bounded, we propose a set of INRG. For INRGs, we use the concept of a sequential game (as opposed to a simultaneous, or normal-form game), which enables asynchronous moves and observation of other players' actions. We argue that sequential games are better suited for critical recovery scenarios in the general case because actions by players are not necessarily taken in lockstep with other players in the game. Also, if interdependencies exist, it may be possible to observe the recovery action of another player. For instance, if a water pump is dependent on power, the water network utility controller can observe if there is power coming to the pump or not. To incorporate *ad hoc* strategies in our sequential INRG model, each player will compute its utility function using td-DINDP restricted to only one timestep ($|T| = 1$). The entire recovery time domain can be represented by incrementally applying the repairs resulting from the strategy played by each player for a single timestep until the network is fully recovered.

Imperfect information games are a subset of sequential games that allow constraints on how much history (or players' past actions) can be observed. We use a sequential model to prescribe an ordering of player actions. The sequence of players' recovery actions can be shown in a game tree, where each choice node, $\eta \in \mathcal{H}$, in the game tree represents a point where a player chooses an action to take. Each choice node, η, has a unique player, $p \in \mathcal{P}$, associated with it, and a set of possible actions from $\mathcal{Q} = (\mathcal{N}' \cup \mathcal{A}')$ that p can take from choice node η. The choice node, η, will have $|\mathcal{Q}|$ successor nodes, which represent a choice node $\hat{\eta} \in \mathcal{H}$ for another player $p' \in \mathcal{P}$ (where $p' \neq p$), or a terminal node, $z \in \mathcal{Z}$. A terminal node, z, assigns a payoff to all players for the path of actions from the root of the game tree to z. In our model, the game tree will always have a height $|\mathcal{P}|$, since each sequential INRG models one timestep in the recovery process, and we assume all players will perform an action. However, there are $|\mathcal{P}|!$ possible player orderings, or game tree configurations, for any timestep in a recovery process.

Information sets in sequential games are used to model a player's lack of information. If two choice nodes belong to the same information set of player $p \in \mathcal{P}$, then p cannot tell the two choice nodes apart (i.e., player p has the same knowledge at both choice nodes). We use interdependencies to form information sets. From the perspective of a player p, any recovery action that another player p' makes on a node that is not a dependency to a failed node in $\mathcal{N}_p'$ belongs to the same information set. In other words, players only gain information if they observe an interdependency being repaired.

The most popular solution concept for sequential games is the *subgame perfect equilibrium* (SPE).[32] SPE requires there to be a Nash equilibrium at every subgame, or game rooted at every choice node, in the game tree. We calculate SPE by use of backward induction; in other words, finding the Nash equilibrium of each subgame from the bottom of the tree, and propagating the equilibrium strategies of each subgame up to the root. For games of imperfect information, the *sequential equilibrium* is used, which requires a distribution of (reasonable) beliefs over choice nodes within information sets. However, the utility of choice nodes in information sets preceded by recovery actions at nodes without interdependencies is equivalent. As such, the best-response action would be identical, enabling the easier to compute SPE concept to be used sensibly for sequential INRGs. Even though there is a discrepancy in utility with respect to player ordering, SPE is guaranteed to exist for sequential INRGs, since this is a finite game where no player forgets their previous move (perfect recall).[32]

3.4.1. Uncertainty over Player Ordering

In general, we cannot assume any prescribed player ordering in infrastructure recovery scenarios. Even if there were a sense of one network being "dominant" over another, this may not be known to the players. In this case, players have incomplete information about what game they are actually playing rather than imperfect information regarding where they are in a game's tree. We can account for this uncertainty by introducing a new player: Nature. Nature determines, with some probability known to all players, what the player ordering will be, and players take this into account when calculating their utilities. Realistically, scenario-specific events could govern this probability. For instance, a gas utility controller may need to contain a natural gas leak before electricity can be restored (thus, the gas utility controller would be required to move first). In our empirical experiments, we model Nature by selecting a player ordering uniformly at random for each timestep in the recovery process.

3.5. Best Response in INRGs

In the previous sequential game models, we make the assumption that each player's utility function (namely, the network structure, failed components, supply/demand penalties, etc.) is known by all other players. However, in the worst case, one network controller knows nothing about the input to another network controller's utility function (aside from dependencies that may have been repaired in previous moves). To model this lack of knowledge, we can assume that players simply best respond to actions that have already been made, or calculate their utility based solely on the information they have up to their choice node, and make no assumptions about future moves.

Unfortunately, best-response dynamics in INRGs may not result in an SPE. Using Condition 1, we will prove this below in Theorem 2.

Condition 1. *Consider a sequential (imperfect) INRG for timestep τ between a leading player $p_1 \in \mathcal{P}$ and a following player $p_2 \in \mathcal{P}$. Using best-response dynamics, player p_1 is able to make a recovery action, and player p_2 can observe this action and best respond. Player p_1 controls failed nodes $i, i' \in \mathcal{N}_1'$. Player p_2 controls a failed node j, $j' \in \mathcal{N}_2'$, and i is dependent on j. Player p_2 has a dominant strategy to repair node j (player p_2 will always recover node j regardless of player p_1's decisions). Initially, for player p_1, we assume $\rho_\tau(i) > \rho_\tau(i') > \rho_\tau(n), \forall n \in \mathcal{N}_1'$, and, if node j is not functional, $\alpha_\tau(i') > \alpha_\tau(i)$ and $\alpha_\tau(i') > \alpha_\tau(n), \forall n \in \mathcal{N}_1'$ (i.e, i' is player p_1's best choice to repair if node j is not functional).*

Theorem 2. *An SPE is not guaranteed to be found using best-response dynamics in imperfect information INRGs.*

Proof. Since we do not assume the leading player p_1 has any knowledge of the following player p_2's utility function, p_1 will select to recover node i'. As a result, player p_2 will best respond by recovering node j. If we now imagine that players can examine each other's strategies, player p_1 would wish to deviate and repair node i; thus, a solution found using best-response dynamics is not guaranteed to be an SPE. □

We can, however, bound the quality of the solution obtained from best-response dynamics relative to an SPE. We use the notion of *regret* for a player, or the difference in a player's utility between the played strategy in the best-response solution and the optimal strategy, given all other players' strategies, they would wish to unilaterally deviate to. A loose upper bound on regret, neglecting network structure, can be adopted using supply and demand penalties. For any scenario, the highest possible utility node, $\chi \in \mathcal{N}_p'$, that could be repaired for player p is the lowest cost node that restores supply/demand to all nodes in $\mathcal{N}_p \setminus \mathcal{N}_p'$. The lowest possible utility node, $\psi \in \mathcal{N}_p'$, is the node that, when repaired, satisfies the least amount of demand to only itself, with the highest cost. If χ is in the SPE solution and ψ is in the best-response solution for player p at time τ, we have an absolute worst-case bound on regret for time τ:

$$\epsilon_\tau' = \max_{p\in\mathcal{P}} \left(\max_{\chi\in\mathcal{N}_p'} \left(\sum_{n}^{\mathcal{N}_p\cup\{\chi\}} \left(M_n^+ \delta_{np\tau}^+ + M_n^- \delta_{np\tau}^- \right) - q_{\chi p\tau} \right) - \min_{\psi\in\mathcal{N}_p'} \left(\delta_{\psi p\tau}^+ M_{\psi p\tau}^+ + \delta_{\psi p\tau}^- M_{\psi p\tau}^- - q_{\psi p\tau} \right) \right).$$

While general and easy to compute, this bound may be too loose to gain much insight into the quality of an equilibrium strategy for a particular scenario. In Theorem 3, we use nodes' potential and actual values (which can be calculated via flow optimization in polynomial time) to prove tighter bounds on regret that consider network and interdependency structure.

Theorem 3. *The maximum regret for any player in a sequential INRG using best-response dynamics is at most $\epsilon_\tau = \max_{p\in\mathcal{P}} \left(\rho_\tau(i' \in \mathcal{N}_p') - \alpha_\tau(i \in \mathcal{N}_p')\right)$, given that node i' is the optimally recovered node (given all other players' best-response strategies) and node i is recovered via best response.*

Proof. Assume that a player p chooses to fix a node $i \in \mathcal{N}_{p_1}'$ via best response at time τ. The regret player p experiences depends purely on the strategies of other players:

- If there *does not exist* a node $j \in \mathcal{N}_{p'}', \forall p' \in \mathcal{P}$, where $p' \neq p$, such that some node $x \in \mathcal{N}_p'$ is dependent on j that is chosen to be recovered in the best-response solution, then $\rho_\tau(x) = \alpha_\tau(x), \forall x \in \mathcal{N}'_p$, and $\alpha_\tau(i) \geq \alpha_\tau(x), \forall x \in \mathcal{N}'_p$ (otherwise, player p would not have chosen to recover i in best response). Thus, the *regret of player p*, $\epsilon_{\tau,p} = 0$.

- If such a node j *does exist*, $i' \in \mathcal{N}'_p$ is dependent on j (and $i \neq i'$), then two possibilities exist:
  - If $\alpha_\tau(i) \geq \rho_\tau(i')$, then there is no utility to be gained by deviating; thus, $\epsilon_{\tau,p} = 0$.
  - If $\alpha_\tau(i) < \rho_\tau(i')$, then the utility gained by deviating to recover node i' is $\epsilon_{\tau,p} = \rho_\tau(i') - \alpha_\tau(i)$.

Generally, in a $|\mathcal{P}|$ player game, there could be $|\mathcal{P}| - 1$ regret values (note, the last player in the ordering must not want to deviate, since all relevant information is available to it). Therefore, the maximum regret, $\epsilon_\tau = \max_{p \in \mathcal{P}}(\epsilon_{\tau,p})$. □

With Theorem 3, we can now say that each solution provided by best-response dynamics in INRGs is at most ϵ_τ away from an SPE, or, more generally an ϵ-Nash equilibrium (where $\epsilon = \epsilon_\tau$). Note that ϵ_τ could be quite large if there is a disparity in the utility between those components that can be repaired versus those that cannot be repaired due to failed dependencies. In Section 4, we analyze the empirical regret values as they compound over the entire recovery process.

3.5.1. Computing Best Response

While best-response dynamics calculates an inexact ϵ-Nash equilibrium, it may be a desirable solution due to its computational efficiency relative to computing an SPE. Calculating an exact SPE requires computing td-DINDP (for one timestep) at each choice game node for each player via backward induction. Since the depth of each tree is $|\mathcal{P}|$, and at every choice node each player has at most $|Q|$ choices, this results in $|Q|^{|\mathcal{P}|}$ computations (note that the recovery decisions have already been made at the bottom of the game tree, so these reduce to pure linear programs for optimizing flow in the repaired subnetwork). Using best-response dynamics at any given timestep, td-DINDP is computed once for each player, which results in only $|\mathcal{P}|$ computations for each timestep.

3.6. Applying Decentralized Models to an Example Scenario

We will use an example scenario (shown in Fig. 1) to illustrate each of the decentralized models discussed in Section 3. In this scenario, there are two players, $p_1, p_2 \in \mathcal{P}$. Player p_1 controls nodes $\{1, 2, 3, 4\} \in \mathcal{N}_1$ (and all arcs between them), where nodes $\{1, 2, 4\} \in \mathcal{N}'_1$ have failed. Player p_2 controls nodes $\{5, 6, 7, 8\} \in \mathcal{N}_2$ (and all arcs between them), where nodes $\{6, 7, 8\} \in \mathcal{N}'_2$ have failed. Interdependencies are drawn between layers. We assume undersupply and oversupply penalties per unit flow to be equal ($M^+_{ip} = M^-_{ip} = \Delta, \ \forall i \in \mathcal{N}_p, \ \forall p \in \mathcal{P}$). Supply and demand, $b_i, \forall i \in \mathcal{N}$, is shown in Table II, where we set $|b_i| = q_i, \forall i \in \mathcal{N}$ with the assumption that components with higher capacity/demand are more costly to repair.

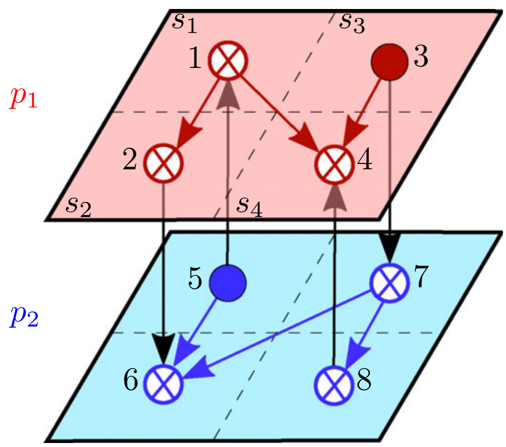

Fig. 1. An example two-player eight-node interdependent network recovery scenario. An "X" over a node indicates a failed node; arcs drawn between nodes in different layers indicate interdependence (i.e., node 6 is dependent on node 2, and thus node 6 cannot function until node 2 has been repaired). Supply (demand, if negative) for each node, b_n, for this scenario is shown in Table II, which we will assume is proportional to the reconstruction cost. Table III shows the approximate payoff matrix, assuming Δ, the undersupply and oversupply penalty, is sufficiently large. Table IV represents four planning iterations of the information sharing process, where each row represents the corresponding player's recovery sequence for each planning stage. After stage 2, a cycle forms, and thus the planning process will fail to converge.

Table II. Supply/Demand for Each Node in the Example Scenario in Fig. 1

p_1	p_2
$b_1 = 3$	$b_5 = 4$
$b_2 = -1$	$b_6 = -6$
$b_3 = 4$	$b_7 = 3$
$b_4 = -6$	$b_8 = -1$

3.6.1. *Equilibrium Analysis of Sample Scenario*

To illustrate the use of td-DINDP in a game setting, we first look at a normal-form recovery game for the first timestep in the sample scenario (i.e., $|T| = 1$). A normal-form game corresponds to a setting where both players make their recovery action simultaneously. The traditional solution of such a game, known as the *pure strategy* Nash equilibrium, is a choice of action for each player such that no player would benefit from changing her strategy (given the other player's action choices remain constant). The payoff matrix for every possible recovery action can be seen in Table III, with player p_1 as the column player (payoff shown as the first entry in each cell), and player p_2 as the row player (payoff shown as the second entry in each cell). Utilities shown are actually recovery *costs*, so each player aims to *minimize* this value. Since we assume Δ dominates the cost, each player's utility for each recovery action is computed using decentralized td-INDP, and written in terms of Δ (except when reconstruction cost is required to break a tie in equilibrium calculation).

For instance, consider if p_2 chooses to recover node 6 and p_1 chooses to recover node 2 in timestep 1 (payoff shown in row 1, column 2 of Table III). Assuming that oversupply and undersupply penalties dominate the cost (using node repair costs as a "tiebreaker" for equilibrium computation), player p_2's recovery cost at timestep 1, using Equation (3), would evaluate to (without considering geographical constraints):

$$\pi_{2,1}(\{6\}) = q_{6,2,1} + \sum_{i \in \mathcal{N}_2} \left(M^+_{i,2,1} \Delta^+_{i,2,1} + M^-_{i,216} \Delta^-_{i,2,1} \right)$$
$$= |b_6| + (2\Delta + 3\Delta + 1\Delta) = 6\Delta + |b_6|.$$

Table III. Payoff Matrix for First Timestep of the Example Scenario in Fig. 1 (Nash Equilibria in Bold)

	1	2	4
6	$(14\Delta + \|b_1\|, 14\Delta)$	$(\mathbf{14\Delta + \|b_2\|, 6\Delta})$	$(14\Delta + \|b_4\|, 14\Delta + \|b_6\|)$
7	$(14\Delta, 14\Delta)$	$(14\Delta, 14\Delta)$	$(14\Delta, 14\Delta + \|b_7\|)$
8	$(14\Delta, 14\Delta)$	$(14\Delta, 14\Delta)$	$(\mathbf{6\Delta, 14\Delta + \|b_8\|})$

Note: Costs are shown with respect to the dominant over/undersupply terms (Δ), and reconstruction cost terms are added if needed for equilibrium calculation.

Similarly, player p_1's payoff would evaluate to:

$$\pi_{1,1}(\{2\}) = q_{2,1,1} + \sum_{i \in \mathcal{N}_1} \left(M^+_{i,1,1} \Delta^+_{i,1,1} + M^-_{i,1,1} \Delta^-_{i,1,1} \right)$$
$$= |b_2| + (3\Delta + 1\Delta + 4\Delta + 6\Delta) = 14\Delta + |b_2|.$$

The other cells in Table III are calculated similarly. To better understand how these values are calculated, we further describe the scenario with the given recovery strategy. Player p_2 would continue to pay the total of 3Δ for oversupply at node 7 and 1Δ for undersupply at node 8 (since these nodes are still not functioning), but only 2Δ for undersupply at node 6, since player p_1 fixed node 6's dependency at node 2, which allows four units of node 6's demand to be met by node 5; thus, player p_2's cost for this action is 6Δ. Player p_1 continues to pay all oversupply and undersupply costs, since, even though node 2 is recovered, the only possible supply node is still nonfunctioning; thus player p_1's cost for this actions is 14Δ. To distinguish this action from the other choices player p_1 could have made given player p_2's choice, we display the node reconstruction term $|b_2|$. In this case, recovering node 2 is player p_1's best option, given player p_2's recovery of node 6, since node 2 is the cheapest node to repair. Recovering node 6 is also player p_2's best option, given player p_1's recovery of node 2; this pure strategy is a Nash equilibrium. There are actually two pure strategy Nash equilibria in this example (shown as bold cells in Table III). Strategy (4, 8) is much better for player p_1, and strategy (2, 6) is better for player p_2. We will refer back to these equilibria in analyses of other game-theoretic models.

3.6.2. *Applying Information Sharing to Example Scenario*

Using information sharing, each player first calculates decentralized td-INDP using some initial prior belief about the other player's recovery strategy (we will assume each player is initially optimistic for this example). The time horizon $|T| = 3$, since it will take three timesteps to fully recover the network. After solving td-INDP, each player now has a plan of recovery for each timestep in the time horizon. The players then look at each others' plans, and update their beliefs accordingly. Table IV shows the solution using information sharing for the example scenario. This shows that a cycle forms in the plans, starting at Plan 1, indicating that neither player agrees on a plan, and the sharing process continues indefinitely. Note that, for this scenario, other equivalent cost solutions

Table IV. Information Sharing Solution with Nonconvergent Planning Cycle for the Example Scenario in Fig. 1

	Plan 1	Plan 2	Plan 3	Plan 4	...
p_1	(4, 1, 2)	(2, 1, 4)	(4, 1, 2)	(2, 1, 4)	
p_2	(6, 7, 8)	(7, 8, 6)	(6, 7, 8)	(7, 8, 6)	

Note: Each column represents a planning stage, and each cell represents a sequence of recovery for the corresponding player.

exist (such as player p_1 recovering nodes in the order (1, 4, 2) in Plan 3) that result in more complex cycles.

3.6.3. Imperfect Information Equilibrium Analysis of Sample Scenario

To better illustrate the concept of sequential INRGs, Fig. 2 shows both possible game tree configurations for the example recovery scenario in Fig. 1. In the top game tree, the player ordering is $\{p_1, p_2\}$. At the root choice node, p_1 can choose to recover node 1, 2, or 4. After this choice is made, p_2 can observe p_1's action (if an interdependency exists between the node player p_1 recovered and any of player p_2's failed nodes), and choose a recovery action. After each of p_2's choices, there is a set of terminal nodes with utilities that correspond to their respective recovery action path. For instance, the far-left path would indicate p_1 repairing node 1, and p_2 repairing node 6, resulting in a cost of 14Δ for p_1 and 14Δ for p_2. The dotted line joining p_2's choice node after the edge labeled 1 and the choice node after the edge labeled 4 indicates that these belong to the same information set. This is because p_2 has no dependencies on node 1 or node 4, and thus cannot observe these actions. Both of these choice nodes result in the same payoff for p_2. This will always be true, since nodes that have no dependencies also have no effect on a player's utility function. We will use this fact to compute equilibria.

For the example scenario recovery, SPEs are shown in Fig. 2 in bold. Following the only path from the root to a terminal node, the utility of these equilibria can be calculated. In the case where $\mathcal{P}_1$ moves first, the utility is calculated from p_1 recovering node 4 and p_2 recovering node 8, or the action vector $(4, 8)$ (a Nash equilibria in Table III that favors p_1). When p_2 leads, the action vector is $(6, 2)$ (also, a Nash equilibria in Table III that favors p_2). While this example is contrived to illustrate the complexities of interdependencies in recovery games, we see similar advantages of leading in our empirical studies.

3.6.4. Best-Response Analysis of Sample Scenario

We will use the sample scenario in Fig. 1 to illustrate how the difference in information between the SPE and best-response solution concepts affects

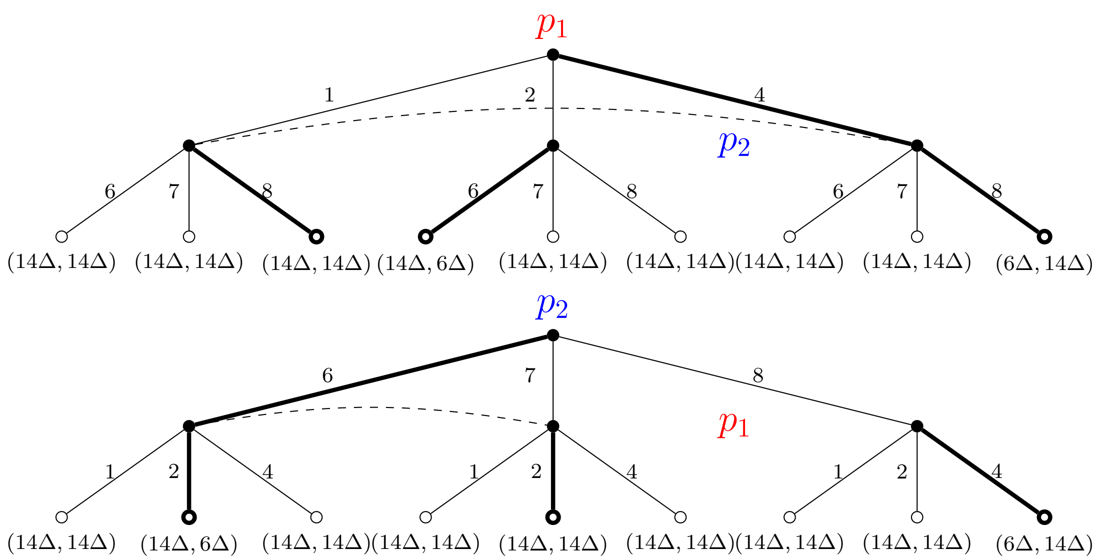

Fig. 2. An extensive form game tree of the example game illustrated in Fig. 1. Actions beside each edge indicate which node is chosen to be repaired at the parent choice game node. Payoffs at leaf nodes indicate the costs of playing the preceding actions to player p_1 and p_2, respectively, Top: player p_1 is leading; Bottom: player p_2 is leading.

players' strategies and payoffs. For instance, the SPE in Fig. 2 (top), p_1 chooses to recover node 4 in equilibrium because they know p_2 will repair node 8. For Fig. 2, best-response dynamics would have resulted in the action sequences $(2, 6)$ (top game tree) and $(8, 4)$ (bottom game tree). While these sequences are not SPE, both correspond to pure strategy Nash equilibria in the normal-form game, favoring the nonleading player instead.

4. EMPIRICAL ANALYSIS OF INRGs

We now investigate applications of INRGs to recovery scenarios on a real-world infrastructure network. In the following sections, we compare and contrast the performance of the centralized td-INDP, decentralization with information sharing (InfoShare), sequential INRG solved using backward induction (INRG-BI), and sequential INRG using best response (INRG-BR), with various player orderings.

4.1. Shelby County Infrastructure Data

For our empirical experiments, we use the power and water network data gathered from Shelby County, TN. Interdependencies between each network were determined in Hernandez *et al.*[25,33] A summary of each network, along with the geographical layout, is shown in Fig. 3. For our failure scenarios, an earthquake simulation technique is used[34,35] to generate failed components in each network. We use 1,000 sampled scenarios for earthquake moment magnitudes ranging from $M_w = 6$ to $M_w = 9$ (unless otherwise noted).

4.2. Nonconvergence in Information Sharing Recovery Games

In Section 3.3, we prove there exist situations in which information sharing may never converge. A natural question to ask then, is how often do these (or other, more complex) nonconvergence properties occur in realistic recovery scenarios? For our empirical study of information sharing, we perform a similar analysis, as in Sharkey *et al.*,[23] on the Shelby County data. Our analysis is done over many more disaster scenarios so that we may investigate patterns of nonconvergence and compare with the other decentralized network recovery methods presented. Fig. 4 shows the percentage of 1,000 simulated earthquakes, for each magnitude, where convergence was never reached after six iterations of information sharing. As shown, nonconvergence becomes exponentially more common in larger magnitude earthquakes (i.e., more damaged components and more decisions to make increase the probability of dependency cycles). A naive approach to resolving a recovery strategy from these nonconvergent samples is to simply threshold the number of rounds of information sharing in which each player can participate, and select the recovery strategy for each

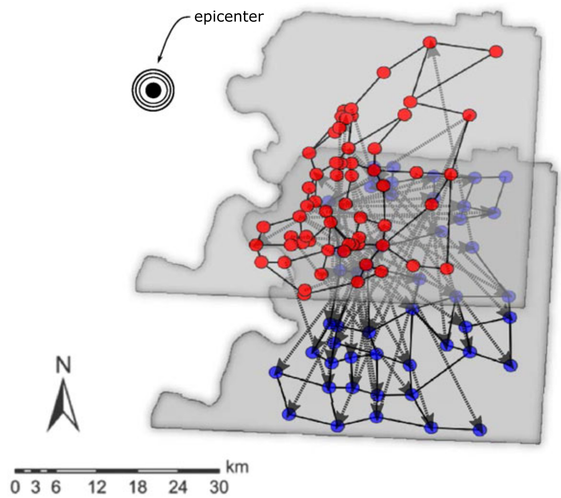

	Power	Water
# of nodes	60	49
# of arcs	152	142
# of dependencies	9	45

Fig. 3. Shelby County infrastructure network, with power network (top, red), water network (bottom, blue), and interdependencies.

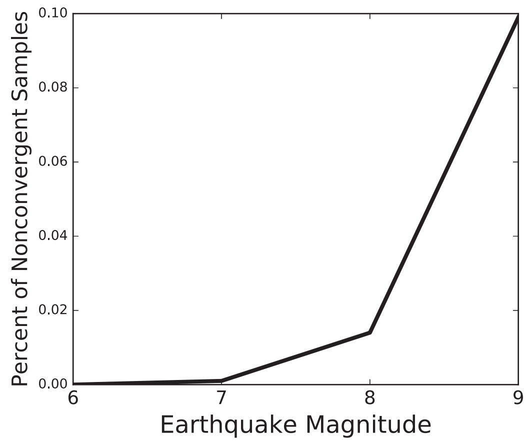

Fig. 4. Number of samples that remain nonconvergent after six iterations of information sharing, for each earthquake magnitude.

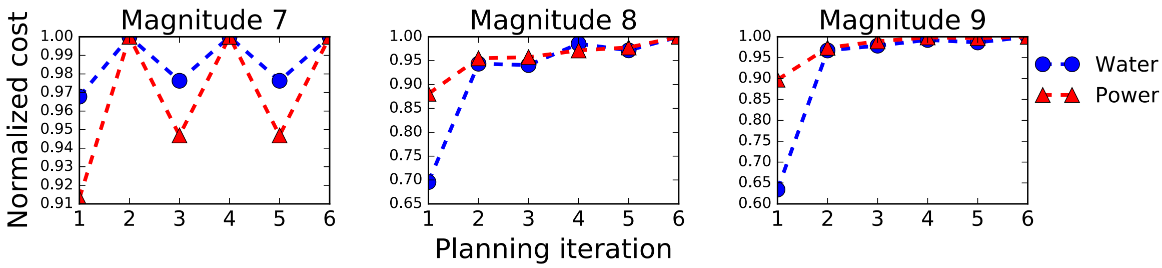

Fig. 5. Total cost of recovery versus iterations of planning in InfoShare for each player.

player that corresponds to the last round. However, to set this threshold, and to see if the threshold is reasonable, we must examine the evolution of cost throughout the information sharing process.

Fig. 5 shows the fluctuation of costs for the nonconvergent samples. Each player's cost is normalized by the highest cost observed for all samples. For the magnitude 7 case, we only observed one nonconvergent sample out of 1,000. Here, we observe the exact behavior described by the proof in Section 3.3.1—there is a cycle in recovery action ordering throughout the information sharing iterations, and this is reflected in the cost. The risk of selecting a suboptimal cutoff for the information sharing process is up to 5% for the power network controller, and 2% for the water network controller.

As the magnitude increases, strategic behavior is no longer explained simply by the scenario described in the proof. In fact, we see gradually increasing cost for both players as the planning process proceeds. This is likely as a result of more complex interdependencies between failed nodes with more severe destruction in the network. The more destructive scenarios motivate a much lower cutoff for the planning process, but it may be difficult to determine the effectiveness of such a cutoff in real time.

4.3. Efficiency of Best Response in Sequential Recovery Games

An alternative solution to information sharing is to treat each timestep in the recovery process as a separate, sequential game. A Nash equilibrium can be reached by solving each sequential game using backward induction. However, Section 3.5 provides a more computationally efficient and realistic heuristic to solve such games. Fig. 6 illustrates the actual recovery cost efficiency of the best-response solution compared to Nash equilibrium over the entire recovery time domain. An exponential improvement in the

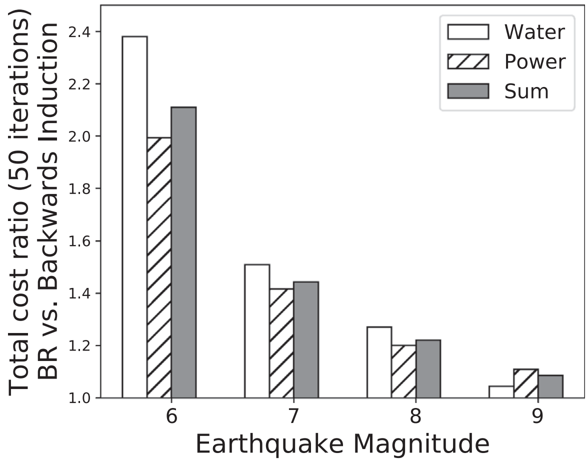

Fig. 6. Total cost ratio between best-response (BR) and exact computation of SPE using backward induction (over 100 samples), with random player ordering.

total cost ratio is seen as the disaster applied to each network increases in magnitude. Recall that using backward induction requires increasing computation time as the disaster becomes more severe, since there are more components to recover. This creates an interesting decision point for infrastructure network controllers, weighing the benefits of recovery cost versus computation time. Specifically, moderate severity disasters (i.e., magnitude 6 earthquakes) may allow time to compute exact solutions (in addition to the time it would take to exchange information between various players, not accounted for in our data) in exchange for a much better savings in recovery cost, whereas with catastrophic disasters (i.e., magnitude 8+ earthquakes), the computation time is not likely to be worth the recovery cost savings.

Since backward induction requires each player to know the utilities of the other players (which requires some form of exchange of information between players beforehand), we will use the best-response solution for INRGs in the following results.

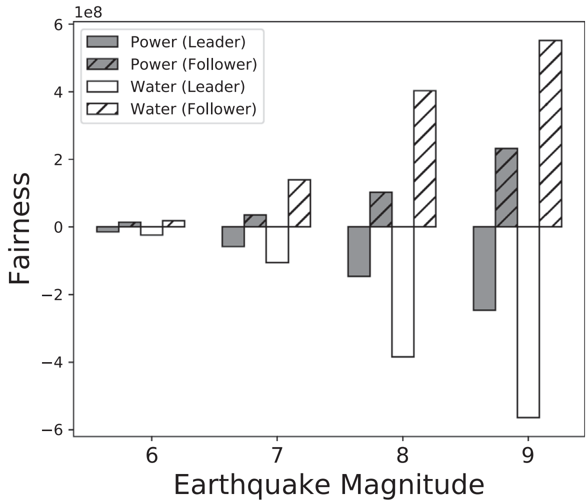

Fig. 7. Fairness in player ordering for INRG. Positive fairness value for a player indicates a beneficial ordering to that player compared to random.

4.4. Fairness in Sequential Recovery Games

As discussed in Section 3.4, player ordering in sequential INRG can have a significant impact on the cost incurred by each player in equilibrium. In the example scenario, the utility difference between the two Nash equilibria that emerged with different player orderings was relatively significant. In order to measure the magnitude of this concept in our empirical results, we define the metric *fairness* to be how much better (or worse) off a player is using a fixed player ordering over a random player ordering. Fig. 7 shows the fairness values of INRG for various magnitudes of earthquake, where positive values indicate that the player is better off with the respective ordering than with a random ordering. As noted in Section 3.5 and reflected in the example scenario, the following player in a fixed ordering is better off. In fact, as the earthquake magnitude increases, the disparity in costs relative to random ordering between the players increases significantly. This can be attributed to the consequences of suboptimal recovery decisions persisting for a longer duration, since there are more damaged components to repair, thus making the recovery process longer.

4.5. Comparison of Algorithms

Decentralized methods are often measured in comparison to their optimal, centralized solutions. For our centralized method, we use td-INDP as our benchmark, which optimizes over the entire time horizon. We set the number of resources, $v_r = 2$, so it is equal to the total number of resources available to each player in the decentralized case. The time horizon is set to be sufficiently large, so the network is fully recovered in one solve of the MILP. For sequential INRG (or simply INRG), we iteratively use best-response dynamics until the network is fully restored. For information sharing (InfoShare), each player is initially optimistic, and solves an approximate td-INDP (for computational purposes) with its current belief of the other player's recovery strategy as input. Since InfoShare is not guaranteed to converge, we restrict the number of planning iterations to six, and record the utilities of each player at the final stage of planning. For both decentralized methods, we set $v_r = 1$ for each player.

Fig. 8 (top) shows the cost breakdown for INRG with random player ordering over the recovery time horizon per player. Over time, INRG costs the players increasingly more than the centralized solution. This effect appears magnified as the earthquake increases in magnitude. Interestingly, in all magnitudes, there is a point in the recovery process where the player controlling the power network in the INRG solution is paying more than the cost of recovering the entire network in the td-INDP solution. This result is likely due to timesteps where it would have been optimal for the water network controller to share its recovery resource with the power network controller. In the extreme case of a magnitude 9 earthquake, at timestep 20, the cost of recovering both players' individual networks in INRG is more costly than the entire centralized solution.

Fig. 8 (bottom) shows similar cost per iteration results for InfoShare. While the differences between InfoShare and td-INDP are comparable to those of INRG, we can notice that the difference in the integrals between the cost functions of td-INDP and InfoShare is slightly smaller than that of td-INDP and INRG. To better examine this difference, we plot a comparison of the total cost of recovery over the entire time domain for each algorithm in Fig. 9. This shows the recovery cost gap between InfoShare and INRG grows as the disaster scenario becomes more severe. Recall, however, that InfoShare may take several rounds of plan sharing (possibly an infinite number of rounds in certain scenarios) to achieve these improved recovery costs. Similar to the tradeoff between optimality and computational efficiency of SPE versus best response, decisionmakers may have to decide how many rounds of planning

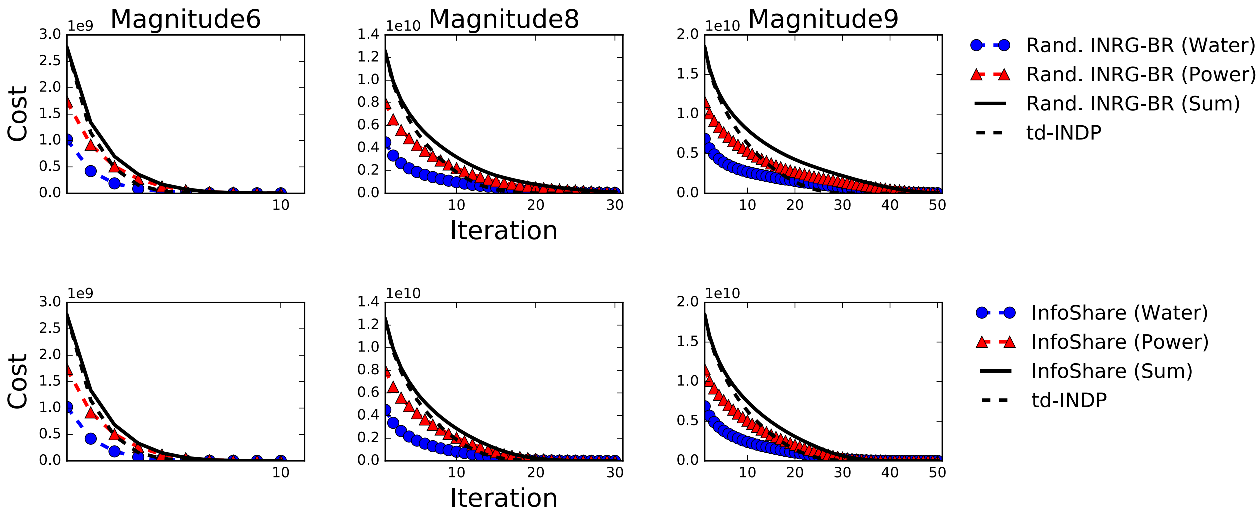

Fig. 8. Average cost comparison per iteration between INRG (top) and InfoShare algorithms, relative to the optimal td-INDP solution, for earthquake magnitude $M_w = 6$, 8, and 9.

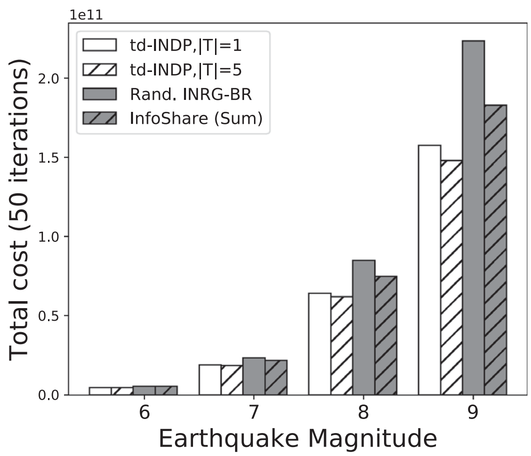

Fig. 9. Total cost comparison between centralized and decentralized recovery methods.

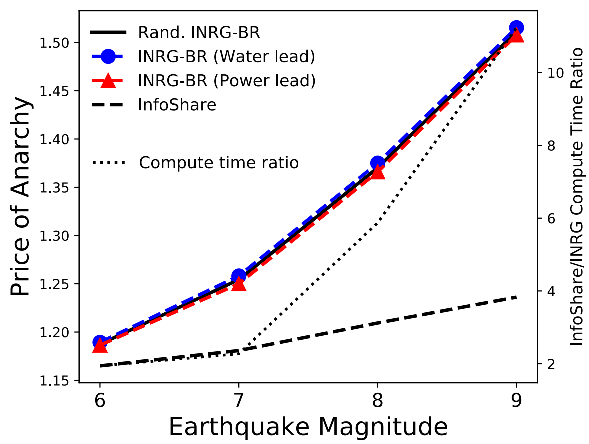

Fig. 10. Average price of anarchy for decentralized recovery methods.

they can afford (if the disaster context allows for it) in order to still be better off than when simply best responding.

4.5.1. Price of Anarchy (PoA) and the Efficiency of Decentralization

For a direct metric of how well each decentralization method does relative to the centralized solution, we use the PoA. Typically, PoA in game theory is used to calculate the efficiency of the (worst Nash equilibrium) decentralized solution compared to the centralized, optimal solution.[(36)] In our experiments, we do not necessarily have an equilibrium (in the case of InfoShare), and we are averaging over 1,000 samples. Nevertheless, this ratio is enlightening in understanding the expected inefficiency of decentralization in a variety of recovery scenarios. We slightly redefine PoA for our model in Equation (5).

$$\text{PoA} = \frac{\text{Avg. decentralized cost}}{\text{Avg. td-INDP cost}} \tag{5}$$

We calculate PoA for INRG and InfoShare, and plot it in Fig. 10. Also shown is the average compute time ratio between InfoShare and INRG. In this comparison, InfoShare is significantly more efficient in terms of recovery cost, at the cost of much more computation and coordination between players. Furthermore, as earthquake magnitude increases, InfoShare increases almost linearly over the data points provided, while INRG has exponential inefficiency growth. The required computation for InfoShare (at a cutoff of 7 plan sharing iterations), however, increases exponentially on a larger scale. Regarding player ordering in INRGs, we compare fairness from Fig. 7 to PoA in Fig. 10 and see that there is very little change in the PoA for any player ordering, despite the disparity between players.

To summarize, our results demonstrate the differences between two strategic network recovery models: InfoShare, which prioritizes optimizing recovery costs in lieu of communication efficiency, and our proposed noncooperative sequential INRG model, which performs recovery actions in an *ad hoc* manner when using a best-response heuristic. Specifically, sequential INRG, while underperforming in physical recovery costs when compared with Info Share, provides an always convergent solution that quickly restores components in a destroyed network. This contrast between computational efficiency and recovery cost efficiency provides network controllers with a decision point about which model to use, given the severity of the disaster or availability of communication resources.

5. CONCLUSION

In this article, we develop a game-theoretic model of resource allocation in interdependent infrastructure recovery (sequential INRGs), considering the limited available information to network controllers, and the *ad hoc* nature in which the recovery process realistically takes place. Furthermore, we adapt a best-response solution technique for solving sequential INRGs, and bound the worst-case recovery cost difference from an exact equilibrium. We compare INRGs to a best-case decentralization method (InfoShare) that considers the entire recovery time horizon and complete visibility of other players' plans, which we prove to be unbounded in convergence time in the worst case. Empirical results show that there is a tradeoff between minimal recovery cost and computation time, specifically in large magnitude disasters, which should be taken into account by decisionmakers in postdisaster infrastructure recovery scenarios. If the urgency to repair components in a network is relatively high, best responding is a more efficient and tractable method than finding communication channels with other players and negotiating. Our work provides a basis for modeling resilience of communities, especially as restoration processes are inherently decentralized, yet coordinated; such a balance cannot be captured with prevailing centralized optimization.

The INRG framework considers that players have no foresight, communication opportunities (aside from the visibility of dependencies in other networks being repaired), and no knowledge of other players' network structure or utility function. In this respect, INRGs are worst-case scenarios; in reality, players may be partially informed about the structure of others' networks from past interactions/observations, census data, natural resource locations, etc. Further sensitivity analysis to network and interdependency structure could also provide valuable insight into expected upper and lower bounds of efficiency in decentralized models. For future work, we leave the extension of this analysis to include imperfect information (via Bayesian games) and develop incentives that guide players closer to an optimal solution. We also wish to apply our models to more diverse network topologies.

ACKNOWLEDGMENTS

We gratefully acknowledge support from the U.S. Army Research Laboratory and the U.S. Army Research Office under MURI award number W911NF-13-1-0340. We would also like to thank Burkhard C. Schipper in the University of California, Davis, Department of Economics for helpful discussions on exploring various related game-theoretic models.

REFERENCES

1. Haimes YY, Crowther K, Horowitz BM. Homeland security preparedness: Balancing protection with resilience in emergent systems. Systems Engineering, 2008; 11(4):287–308.
2. Santos JR, Haimes YY. Modeling the demand reduction input-output (I-O) inoperability due to terrorism of interconnected infrastructures. Risk Analysis, 2004; 24(6):1437–1451.
3. MacKenzie CA, Zobel CW. Allocating resources to enhance resilience, with application to Superstorm Sandy and an electric utility. Risk Analysis, 2015; 36(4):847–862.
4. Park J, Seager TP, Rao PSC, Convertino M, Linkov I. Integrating risk and resilience approaches to catastrophe management in engineering systems. Risk Analysis, 2013; 33(3):356–367.
5. Gilbert SW. Disaster Resilience: A Guide to the Literature. Gaithersburg, MD: National Institute of Standards and Technology, 2010.
6. NIST. Community Resilience Planning Guide for Buildings and Infrastructure Systems, Vol. I and II. Gaithersburg, MD: National Institute of Standards and Technology, 2015.

7. Dueñas-Osorio L, Craig JI, Goodno BJ. Seismic response of critical interdependent networks. Earthquake Engineering & Structural Dynamics, 2007; 36(2):285–306.
8. Cavallaro M, Asprone D, Latora V, Manfredi G, Nicosia V. Assessment of urban ecosystem resilience through hybrid social–physical complex networks. Computer-Aided Civil and Infrastructure Engineering, 2014; 29(8):608–625.
9. Kunreuther H, Heal G. Interdependent security. Journal of Risk and Uncertainty, 2003; 26(2-3):231–249.
10. Heal G, Kunreuther H. Modeling interdependent risks. Risk Analysis, 2007; 27(3):621–634.
11. Baroud H, Barker K, Ramirez-Marquez JE, Rocco CM. Inherent costs and interdependent impacts of infrastructure network resilience. Risk Analysis, 2015; 35(4):642–662.
12. Santos JR, Herrera LC, Yu KDS, Pagsuyoin SAT, Tan RR. State of the art in risk analysis of workforce criticality influencing disaster preparedness for interdependent systems. Risk Analysis, 2014; 34(6):1056–1068.
13. Cavdaroglu B, Hammel E, Mitchell JE, Sharkey TC, Wallace WA. Integrating restoration and scheduling decisions for disrupted interdependent infrastructure systems. Annals of Operations Research, 2013; 203(1):279–294.
14. González AD, Dueñas-Osorio L, Sánchez-Silva M, Medaglia AL. The interdependent network design problem for optimal infrastructure system restoration. Computer-Aided Civil and Infrastructure Engineering, 2016; 31(5):334–350.
15. González AD, Dueñas-Osorio L, Sánchez-Silva M, Medaglia AL. The time-dependent interdependent network design problem (TD-INDP) and the evaluation of multi-system recovery strategies in polynomial time. Pp. 544–550 in The 6th Asian-Pacific Symposium on Structural Reliability and its Applications. Shanghai, China: Tongji University Press, 2016.
16. Reilly AC, Samuel A, Guikema SD. "Gaming the System": Decision making by interdependent critical infrastructure. Decision Analysis, 2015; 12(4):155–172.
17. Zhang P, Peeta S, Friesz T. Dynamic game theoretic model of multi-layer infrastructure networks. Networks and Spatial Economics, 2005; 5(2):147–178.
18. Lou J, Smith AM, Vorobeychik Y. Multidefender security games. IEEE Intelligent Systems, 2017; 32(1):50–60.
19. Heal G, Kearns M, Kleindorfer P, Kunreuther H. Interdependent security in interconnected networks. Pp. 258–275 in Seeds of Disaster, Roots of Response: How Private Action Can Reduce Public Vulnerability. New York, NY: Cambridge University Press, 2006.
20. Coles J, Zhuang J. Decisions in disaster recovery operations: A game theoretic perspective on organization cooperation. Journal of Homeland Security and Emergency Management, 2011; 8(1).
21. Guan P, Zhuang J. Modeling public–private partnerships in disaster management via centralized and decentralized models. Decision Analysis, 2015; 14(4):173–189.
22. Mendonça D, Wallace WA. Impacts of the 2001 World Trade Center attack on New York City critical infrastructures. Journal of Infrastructure Systems, 2006; 12(4):260–270.
23. Sharkey TC, Cavdaroglu B, Nguyen H, Holman J, Mitchell JE, Wallace WA. Interdependent network restoration: On the value of information-sharing. European Journal of Operational Research, 2015; 244(1):309–321.
24. Sharkey, TC, Nurre, SG, Nguyen H, Chow, JH, Mitchell, JE, Wallace WA. Identification and classification of restoration interdependencies in the wake of Hurricane Sandy. Journal of Infrastructure Systems, 2015; 22(1).
25. Hernandez-Fajardo I, Duenas-Osorio L. Probabilistic study of cascading failures in complex interdependent lifeline systems. Reliability Engineering & System Safety, 2013; 111:260–272.
26. Rao NS, Poole SW, Ma CY, He F, Zhuang J, Yau DK. Defense of cyber infrastructures against cyber-physical attacks using game-theoretic models. Risk Analysis, 2015; 36(4):694–710.
27. Lee II EE, Mitchell JE, Wallace WA. Restoration of services in interdependent infrastructure systems: A network flows approach. IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews), 2007; 37(6):1303–1317.
28. Matisziw TC, Murray AT, Grubesic TH. Strategic network restoration. Networks and Spatial Economics, 2010; 10(3):345–361.
29. Nurre SG, Cavdaroglu B, Mitchell JE, Sharkey TC, Wallace WA. Restoring infrastructure systems: An integrated network design and scheduling (INDS) problem. European Journal of Operational Research, 2012; 223(3):794–806.
30. Nurre SG, Sharkey TC. Integrated network design and scheduling problems with parallel identical machines: Complexity results and dispatching rules. Networks, 2014; 63(4):306–326.
31. Rudnick H, Mocarquer S, Andrade E, Vuchetich E, Miquel P. Disaster management. IEEE Power and Energy Magazine, 2011; 9(2):37–45.
32. Mas-Colell A, Whinston MD, Green JR. Microeconomic Theory. Oxford, UK: Oxford University Press, 1995.
33. Hernandez-Fajardo I, Dueñas-Osorio L. Sequential propagation of seismic fragility across interdependent lifeline systems. Earthquake Spectra, 2011; 27(1):23–43.
34. Adachi T, Ellingwood BR. Serviceability assessment of a municipal water system under spatially correlated seismic intensities. Computer-Aided Civil and Infrastructure Engineering.2009; 24(4):237–248.
35. Adachi T, Ellingwood BR. Comparative assessment of civil infrastructure network performance under probabilistic and scenario earthquakes. Journal of Infrastructure Systems, 2010; 16(1):1–10.
36. Roughgarden T. Selfish Routing and the Price of Anarchy. Cambridge, MA: MIT Press, 2005.