



Innovative Applications of O.R.

Restoring infrastructure systems: An integrated network design and scheduling (INDS) problem

Sarah G. Nurre^{a,1}, Burak Cavdaroglu^{a,2}, John E. Mitchell^b, Thomas C. Sharkey^{a,*}, William A. Wallace^{a,2}

^a Department of Industrial and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180, United States

^b Department of Mathematical Sciences, Rensselaer Polytechnic Institute, Troy, NY 12180, United States

ARTICLE INFO

Article history:

Received 10 June 2011

Accepted 5 July 2012

Available online 16 July 2012

Keywords:

Infrastructure restoration

Extreme events

Network design

Scheduling

Dispatching rules

ABSTRACT

We consider the problem of restoring services provided by infrastructure systems after an extreme event disrupts them. This research proposes a novel integrated network design and scheduling problem that models these restoration efforts. In this problem, work groups must be allocated to build nodes and arcs into a network in order to maximize the cumulative weighted flow in the network over a horizon. We develop a novel heuristic dispatching rule that selects the next *set* of tasks to be processed by the work groups. We further propose families of valid inequalities for an integer programming formulation of the problem, one of which specifically links the network design and scheduling decisions. Our methods are tested on realistic data sets representing the infrastructure systems of New Hanover County, North Carolina in the United States and lower Manhattan in New York City. These results indicate that our methods can be used in both real-time restoration activities and long-term scenario planning activities. Our models are also applied to explore the effects on the restoration activities of aligning them with the goals of an emergency manager and to benchmark existing restoration procedures.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

The restoration of services provided by infrastructure systems is critical for society to recover from extreme events. Therefore, the managers of these systems are faced with demanding choices in formulating their restoration efforts after the event. This research proposes a novel class of integrated network design and scheduling problems that can be used to model the formulation of these restoration efforts. The infrastructure system is modeled as a network, where flows in it represent the services provided by the system (see, e.g., Lee et al., 2007). Damage to the components of the infrastructure is then modeled as the removal of nodes and arcs in the network. The restoration efforts associated with the infrastructure will focus on installing or repairing physical components within the system and can be modeled as installing nodes and arcs into the network. Therefore, we can view this selection of nodes and arcs as *network design* decisions. Traditional network design problems are often only concerned with the performance of the end

design of the network; however, the driving performance metric, especially in the eyes of the public, in evaluating the restoration efforts is how well the services provided by the system come back online. This means that the network design decisions will be evaluated as they are being implemented, so that the scheduling decisions associated with them will have a significant impact on the objective. In particular, the performance of the network at time t , which is composed of the original network plus the nodes and arcs completed by t , will be evaluated by determining the (weighted) amount of flow that can be sent from supply nodes to demand nodes. This directly models the focus of the restoration efforts: it is on restoring services rather than on the monetary cost of the efforts.

The analysis of civil infrastructure systems is complex since they are interdependent (see O'Rourke, 2007); disruptions in one can spread to others causing cascading failures (see Wallace et al., 2003; Mendonca and Wallace, 2006; Chang et al., 2007). Rinaldi et al. (2001) note that managers of the infrastructure systems have become inclined to consider these interdependencies; however, the managers of a particular infrastructure will have little knowledge of the structure and operations of the other systems. We can expect that the managers of an individual infrastructure will understand the direct connections of it with other infrastructures and, therefore, can weigh the services provided to certain connections more heavily (e.g., a hospital has a higher weight than a residential household). Our class of problems can be used to model this situation and, further, explore the effects on the

* Corresponding author. Tel.: +1 518 276 2958.

E-mail address: sharkt@rpi.edu (T.C. Sharkey).

¹ The work of this author was supported by a Sandia National Laboratories/Rensselaer Polytechnic Institute (RPI) Excellence in Engineering Research Fellowship.

² The work of this author was supported by the US Department of Homeland Security under Award Number: 2008-ST-061-ND 0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the US Department of Homeland Security.

restoration plan of an infrastructure when we consider these interdependencies.

There has been some research done on problems related to our class of integrated network design and scheduling problems. Guha et al. (1999) develop approximation algorithms for problems concerned with installing nodes into a power network to recover from disruptions; however, this work assumes that demand nodes simply need to be connected to supply nodes and thus does not model the capacity limitations of the system. Ang (2006) develops an integer programming formulation of the problem of scheduling the installation of a set of nodes and arcs into a power network in order to minimize the cost of power-shedding over time. This model focuses on the recovery efforts and therefore requires that all nodes and arcs are installed, meaning that it does not model the network design decisions often associated with the restoration efforts of an infrastructure after an extreme event. The integer programming formulation of Ang (2006) was only able to find solutions within 7% of optimal within an hour for a network with 300 nodes and, therefore, may not support real-time decision-making, especially on large-scale networks. Xu et al. (2007) apply a genetic algorithm to a problem associated with restoring power after an earthquake. The objective of this problem minimizes the average time that each customer is without power; therefore, this problem does not prioritize demand to critical points within the infrastructure. Matisziw et al. (2010) propose an integer programming model in order to restore networks where connectivity between pairs of nodes is the driving performance metric associated with the network. Cavdaroglu et al. (forthcoming) examined a model to determine the restoration efforts of a set of interdependent infrastructure systems. However, this work assumes that an individual infrastructure system fully understands the operations of the other systems. This means that it knows the exact network structure and operational decisions of the infrastructure based on the amount of services provided to it. It would be more reasonable to assume that an individual infrastructure system understands its connections to other systems and can thus prioritize services provided to these connections. It may not be appropriate to apply the model of Cavdaroglu et al. (forthcoming) to the restoration efforts of a single infrastructure, especially when there is not a centralized coordinator in the emergency response efforts.

This paper proposes an integrated network design and scheduling problem that can be applied to the restoration efforts of a variety of infrastructure systems. This problem focuses on selecting a set of nodes and arcs for installation into a network and then scheduling them on a set of parallel identical work groups. Once a node or arc is processed by one of these work groups, it becomes operational in the network and flow can be routed through it or over it. At each time period, the performance of the network is evaluated by determining the maximum amount of (weighted) demand that can be sent from a set of supply nodes to a set of demand nodes over the *operational network* (i.e., the original network plus those nodes and arcs completed by the work groups by this time). The objective function of this problem then focuses on maximizing the cumulative performance of the network over a finite horizon.

We propose an integer programming formulation of this problem and discuss valid inequalities for it that improve the effectiveness of solving it with a commercial software package. One of these inequalities specifically links the network design and scheduling decisions. This integer programming formulation can be useful in the long-term scenario planning activities for the managers of the infrastructure systems. These activities increase the preparedness of the managers and, therefore, will result in more effective decision-making after an actual extreme event. However, the time required to solve the integer programming formulation may

prohibit it from being useful in the real-time restoration activities after an extreme event. Therefore, we propose a heuristic dispatching rule for this problem that integrates fundamental concepts from network flows (the residual network optimality conditions) and scheduling (the weighted shortest processing time dispatching rule). The rule focuses on selecting the next set of tasks to be processed by the resources rather than traditional dispatching rules that simply focus on selecting the next individual task to be processed. This dispatching rule is shown to provide near-optimal solutions (typically on the order of 2–3% of optimal for case studies where the optimal solution is known) to realistic data sets and can be utilized in real-time restoration activities. Both the dispatching rule and integer programming formulation (with its valid inequalities) are tested extensively on realistic data sets representing infrastructure systems in New Hanover County, North Carolina in the United States and lower Manhattan in New York City. These results demonstrate the power of our proposed methods along with providing insight into the impact on the restoration plan of an infrastructure when it is aligned with the priorities of the emergency managers of the region.

The remainder of this paper is organized as follows. Section 2 provides the mathematical formulation of our integrated network design and scheduling problem. Section 3 focuses on optimization methods to solve the problem including the heuristic dispatching rule (Section 3.1) and valid inequalities (Section 3.2). Section 4 applies our optimization methods for the integrated network design and scheduling problem on realistic case studies associated with infrastructure systems of New Hanover County, North Carolina. These case studies were created through collaborations with managers of these infrastructure systems and the emergency manager of the County. Section 5 focuses on case studies associated with the power infrastructure of lower Manhattan in New York City. We provide concluding remarks in Section 6.

2. The integrated network design and scheduling problem

The mathematical model of our integrated network design and scheduling (INDS) problem involves a network $G = (N, A)$ where N is the set of nodes and A is the set of arcs. The node set N and arc set A can be viewed as the 'operational network' immediately after the extreme event, i.e., these sets represent the components of the infrastructure that are unaffected by the event. There is a set of supply nodes, $S \subseteq N$, and a set of demand nodes, $D \subseteq N$. Each arc $(i, j) \in A$ has an associated capacity u_{ij} while each supply node $i \in S$ has a supply capacity s_i per time period and each demand node $i \in D$ has a demand d_i per time period. We are interested in sending flow (respecting the flow capacities of the arcs and the supply/demand capacities of the nodes) from the supply nodes to the demand nodes where each unit of flow that arrives at demand node $i \in D$ is given a weight of w_i . The performance of the network is evaluated by determining the maximum amount of weighted flow that can be sent from the supply nodes to the demand nodes. There is a set of arcs, A' , that we can install into the network. Without loss of generality, this can model problems where we can install both nodes and arcs into the network since a node can be 'split' into two nodes and an arc. We are interested in scheduling a subset of the arcs in A' onto a series of parallel identical work groups, $k = 1, \dots, K$, in order to install them into the network. The identical work group assumption is practical in the context of single infrastructure restoration since the differences between the work crews are often negligible with respect to the units used to measure the processing times. Each arc $(i, j) \in A'$ has an associated processing time, p_{ij} , and capacity, u_{ij} . We assume, without loss of generality, that the processing times are integral. We further assume that we are in a non-preemptive environment so that a task

infrastructure (see the Appendix for its details) to model its operations. We will examine the INDS problem with the DC model in our computational testing. The INDS problem can be directly applied to water, waste water, and supply chain systems since the flow in these systems can be directly controlled. We note that the types of components in these systems that are vulnerable to damage will vary based upon the type of extreme event (e.g., the pipes of the waste water system are vulnerable to earthquakes while its pump stations are more vulnerable to flooding).

It is possible to perform pre-processing on this integer programming formulation in order to reduce the capacities of the arcs in the network. This is important since, for arcs in A' , a reduction in the capacity will require a larger β_{ijt} value in the relaxation in order to achieve a particular flow. This in turn requires more scheduling resources to be dedicated to processing this arc. The capacity reduction techniques seek to tighten the capacity of an arc, u_{ij} , where $(i,j) \in A'$. It is clear that we could determine the maximum flow placed on arc (i,j) in any potentially feasible flow in the network $(N, A \cup A')$ and reduce u_{ij} to this value. However, if arc (i,j) belongs to a directed cycle, this value could be arbitrarily high since flow may move through this cycle without ever arriving at a demand node. We can thus restrict ourselves to feasible flows where every unit of flow reaches some demand node (i.e., the flow can be decomposed into a series of paths). Therefore, we can reduce u_{ij} to be equal to the minimum cut value that separates j from the set of demand nodes D , since every unit of flow on arc (i,j) must cross this cut to reach a demand node. We can further reduce the capacity of arc $(i,j) \in A'$ when all demand node weights are equal. We will assume that, without loss of generality, the network has a single supply and demand node. This further reduction focuses on the additional flow sent through the network by installing some subset of arcs in A' . It relies on the following result.

Lemma 2.1. *There exists a maximum flow in the network $\bar{G} = (N, A \cup \bar{A})$, where $\bar{A} \subseteq A'$ such that the flow on arc $(i,j) \in \bar{A}$ is less than or equal to $\bar{v} - v$ for all $(i,j) \in \bar{A}$, where \bar{v} is the maximum flow in the network \bar{G} and v is the maximum flow in the network G .*

Proof 1. We will construct a solution that satisfies this property by applying the augmenting path algorithm (see Ahuja et al., 1993) starting from the maximum flow in G in order to determine the maximum flow in \bar{G} . The proof will be by induction on iteration ℓ of the augmenting path algorithm. We let x_{ij}^ℓ denote the flow on arc (i,j) and \bar{v}^ℓ represent the flow into the demand node after iteration ℓ of the algorithm. It is clear that for $\ell = 0$, that $x_{ij} \leq \bar{v}^\ell - v$ for all $(i,j) \in \bar{A}$. We assume the claim holds up to iteration ℓ and show it holds for $\ell + 1$. Let P be the augmenting path found in iteration $\ell + 1$. By definition, the algorithm pushes $\bar{v}^{\ell+1} - \bar{v}^\ell$ flow along this path. The only way for the flow on arc $(i,j) \in A'$ to increase is for it to be part of this path. For this type of arc, we know that $x_{ij}^{\ell+1} = x_{ij}^\ell + (\bar{v}^{\ell+1} - \bar{v}^\ell) \leq \bar{v}^\ell - v + \bar{v}^{\ell+1} - \bar{v}^\ell = \bar{v}^{\ell+1} - v$. Our desired result holds by induction. \square

Lemma 2.1 holds for any set $\bar{A} \subseteq A'$. We can, therefore, reduce the capacity of arc $(i,j) \in A'$ to $v' - v$ where v' is the maximum flow in $(N, A \cup A')$. Note that there is overhead in these pre-processing steps to reduce the capacity of the arcs and, therefore, we will only apply them to problems where the initial integer programming formulation cannot be solved quickly.

3. Solution methods: Dispatching rules and valid inequalities

This section is focused on the development and analysis of solution methods for the integrated network design and scheduling problem. We first develop a novel heuristic dispatching rule for

the problem that selects the next set of arcs to be processed by the work groups. This dispatching rule will integrate fundamental concepts from the fields of network flows and scheduling. We then discuss families of valid inequalities for the integer programming formulation of the INDS problem.

3.1. Dispatching rules

There has been a significant amount of research in parallel machine scheduling on so-called dispatching rules (Pinedo, 2008). These rules often characterize the desirability of scheduling a certain task by estimating its contribution to the objective function and then greedily schedule the unscheduled task with the 'best' desirability. The key in developing a heuristic dispatching rule for our INDS problem is to understand how completing a task or, equivalently, an arc impacts the objective function of the problem. The installation of an arc (i,j) can improve the performance of the network by increasing the amount of weighted flow in it. This will then impact the objective function for the remainder of the horizon of the problem. We could attempt to develop a more traditional dispatching rule for our INDS problem by examining, for each uninstalled arc $(i,j) \in A'$, the improvement in the performance of the network by installing arc (i,j) into it. A modification of the classic weighted shortest processing time (WSPT) first rule (see Smith, 1956) would then select and schedule the arc that maximizes the ratio of the improvement by installing the arc and the processing time of the arc. However, this is short-sighted in the sense that certain arcs will not increase the weighted flow if they are not performed in sequence with other arcs. For example, there may exist a path of arcs from a supply node to a demand node that needs to be installed in order to increase the weighted flow. The installation of each of these arcs may not increase the weighted flow but the installation of this path may significantly increase the weighted flow. Therefore, our dispatching rule should be concerned with the contributions to the weighted flow of installing a set of arcs.

It is easiest to motivate the dispatching rule for the INDS problem by focusing on a special class of the problem where each demand node is given the same weight. In other words, the performance of the network is simply concerned with maximizing the total flow from the supply nodes to the demand nodes. We can assume that the performance of the network is evaluated by determining the maximum flow from a single supply node s to a single demand node τ . This assumption is without loss of generality by applying a standard network expansion technique to multiple supply/demand node networks. It is well-known that the residual network associated with a maximum flow in the network does not contain an augmenting path from the supply node s to the demand node τ (see Ahuja et al., 1993). Therefore, in order to increase the amount of flow sent from s to τ in the current operational network, we must install a set of arcs that form some residual path between the source and the sink node. Our dispatching rule for this INDS problem will then select a set of uninstalled arcs that belong to some residual path and that maximizes the ratio of the residual capacity of the path and the cumulative processing times of the uninstalled arcs in it. Mathematically, suppose that x^* is the current optimal flow in the network composed of the original arcs and the installed arcs from A' . The arcs in the residual network associated with x^* will have a residual capacity of r_{ij} and a processing time of $p_{ij} = 0$ since they are already installed. The uninstalled arcs in A' will have a residual capacity of u_{ij} (their original capacity) and a processing time of p_{ij} (their original processing time). We then define the residual capacity of path P as $r(P) = \min_{(i,j) \in P} r_{ij}$ and the processing time of a path as $p(P) = \sum_{(i,j) \in P} p_{ij}$. We are then interested in scheduling the uninstalled arcs in the path that is an optimal solution to the problem

$$\max_{P \in \Phi: p(P) > 0} \frac{r(P)}{p(P)} \quad (12)$$

where Φ is the set of all paths from s to τ in the network composed of the residual network and all uninstalled arcs. We further note that if $p(P) = 0$, this implies that all arcs are already installed in the network and that $r(P) = 0$ since x^* is the maximum flow in the network. The numerator of (12) provides a measure of the amount of additional flow in the network by installing arcs in P while the denominator provides a measure of the resources required to process P . For a single work group, $p(P)$ is precisely the ‘makespan’ required to complete all uninstalled arcs in the path. For multiple work groups, we could alter the denominator to approximate the makespan of the path by including the leading term $\frac{1}{k}$. This provides a lower bound on the actual makespan for the path. In terms of the optimization problem, this leading term is a constant for all paths, so optimizing (12) is equivalent. Although this is only a lower bound on the makespan of a path for multiple work groups, the dispatching rule is shown to provide solutions of high-quality.

It is more difficult to determine an optimal path to (12) than it is to determine the next job according to the WSPT rule since we cannot decompose (12) by arcs. We will now discuss a combinatorial algorithm to determine an optimal solution to (12). The idea for the algorithm is motivated by the following observation: if we know that $r(P^*)$ is the numerator in an optimal solution to (12), then P^* is the path with the shortest processing time in the network where we only include arcs whose residual capacities are greater than or equal to $r(P^*)$. This immediately leads to an algorithm to solve (12): for each potential value of the numerator (i.e., the residual capacity of a path), we determine the shortest processing time path in the network containing only arcs whose residual capacities are above the numerator. An optimal solution is then the path obtained in this procedure that has the maximum ratio of residual capacity to processing time. We note that this procedure is easily adapted to situations where a constraint is placed on the denominator (for example, if we are in a single work group setting and at time t , we do not want to select a path with a processing time greater than $T - t$) by terminating the procedure when the shortest path in the network exceeds the threshold of the constraint. Note that the residual capacity of a path is the minimum residual capacity of the arcs in the path, so there are at most $2(|A| + |A'|)$ distinct values to be considered. This means that we can determine the next set of arcs to be processed by solving $O(|A| + |A'|)$ shortest path problems.

The dispatching rule for this INDS problem will determine the first set of tasks to be processed by solving the problem (12). We will assign these tasks (according to the longest processing time first rule) to the available work groups until all tasks from this set are processed. In other words, we can view the tasks that need to be processed as a queue and we will process the next task in the queue whenever a work group becomes available. If no tasks are in the queue, we will then determine the next set of arcs to be processed by considering the residual network associated with an optimal solution to the maximum flow problem where all arcs that are currently being processed are assumed to be available in the network. This process will continue until either all tasks are processed or we reach the end of the horizon.

We now discuss the dispatching rule for the INDS problem where the demand nodes can have different weights. In particular, we will select a residual path from some source node $j \in S$ to some demand node $i \in D$ that maximizes the ratio of w_i times the residual capacity of the path and the cumulative processing times of the uninstalled arcs in it. We can determine this set of uninstalled arcs by solving $O(|D|)$ problems of the form (12)–one for each distinct node $i \in D$, viewing that node as the ‘super-demand’ node τ .

However, in the method for solving (12), note that for a fixed residual capacity, we are solving a shortest path problem on the same set of arcs regardless of the node $i \in D$. Therefore, because of the structure of Dijkstra’s algorithm, we can determine the relevant information (i.e., shortest path from the super-source node to each $i \in D$) for each of the $O(|D|)$ problems of the form (12) by solving $O(|A| + |A'|)$ shortest path problems. We will now formally present the algorithm to determine the next set of arcs to be processed according to the dispatching rule for the core INDS problem. This algorithm assumes that we are working with the network where we have created a ‘super-supply’ node s and have calculated the residual network, $G(x^*)$, associated with the current optimal weighted flow, x^* , in this network. We denote r_{ij} as the residual capacity of arc (i, j) where if arc (i, j) is uninstalled in the network $r_{ij} = u_{ij}$. The notation $G(N, A(r))$ is used to denote the network composed of only arcs with a residual capacity such that $r_{ij} \geq r$ where $(i, j) \in A \cup A'$ or $(j, i) \in A \cup A'$. Algorithm 1 provides the pseudo-code for determining the path selected by the dispatching rule.

Algorithm 1. Algorithm for Path Selection in the Dispatching Rule

-
- 1: Set MaxRatio = 0 and $P^* = \text{null}$.
 - 2: Sort the residual capacities of all residual arcs (i.e., arcs in the set $\{(i, j), (j, i) : (i, j) \in A \cup A'\}$) in non-decreasing order and put them into array R .
 - 3: **for** $\ell = 1, \dots, 2(|A| + |A'|)$ **do**
 - 4: Determine the shortest path distance labels, $d(i, R[\ell])$ for $i \in D$, from the source node s in the network $G(N, A(R[\ell]))$.
 - 5: **for all** $i \in D$ **do**
 - 6: **if** $\frac{w_i R[\ell]}{d(i, R[\ell])} > \text{MaxRatio}$ **then**
 - 7: Set MaxRatio = $\frac{w_i R[\ell]}{d(i, R[\ell])}$.
 - 8: Set P^* to be the shortest path from s to i .
 - 9: **end if**
 - 10: **end for**
 - 11: **end for**
 - 12: Return P^* .
-

3.2. Valid inequalities for IP formulation

The purpose of this section is to present valid inequalities in order to strengthen the bounds provided by the linear programming relaxation of the IP formulation. We focus on three families of valid inequalities: (i) shortest processing time path inequalities, (ii) flow cover inequalities, and (iii) β -conservation inequalities. The main idea behind the latter two families is to increase the β_{ijt} variables in order to achieve the desired level of flow on arc (i, j) in time period t . For example, the capacity reduction techniques from Section 2.1 focus on reducing the capacity of an arc to $\bar{u}_{ij} < u_{ij}$ in order to drive up β_{ijt} . This implies that $\beta_{ijt} \geq x_{ijt} / \bar{u}_{ij} > x_{ijt} / u_{ij}$ and thus β_{ijt} will need to be at a higher level to achieve the flow of x_{ijt} . This, in turn, requires more scheduling resources to be dedicated to arc (i, j) prior to time period t . The valid inequalities in this section are similar in nature and, therefore, help to ‘link’ the flow variables with the scheduling decisions.

The shortest processing time path constraints are a generalization of constraints (9) where we also consider the time required to build a path of arcs to node i . In particular, we know that the flow on arc $(i, j) \in A'$ cannot be positive until we install some path from a supply node to node i and arc (i, j) into the network. We are, therefore, concerned with the shortest processing time path from any supply node to node i plus the process time of arc (i, j) . We then know that arc $(i, j) \in A'$ cannot have flow on it until this value

divided by the number of work groups (i.e., the lower bound on the makespan to complete this path and arc (i, j)). This helps rule out 'partial' installations of a path to arc (i, j) before this time in the linear programming relaxation. Formally, let \mathcal{P}_{ij} be equal to the shortest processing time path value (where an arc in A has a processing time of 0) from some supply node to node i plus the processing time of arc (i, j) , p_{ij} . The shortest processing time path constraints are then defined as:

$$\sum_{t=1}^{(\mathcal{P}_{ij}/K)-1} \beta_{ijt} = 0 \quad \text{for } (i, j) \in A'. \tag{13}$$

The flow cover inequalities of Gu et al. (1999) can be applied to the network during each time period t for each arc $(i, j) \in A \cup A'$ and each cutset $C(i)$ separating all supply nodes from i , provided that the cutset contains only arcs in A' , i.e., $C(i) \subseteq A'$. This would be especially applicable, for example, in disruptions to infrastructures that wipe out entire portions of the network. These inequalities link the flow on arc (i, j) in time period t with the availability of arcs in the cutset, since any flow on arc (i, j) must be on some arc in $C(i)$. We can define these inequalities as

$$x_{ijt} \leq u_{ij} \sum_{(\kappa, i) \in C(i)} \beta_{\kappa it} \quad \text{for } t = 1, \dots, T. \tag{14}$$

Let $O(i)$ be the set of nodes such that there exists an arc $(\kappa, i) \in A'$. If the set of arcs $\{(\kappa, i): \kappa \in O(i)\}$ is a valid choice for $C(i)$ then the corresponding inequality focuses on those arcs that are 'one step back' from node i in the network, i.e., they come directly into node i . We can apply the same type of inequalities to arcs that are 'two steps back' from node i in the network, i.e., those arcs that come into some node $\kappa \in O(i)$, as long as all arcs two steps back are in A' . This family of inequalities can be extended to an arbitrary number of steps back from node i as long as all arcs that far back are in A' .

The β -conservation inequalities are motivated by the fact that if arc (i, j) is operational (e.g., has flow on it), then some arc into node i must also be operational. They are similar to the one step back flow-cover inequalities since they can only be applied to an arc where all arcs coming into node i belong to the set A' . These inequalities are

$$\beta_{ijt} \leq \sum_{\kappa \in O(i)} \beta_{\kappa it} \quad \text{for } t = 1, \dots, T. \tag{15}$$

These constraints help link the scheduling decisions of a path of arcs by forcing those that are earlier in the path to be operational when the arcs later in the path are operational. We now prove a relationship between these inequalities and the flow-cover inequalities.

Theorem 3.1. *The flow cover inequalities for one-step and two-step back cutsets are implied by the β -conservation inequalities, if the network has no parallel arcs (where two arcs are considered parallel if they both leave node i and enter node j).*

Proof 2. Note that

$$x_{ijt} \leq u_{ij} \beta_{ijt} \leq u_{ij} \sum_{\kappa \in O(i)} \beta_{\kappa it},$$

which shows that the β -conservation inequalities imply the one-step back flow cover inequalities. For the two-step back flow cover inequalities, we apply the β -conservation inequalities to each node $\kappa \in O(i)$ to yield

$$u_{ij} \sum_{\kappa \in O(i)} \beta_{\kappa it} \leq u_{ij} \sum_{\kappa \in O(i)} \sum_{q \in O(\kappa)} \beta_{q\kappa t},$$

where no β term is repeated because there are no parallel arcs. Note that if there was some arc into $\kappa \in O(i)$ that is in A , the two-step

back inequalities would not be valid for arc (i, j) as well as the β -conservation constraint for κ . \square

4. Computational testing on the New Hanover County data set

We will now discuss the results of case studies of applying the INDS problem to realistic data sets representing infrastructure systems in New Hanover County, North Carolina in the United States. New Hanover County is a coastal county in southern North Carolina that includes the city of Wilmington and the Cape Fear beaches. These data sets were created through extensive collaborations with the managers of the infrastructure systems in New Hanover County as well as collaborations with the emergency manager of the County. This section focuses on applying the INDS problem to three separate infrastructure systems: (i) the power infrastructure of the County, (ii) the waste water infrastructure of the county, and (iii) an emergency supply chain infrastructure in the city of Wilmington. The disruptions to the power and waste water infrastructures model the effects of a strong hurricane whose eye passes to the south of the county so that both extensive flooding and wind damage are possible, which is similar to the scenario caused by Hurricane Ophelia in Septemeber 2005. The emergency supply chain infrastructure models situations in which the emergency manager of the county must deliver critical goods to the affected population in Wilmington that were unable to leave the city prior to the hurricane.

In the analysis of these case studies, we are interested in examining the performance of the heuristic dispatching rule and the integer programming formulation of the INDS problem. It is especially important to determine the potential for these methods to be used in real-time restoration activities. Therefore, our testing was performed on a laptop with a 2.16 GHz Intel Core 2 Duo Processor with 3 GB of RAM, which would be similar to the computing resources available during real-time restoration activities. We have used CPLEX 12.0 in order to solve the integer programming formulation of the INDS problem. We have chosen to only examine the core integer programming formulation of the INDS problem since it is solved rather quickly (so the capacity reduction techniques and the shortest processing time path constraints were not necessary) and the disruptions do not wipe out large portions of the network (so the flow cover and β -conservation inequalities are not applicable).

The weights associated with the demand nodes in the network can be utilized so that the restoration plan is aligned with certain goals. The most straightforward goal would be to focus on evaluating the performance of the network by measuring the amount of demand met in it. This would be modeled by setting all weights to be equal across demand nodes and can be applied when an infrastructure is interested in maximizing its revenues (since each unit of met demand provides the same revenue). An alternative goal would be to focus on aligning the restoration plan with the goals of the overall response efforts to the extreme event. This is especially important for New Hanover County since the emergency manager of the county has certain priorities in the overall restoration efforts of the county that may differ from simply meeting demand in the infrastructure system. The emergency manager of the county would view the demand at certain 'high-priority' nodes as being more important (especially those nodes that feed critical facilities in other infrastructures) than demand at 'low-priority' nodes. Therefore, they will measure the performance of the network as the maximum amount of weighted demand met (or, equivalently, the weighted maximum flow), where each node is assigned a weight based on its priority level. For New Hanover County, the priorities of the emergency manager of the county, in decreasing order of their importance, are: (1) emergency

communications centers, (2) hospitals and emergency shelters, (3) police and fire departments, (4) all components in the water and wastewater infrastructures, and (5) all other components (including residential households).

This research will explore the impact to the restoration plan when it is aligned with the goals of the county. This will involve determining the restoration plan by solving an INDS problem where the weights of the nodes in the network reflect the priorities of the county and then evaluating its performance under the objective function where all demand weights are equal. The deviation of this restoration plan from the optimal objective function value to this 'equal demand weight' INDS problem measures the impact on the restoration plan of aligning it with the priorities of the county.

4.1. Power infrastructure case study

The focus of this section is on examining a case study representing the power infrastructure of New Hanover County. The network model of this infrastructure has 377 nodes and 386 arcs under normal operations. It may seem that there is not much redundancy in the system but we note that there is a good level of redundancy in the *transmission* network in the county. The managers of the power infrastructure suggested that the model should focus on the transmission network of the county—so that the distribution networks to the demand points are rather simple. In particular, we model the distribution networks as point-to-point connections with the transmission substations. There are 37 nodes and 46 arcs in the transmission network. We have added 340 demand nodes and 340 arcs connecting these nodes to appropriate distribution substations in our model of this infrastructure. These demand nodes and arcs then model the distribution network in the county. The INDS problem has $|N| = 377$ nodes, $|A| = 346$ arcs, and $|A'| = 40$ arcs as network design decisions. The horizon of the problem is equal to $T = 30$ where, roughly, each time period represents a six hour block of time so that the horizon is roughly a week. We consider two different types of weights on the performance of the network (i.e., μ_t for $t = 1, \dots, T$) where the first class ('Constant') weighs the performance evenly over the horizon and the second class ('Scaled') weighs the performance more heavily later in the horizon by setting $\mu_t = t/T$. We further consider two different types of weight for the demand nodes: one where each unit of met demand is

Table 1

The performance of the dispatching rule and IP formulation on the power infrastructure case study.

K	μ_t	w_i	Dispatching rule		IP	
			Time (s)	Opt. gap (%)	Time (s)	Opt. gap (%)
1	Constant	Constant	1.77	0.00	4.54	0.00
1	Constant	Priority	2.24	0.00	5.22	0.00
1	Scaled	Constant	1.92	0.00	4.56	0.00
1	Scaled	Priority	2.80	0.00	6.14	0.00
2	Constant	Constant	2.47	0.12	5.89	0.00
2	Constant	Priority	3.44	0.12	6.02	0.00
2	Scaled	Constant	2.98	0.14	6.39	0.00
2	Scaled	Priority	3.35	0.12	6.24	0.00

weighed evenly across the demand nodes ('Constant') and one where the demand nodes with higher priorities according to the emergency manager of the county have larger weights ('Priority'). The number of work groups in this study is equal to $K = 1$ or $K = 2$ which represents the fact that the county itself has very limited resources.

Table 1 provides the computational performance of the heuristic dispatching rule and the integer programming formulation of the INDS problem. Note that the gap of the solution methods was 'normalized' by setting the performance of the network in each time period equal to its maximum weighted flow minus the maximum weighted flow in the network without any arcs from A' . Therefore, the initial performance of the network does not bias the gap of the solution methods. These results indicated that the dispatching rule obtains high-quality (and optimal in some instances) solutions to the INDS problem extremely quickly. CPLEX 12.0 is able to identify an optimal solution to the problem quickly as well, although commercial software packages may not be available to managers of the infrastructure systems during real-time restoration activities.

The restoration efforts formulated by the dispatching rule and the INDS problem can assist in the decision-making of the managers of the power infrastructure. However, these managers may have limited mathematical expertise implying that it would be beneficial to reproduce the restoration efforts using visualization tools. Fig. 1 provides an example of such a visualization tool for the optimal restoration efforts for the problem with $K = 2$, Constant μ_t , and Priority demand node weights. This visualization tool was




Fig. 1. The improvement in the performance over time of the power infrastructure in New Hanover County for the optimal restoration efforts with $K = 2$, Constant μ_t , and Priority demand node weights. Black arcs represent the disrupted services and the bold black arcs are those currently being processed by a work group at a particular point in time. Once a work group completes the processing of an arc, the arc becomes operational and some disrupted services are restored.

Table 2

The performance of the restoration efforts obtained through the dispatching rule and the core INDS problem under the DC model.

K	μ_t	w_i	Dispatching rule		INDS problem		INDS Problem with DC model	
			Time (s)	Opt. gap (%)	Time (s)	Opt. gap (%)	Time (s)	Opt. gap (%)
1	Constant	Constant	1.77	2.21	4.54	2.21	33.13	0.00
1	Constant	Priority	2.24	2.00	5.22	2.00	26.72	0.00
1	Scaled	Constant	1.92	0.80	4.56	0.80	95.01	0.00
1	Scaled	Priority	2.80	0.75	6.14	0.75	323.14	0.00
2	Constant	Constant	2.47	1.09	5.89	0.13	130.70	0.00
2	Constant	Priority	3.44	0.99	6.02	0.18	154.73	0.00
2	Scaled	Constant	2.98	0.24	6.39	0.00	186.88	0.00
2	Scaled	Priority	3.35	0.26	6.24	0.06	66.01	0.00

produced using a geographic information system (GIS) to display data on the infrastructure.

The 'Constant' weights for the demand nodes model the objective of meeting demand within the power infrastructure. Therefore, the operations of the infrastructure under this performance metric would be to deliver the maximum amount of power possible throughout the system. It is clear that the 'Priority' weights for the demand nodes reflect the performance metric that aligns the restoration efforts of the power infrastructure with the overall goals of the response efforts of the emergency manager of the county. We have observed that, for all combinations of work groups (K) and performance weights (μ_t), the optimal restoration plan aligned with the priorities of the emergency manager meets the same level of demand over the horizon of the problem as the optimal solution to the 'demand-based' objective function. This implies that the optimal 'priority-based' restoration plan is also an optimal solution to the INDS problem with the 'demand-based' objective function. Therefore, the restoration plan of the power infrastructure can be aligned with the overall goals of the response effort without any effect on the amount of demand met over the horizon of the restoration efforts. We do note that this observation is a function of the network structure and disruption scenarios; however, it is still interesting that the 'priority-based' and 'demand-based' objective functions do not conflict for these case studies.

4.1.1. Incorporating the DC model into the INDS problem

As discussed in Section 2.1, there is a potential issue in applying the restoration efforts obtained either through the dispatching rule or integer programming formulations of the INDS problem: it was obtained with a primarily 'flow-based' model of the operations of the power infrastructure. As discussed in the Appendix, it is only possible to *indirectly* control power flows since they behave according to certain physical laws. Therefore, we will explore the quality of restoration efforts (i.e., which arcs to process and when to process them) obtained through these methods when the operations of the power infrastructure are modeled according to the DC model discussed in the Appendix. The performance of the network in each time period for the 'scheduling' solution returned by the dispatching rule is calculated through the DC model over the available arcs in the network. We, therefore, associate a value $v(DR)$ with the dispatching rule, which represents the cumulative performance of the DC model applied to the available network over the horizon of the problem. The gap of the dispatching rule is then calculated as

$$100\% * \frac{v^*(DC) - v(DR)}{v(DR)},$$

where $v^*(DC)$ is the optimal objective function value to the INDS problem with the DC model. The gap associated with the optimal scheduling solution to the core INDS problem was calculated in a similar fashion. In other words, the gaps presented in Table 2 can be viewed as the quality of the heuristic that determines a solution

Table 3

Performance (measured through optimality gaps) of current restoration procedures of the power infrastructure in New Hanover County.

K	μ_t	w_i	Power company restoration efforts (%)
1	Constant	Constant	12.06
1	Constant	Priority	10.03
1	Scaled	Constant	65.41
1	Scaled	Priority	50.73
2	Constant	Constant	14.05
2	Constant	Priority	12.67
2	Scaled	Constant	79.01
2	Scaled	Priority	65.82

to the 'INDS problem with DC model' by (i) first determining the restoration decisions using the 'flow-based' (or core) INDS problem and (ii) determining the operations of the network using the DC model in each time period over the operational network. Table 2 presents the performance of the restoration efforts returned by the dispatching rule and the optimal restoration efforts in the core INDS problem when compared with the optimal restoration efforts in the INDS problem with the DC network model. The performance of the restoration efforts obtained using the core INDS problem and the dispatching rule perform extremely well under the DC model of the problem. We further note that, for all combinations of work groups and performance weights the optimal solution to the INDS problem with the DC model and weighted demand nodes is also an optimal solution to the same problem with constant demand node weights.

4.1.2. Examining the INDS problem in the context of the power infrastructure's current restoration procedures

This section will focus on the use of the INDS problem as an alternative method to formulating the restoration efforts of the power infrastructure from their current procedure. The power company will tend to focus on the *unmet demand* in the network when it evaluates its performance since this is a more direct reflection of the disruptions caused by the damage and customers will report outages as opposed to services. Therefore, the objective function will be to minimize the cumulative (weighted) unmet demand over the horizon of the efforts, which is slightly different than the current objective of the INDS problem. This is not an issue because the two objective functions are equivalent (i.e., the same set of decisions optimizes both objective functions): the difference between the weighted flow in the fully restored network and the weighted flow in a partially restored network is exactly the weighted unmet demand in the partially restored network. Therefore, the decisions that maximize the weighted flow over the horizon will also minimize the weighted unmet demand. The distinction between these two objective functions is important in presenting the INDS problem as an alternative method since the managers of the power infrastructure will tend to think in terms of unmet demand. Therefore, in evaluating the current restoration

procedures of the power infrastructure, we will focus on this weighted unmet demand objective.

The power company in New Hanover County is Progress Energy and they discuss their restoration procedure in Progress Energy (Progress Energy, 2010). The company first repairs (or restores) all damaged components within the transmission network and then moves onto damaged components within the distribution networks. The order in which the damaged transmission components are repaired is based on how close they are to supply points in the network. The power infrastructure will repair transmission components close to the supply points first and progress further through the network as repairs are completed. Once the transmission network is repaired, the efforts then focus on the distribution networks. The power company will first work to restore services to critical customers (such as hospitals) and then repair damage that can restore power to large segments of the population. In this analysis, we have viewed the order of the 'critical customers' as the same order as the priorities of the county after the extreme event. Therefore, there is no bias in terms of applying this procedure to a different set of critical customers/priorities than presented in the weights of the INDS problem.

The focus of our comparison will be on situations when the power company utilizes temporary components in its restoration efforts. This situation typically arises in large-scale events when the emergency manager directs the power company to focus on restoring services rather than repairing damage (i.e., temporary components typically have shorter processing times). The procedure to determine the 'Power Company Restoration Efforts' will be to determine which option (i.e., repair the damaged component or install a temporary component) will be utilized to restore disrupted services caused by the damaged component and then implement these choices according to the procedure from Progress Energy (Progress Energy, 2010).

Table 3 provides the performance of these efforts (as measured through their optimality gaps of weighted unmet demand) when compared to the optimal solution to the INDS problem. These results indicate that the utilization of the INDS problem as a method to formulate the restoration efforts can improve the performance of these efforts. We do note that the results for the 'Constant' μ_t weights are, typically, reflective of the focus of the power infrastructure since unmet demand impacts a customer regardless of when it occurs over the horizon of the problem. The instances with the 'Constant' μ_t weights have that there still is room for between 10% and 12% improvements when utilizing the INDS problem to formulate the restoration efforts.

4.2. Waste water infrastructure case study

The focus of this section is on a case study of the waste water infrastructure in New Hanover County. The initial network model of this infrastructure has 543 nodes and 538 arcs. There is no redundancy in this infrastructure—it is composed of 5 separate components, which would correspond to trees in the network representation. Each of these components has a 'root' node that is a treatment plant and all arcs in a particular component are directed away from the plant. This may seem counter-intuitive since the waste water actually flows toward the treatment plant. However, modeling the infrastructure in this fashion allows us to view the requested service at a particular node as demand in the network representation. The disruption scenario in this case study represents damage to 25 pump stations that could occur simultaneously with the disruption scenario for the power infrastructure from Section 4.1. The pump stations are nodes in the network representation of the infrastructure, however, since the infrastructure is composed of 5 trees and all arcs are directed away from the root node of the tree, each node has one incoming arc into it. Therefore,

Table 4

The performance of the dispatching rule and IP formulation on the waste water infrastructure case study.

K	μ_t	w_i	Dispatching rule		IP	
			Time (s)	Optimality gap (%)	Time (s)	Optimality gap (%)
1	Constant	Constant	1.43	1.80	4.46	0.00
1	Constant	Priority	1.89	2.05	4.65	0.00
1	Scaled	Constant	1.50	1.33	4.56	0.00
1	Scaled	Priority	2.42	1.46	5.08	0.00
2	Constant	Constant	2.00	0.94	5.24	0.00
2	Constant	Priority	2.98	1.10	5.64	0.00
2	Scaled	Constant	2.52	0.42	5.07	0.00
2	Scaled	Priority	3.18	0.47	5.63	0.00

we can model the disruption to the pump station as damage to the incoming arc into it. Therefore, we have that $|N| = 543$, $|A| = 513$, and $|A'| = 25$ in this case study. We note that, since the damage occurs at a pump station of the infrastructure, it is necessary to fix that pump station rather than implement some alternative route. Therefore, the design decisions in A' correspond exactly to the arcs representing the damaged pump stations. The horizon of the problem is set equal to $T = 30$, where each time period represents a roughly six hour block of time.

Table 4 provides the computational performance of the heuristic dispatching rule and the integer programming formulation for the INDS problem. We again consider two types of performance weights (μ_t) and demand node weights (w_i). These computational results are very similar to the results for the power infrastructure case study: the dispatching rule provides high-quality solutions in seconds and can thus be used in real-time activities. Further, the application of CPLEX 12.0 to the integer programming formulation of the INDS problem provides an optimal solution to the problem in real-time. We again investigate the effects on the restoration plan of the waste water infrastructure by formulating it considering the priorities of the emergency managers. This investigation demonstrated that the optimal restoration plan to the problem with priority-based weights for the demand nodes meets the same level of demand over the horizon of the problem as the optimal solution to the 'demand-based' objective function (i.e., the INDS problem with constant demand node weights). Therefore, this indicates that, for this case study, the restoration plan of the waste water infrastructure can be aligned with the overall goals of the response effort without any effect on the amount of demand met over the horizon of the restoration efforts. Fig. 2 provides a visualization tool for the restoration efforts of the problem with $K = 2$ work groups, Constant μ_t , and Priority demand node weights for the waste water infrastructure.

4.3. Emergency supply chain infrastructure case study

The focus of this section is on a case study corresponding to setting up an emergency supply chain infrastructure in New Hanover County. This emergency supply chain will deliver a critical good (e.g. food or water) to those affected by the hurricane that did not have the resources to evacuate the area prior to its landfall. The supply chain will set up a number of distribution sites in lower income areas. Therefore, this case study focuses primarily on setting up distribution sites in the city of Wilmington since all other areas have a median income of over \$25,000 dollars. The potential locations for these distribution sites correspond to well-known, central entities in the city (such as malls, schools, and parks) and were determined through input from the emergency manager of the county. In this case study, these critical goods will flow into the county through the American Red Cross, be routed to the

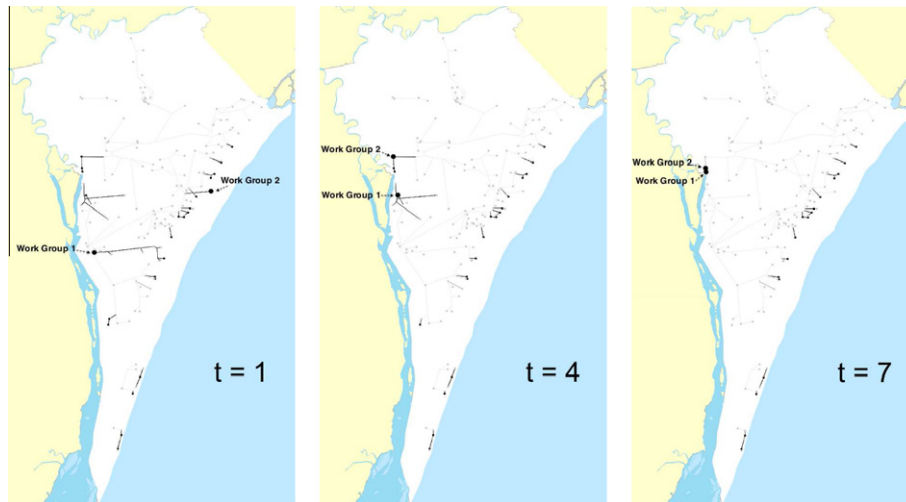


Fig. 2. The improvement in the performance over time of the waste water infrastructure in New Hanover County for the optimal restoration efforts with $K = 2$, Constant μ_t , and Priority demand node weights. Black nodes (squares) and arcs represent the disrupted services and the bold black nodes (circles) are those currently being processed by a work group at a particular point in time. Once a work group completes the processing of a node, the node becomes operational and some disrupted services are restored.

Table 5

The performance of the dispatching rule and IP formulation on the emergency supply chain infrastructure case study.

K	μ_t	w_i	Dispatching rule		IP	
			Time (s)	Optimality gap (%)	Time (s)	Optimality gap (%)
1	Constant	Constant	0.01	2.92	3.51	0.00
1	Scaled	Constant	0.01	2.63	3.30	0.00
2	Constant	Constant	0.04	4.71	7.36	0.00
2	Scaled	Constant	0.04	7.85	4.31	0.00

operational distribution sites, and then ‘delivered’ to the affected population that come to the sites. The network model for the INDS problem representing this case study will have two nodes (say i and i') and an arc $((i, i')$ in A') representing each potential distribution site. There will be an arc from the Red Cross to each potential distribution site and then arcs from each potential distribution site to the different populations to which that site can deliver goods. The demand for these different populations was created using census tract information and the arcs from the potential sites to the populations were included based on whether the site was close to the population. This case study has $|N| = 59$ nodes, $|A| = 85$ arcs, and $|A'| = 23$ design decisions. The problem has a horizon of $T = 12$, where each time period corresponds to 6 h. Table 5 presents the computational performance of the various solution methods on this case study for problems where all demand is weighted equally. The heuristic dispatching rule does provide good quality solutions extremely quickly. CPLEX 12.0 was able to solve the integer programming formulation of this problem in real-time as well.

5. Computational testing on the lower Manhattan data set

We will now discuss the results of a case study of applying the INDS problem to a realistic data set representing the power infrastructure of lower Manhattan in New York City. We will consider a case study which represents the effects of the disruption scenario of the failure of components in and around the Brooklyn-Battery tunnel. Therefore, this disruption scenario has a large portion of the network wiped out, meaning the valid inequalities presented in Section 3.2 are applicable. This data set was created through close collaborations with officials in this system and was first presented in

the work of Lee et al. (2007). The network in the INDS problem has $|N| = 1810$ nodes and $|A| = 2621$ arcs and is significantly larger than the networks from the case studies considered in Section 4. Further, the number of design alternatives is $|A'| = 695$ arcs. We have considered problems with $K = 1, 2, 3$ work groups, where the latter size is representative of the resources considered in Lee et al. (2007). The horizon of the INDS problem is $T = 60$ time periods.

Table 6 presents the computational results for the heuristic dispatching rule and core integer programming formulation for the lower Manhattan data set. The results indicate that CPLEX 12.0 is not capable of determining (or verifying) the optimal solution to the problem within a six hour time limit for the core integer programming formulation of the problem. We note that we have chosen to provide CPLEX 12.0 with the restoration plan determined by the dispatching rule as a warm-start solution. The gap of the dispatching rule is an *overestimate* on its actual gap since CPLEX 12.0 did not determine the optimal solution to the instance for any of the formulations (including ones strengthened with the valid inequalities) tested in this section. The reported optimality gap overestimate for an instance was calculated by comparing the objective function of the dispatching rule solution with the *best known* upper bound identified across all formulations for the instance tested. It is important to note that CPLEX 12.0 did not identify a better solution than the one provided by the dispatching rule for any instance of the problem in the six hour time limit. This means that the dispatching rule is more appealing to apply to the INDS problem than CPLEX 12.0 for these large-scale problems. Note that, due to the high running times of the core INDS problem, we have not examined the INDS problem with the DC Model for this case study since it is imperative in that analysis to obtain the optimal solution to the INDS problem with the DC model.

We now explore the effectiveness of the valid inequalities presented in Section 3.2 on this case study. We will view the capacity reduction techniques discussed in Section 2.1 as ‘standard’ since they do not increase the number of constraints or variables in the formulation. The capacities of the arcs were reduced based on the lowest value obtained through each of the capacity reduction techniques presented in Section 2.1. In other words, the capacity of arc (i, j) was reduced to the lowest of either (i) the largest amount of flow placed on the arc in any feasible flow, (ii) the minimum cut separating node j from the set of demand nodes or (iii) the difference between the maximum flows in the network $(N, A \cup A')$ and the network (N, A) . Table 7 presents the solution

Table 6
The performance of the dispatching rule and IP formulation on the power infrastructure case study of lower Manhattan.

K	μ_t	w_i	Dispatching Rule		IP	
			Time (s)	Opt. gap overestimate (%)	Time (s)	Opt. gap (%)
1	Constant	Constant	8.12	16.47	21,600	25.64
1	Scaled	Constant	13.28	10.57	21,600	13.96
2	Constant	Constant	7.17	6.61	21,600	8.75
2	Scaled	Constant	7.20	2.42	21,600	3.01
3	Constant	Constant	6.15	4.58	21,600	5.23
3	Scaled	Constant	6.12	1.22	21,600	1.40

Table 7
Strength of the valid inequalities on the linear programming relaxation.

K	μ_t	w_i	IP relaxation		Capacity		Shortest path + capacity		Two step flow cover + capacity		β -conservation + capacity	
			Time (s)	Improvement (%)	Time (s)	Improvement (%)	Time (s)	Improvement (%)	Time (s)	Improvement (%)	Time (s)	Improvement (%)
1	Constant	Constant	32.79	299.00	1.53	131.64	6.93	3180.06	9.49	16984.54	10.72	
1	Scaled	Constant	42.65	338.95	0.29	134.66	1.56	2854.35	2.91	17154.55	3.40	
2	Constant	Constant	37.86	106.11	0.00	87.90	0.08	1198.44	2.02	8184.34	2.45	
2	Scaled	Constant	38.71	111.17	0.00	110.15	0.02	1712.72	0.43	13157.84	0.53	
3	Constant	Constant	41.87	46.56	0.00	50.07	0.00	797.58	0.18	5370.55	0.32	
3	Scaled	Constant	46.18	46.98	0.00	47.34	0.00	766.71	0.03	5827.52	0.06	

Table 8
Strength of the two step flow cover and β -conservation inequalities in combination with the capacity and shortest processing time path inequalities on the linear programming relaxation.

K	μ_t	w_i	IP relaxation		Two step flow cover + capacity + shortest path		β -conservation + capacity + shortest path	
			Time (s)	Improvement (%)	Time (s)	Improvement (%)	Time (s)	Improvement (%)
1	Constant	Constant	32.79	1281.76	11.06	9420.33	12.17	
1	Scaled	Constant	42.65	1218.40	3.22	6051.16	3.70	
2	Constant	Constant	37.86	1319.03	2.03	11527.58	2.52	
2	Scaled	Constant	38.71	1301.85	0.43	8761.33	0.54	
3	Constant	Constant	41.87	780.64	0.18	4743.86	0.32	
3	Scaled	Constant	46.18	715.18	0.03	5407.26	0.06	

Table 9
Computational results for the two step flow cover and β -conservation in combination with the capacity and shortest processing time path inequalities on the INDS problem.

K	μ_t	w_i	IP		Two step flow cover + capacity + shortest path		β -conservation + capacity + shortest path	
			Time (s)	Opt. gap (%)	Time (s)	Opt. gap (%)	Time (s)	Opt. gap (%)
1	Constant	Constant	21,600	25.64	21,600	16.47	21,600	18.19
1	Scaled	Constant	21,600	13.96	21,600	10.57	21,600	11.42
2	Constant	Constant	21,600	8.75	21,600	6.61	21,600	7.42
2	Scaled	Constant	21,600	3.01	21,600	2.42	21,600	2.79
3	Constant	Constant	21,600	5.23	21,600	4.58	21,600	4.61
3	Scaled	Constant	21,600	1.40	21,600	1.24	21,600	1.22

time and percentage improvement in the linear programming relaxation for each of the remaining inequalities plus the capacity inequalities. It is interesting to note that the valid inequalities are most effective for problems with $K=1$ and their effectiveness drops off quickly as you increase the number of work groups. The β -conservation constraints improve the quality of the relaxation the most, but this formulation requires significantly longer to solve than the formulation with the two step flow cover inequalities. The shortest processing time path inequalities do not add much in terms of computation time, so we have explored them in combination with the flow cover inequalities and β -conservation inequalities in Table 8.

The inclusion of the shortest processing time path inequalities does improve the effectiveness of both the formulations based on

the flow cover and β -conservation inequalities but the improvement is not as significant as to when the shortest processing time path inequalities are added to the original formulation of the problem. Once again, the results in Table 8 demonstrate the tradeoff between the quality of the two step flow cover and β -conservation inequalities and the solution time required to solve the relaxation. Therefore, we tested both these classes of valid inequalities on the integer programming formulation for the full six hour time limit. These results are presented in Table 9. These results indicate that the formulation based on the two step flow cover inequalities with the capacity and shortest processing time path inequalities provides the best bounds on the optimal solution to the integer program in our six hour time limit. This can be attributed to the fact that the relaxation of this formulation can be solved more

effectively than the one based on the β -conservation inequalities with only a small drop-off in terms of the quality of it.

We note that we have applied CPLEX 12.0 to the IP formulation with the capacity, shortest path, and β -conservation inequalities for the problem with $K = 3$ and constant μ_t on an eight-core computer. After a week of computational time, CPLEX 12.0 still had an optimality gap of 4.12% and had not identified a better solution than the one obtained by the dispatching rule. It is unreasonable to expect the managers of the infrastructure systems to have access to this type of computing resource and be willing to wait for over a week to determine their restoration plan, even in their scenario planning activities. Therefore, the dispatching rule is an important tool for the restoration of larger infrastructure systems.

6. Conclusions

This research has developed a novel integrated network design and scheduling problem that can be used to model the problem of restoring services provided by infrastructure systems after an extreme event disrupts them. The problem is general enough to be applicable to a variety of infrastructures, including the power, water, waste water, and emergency supply chain infrastructures. We have developed a novel heuristic dispatching rule for the INDS problem that focuses on selecting a set of arcs to process by examining the residual path optimality conditions from the area of network optimization. This dispatching rule has been shown to provide near-optimal solutions to small networks and the (current) best-known solution to large-scale networks for realistic case studies in seconds and can thus be used in real-time restoration activities. We have further developed an integer programming formulation of the INDS problem that was able to provide the optimal solutions to the smaller case studies quickly. However, the application of CPLEX 12.0 to this integer programming formulation for the larger case studies is not able to identify an optimal solution to the problem even when using significant computational resources. The valid inequalities developed in this research helped to strengthen the integer programming formulation but were not able to significantly help it solve the large case study. Further, CPLEX 12.0 was not even able to improve on the solution returned by the real-time dispatching rule for this large case study when it is applied to the integer programming formulation for six hours.

This research has applied the optimization models and algorithms developed for the INDS problem to several realistic case studies representing infrastructure systems in New Hanover County in North Carolina and lower Manhattan in New York City. The network models of the infrastructure systems in these case studies were created through careful collaborations with the managers of these systems. The disruptions/damage scenarios in the New Hanover County case studies were created to represent the effects of a hurricane whose eye passed just to the south of the county, implying that the strongest winds and flooding associated with the hurricane will affect the county. For the power infrastructures and waste water infrastructure, we examined the effects of the restoration efforts of the infrastructures when they are aligned with the priorities of the emergency manager of the county. We observed that, for these case studies, the restoration efforts of these infrastructures can be aligned with these priorities with no detrimental effects to the amount of demand met over the horizon of the restoration efforts. This is because the optimal restoration efforts formulated according to the priorities of the county was also optimal to the INDS problem representing the goals of meeting the maximum amount of demand over the horizon of the efforts. These case studies demonstrate that the dispatching rule can serve as a powerful decision support tool in real-time restoration activities. It will be important, in the future, to integrate the dispatching

rule and integer programming methods for the INDS problem into decision support systems for managers of infrastructure systems.

An important direction for future research related to the INDS problem is to determine customized large-scale integer programming methods that can be used to determine the optimal solution to large case studies. This is important since it can help to better understand the *true* performance of our dispatching rule rather than just bounds on the performance. This is further important since infrastructures are becoming larger and more complex due to the increase in demand for their services by society. Therefore, methods that can provide the optimal solution to large case studies and help to benchmark the performance of real-time solution methods to them is a critical area for future research. It may also be of interest to examine different methods to evaluate the performance of the network in order to capture the considerations of different infrastructure. For example, it may be more appropriate to model the operations of a road network by determining the shortest paths from a set of critical source nodes to a set of critical destination nodes.

Appendix A

A.1. Incorporating the DC model for power infrastructures

The DC model is a linear approximation of the behavior of a power system that is typically used in modeling the behavior of the power infrastructure system, especially the transmission network. For example, [Bienstock and Mattia \(2007\)](#) used the DC model in modeling problems related to mitigating grid blackout problems. We note that [Bienstock and Mattia \(2007\)](#) also provide a detailed discussion on the relationship between the linear DC model and the nonlinear AC model for a power system. The DC model includes decision variables at each node of the network that represent the *phase angle* of the node. The flow on arc (i,j) is then a function of the phase angles of nodes i and j along with the reactance of the arc (i,j) . The reactance, b_{ij} , of the arc is dependent on the length of it and the voltage levels. By defining θ_i for $i \in N$ as the phase angle of node i , the flow on arc (i,j) is determined through the equation:

$$b_{ij}x_{ij} = (\theta_i - \theta_j). \quad (16)$$

We note that both the phase angle variables *and* the arc flow variables are unrestricted in the DC model. A negative flow on arc (i,j) corresponds to power flowing from node j to node i . Therefore, it is necessary to include constraints that model Eq. (16) into the core INDS problem. The difficulty, however, is that this constraint should only be enforced at time t for the arcs in the network that have been completed prior to time t . In order to incorporate the DC model into (IP), we will define variables θ_{it} for $i \in N$ and $t = 1, \dots, T$ that represent the phase angle of node i in time period t . For all arcs $(i,j) \in A$, we replace constraints (5) with the constraints

$$b_{ij}x_{ijt} = (\theta_{it} - \theta_{jt}) \quad \text{for } (i,j) \in A, \quad t = 1, \dots, T, \quad (17)$$

$$-u_{ij} \leq x_{ijt} \leq u_{ij} \quad \text{for } (i,j) \in A, \quad t = 1, \dots, T. \quad (18)$$

These constraints enforce the DC flow model for arcs $(i,j) \in A$ and ensure that the flow on the arc does not exceed its capacity (thus preventing failure of these components). For arcs $(i,j) \in A'$, we can model the fact that we only wish to enforce DC flow calculations (16) when arc (i,j) appears in the network at time t by using ‘Big-M’ constraints. In particular, for arcs $(i,j) \in A'$, we replace constraints (6) with

$$b_{ij}x_{ijt} \leq (\theta_{it} - \theta_{jt}) + M(1 - \beta_{ijt}) \quad \text{for } (i,j) \in A', \quad t = 1, \dots, T, \quad (19)$$

$$b_{ij}x_{ijt} \geq (\theta_{it} - \theta_{jt}) - M(1 - \beta_{ijt}) \quad \text{for } (i,j) \in A', \quad t = 1, \dots, T \quad (20)$$

$$-u_{ij}\beta_{ijt} \leq x_{ijt} \leq u_{ij}\beta_{ijt} \quad \text{for } (i,j) \in A', \quad t = 1, \dots, T. \quad (21)$$

If $\beta_{ijt} = 0$, then these constraints force $x_{ijt} = 0$ while not imposing any restrictions on the relationship between the phase angles of nodes i and j due to the big M . If $\beta_{ijt} = 1$, then constraints (19) and (20) guarantee that the DC flow Eq. (16) is satisfied for arc (i, j) in time period t while constraint (21) guarantees that the capacity of the arc is not violated. We will refer to the core INDS problem where we have replaced constraints (5) and (6) with constraints (17)–(21) as the INDS problem with the DC Model.

References

- Ahuja, R., Magnanti, T., Orlin, J., 1993. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, Englewood Cliffs, New Jersey.
- Akker, J.v.d., Hoesel, C.van, Savelsbergh, M., 1999. A polyhedral approach to single-machine scheduling problems. *Mathematical Programming* 85 (3), 541–572.
- Ang, C., 2006. Optimized recovery of damaged electrical power grids. Unpublished master's thesis, Naval Postgraduate School.
- Bienstock, D., Mattia, S., 2007. Using mixed-integer programming to solve power grid blackout problems. *Discrete Optimization* 4 (1), 115–141.
- Cavdaroglu, B., Hammel, E., Mitchell, J., Sharkey, T., Wallace, W., forthcoming. Integrating restoration and scheduling decisions for disrupted interdependent infrastructure systems. *Annals of Operations Research*. <http://dx.doi.org/10.1007/s10479-011-0959-3>.
- Chang, S., McDaniels, T., Mikawoz, J., Peterson, K., 2007. Infrastructure failure interdependencies in extreme events: power outage consequences in the 1998 ice storm. *Natural Hazards* 41 (2), 337–358.
- Correa, J., Schulz, A., 2005. Single-machine scheduling with precedence constraints. *Mathematics of Operations Research* 30 (4), 1005–1021.
- Gu, Z., Nemhauser, G., Savelsbergh, M., 1999. Lifted flow cover inequalities for mixed 0–1 integer programs. *Mathematical Programming* 85 (3), 439–467.
- Guha, S., Moss, A., Naor, J., Schieber, B., 1999. Efficient recovery from power outage. In: Vitter, J., Larmore, L., Leighton, F. (Eds.), *Proceedings of the Symposium on Theory of Computing (STOC)*. Atlanta, GA, USA.
- Lee, E., Mitchell, J., Wallace, W., 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.
- Matisziw, T., Murray, A., Grubestic, T., 2010. Strategic network restoration. *Networks and Spatial Economics* 10 (3), 345–361.
- Mendonca, D., Wallace, W., 2006. Impacts of the 2001 World Trade Center attack on New York City critical infrastructures. *Journal of Infrastructure Systems* 12 (4), 260–270.
- Möhring, R., Schulz, A., Stork, F., Uetz, M., 2003. Solving project scheduling problems by minimum cut computations. *Management Science* 49 (3), 330–350.
- Nurre, S., Sharkey, T., 2010. Restoring infrastructure systems: an integrated network design and scheduling problem. In Johnson, A., Miller, J. (Eds.), *Proceedings of the 2010 Industrial Engineering Research Conference*. Cancun, Mexico.
- O'Rourke, T., 2007. Critical infrastructure, interdependencies, and resilience. *The Bridge: National Academy of Engineering* 37 (1), 22–29.
- Progress Energy, 2010. Storm Information. Retrieved from <<https://www.progress-energy.com/carolinas/home/storms-outages/index.page>>.
- Pinedo, M., 2008. *Scheduling: Theory, Algorithms, and Systems*. Springer, New York, New York.
- Rinaldi, S.M., Peerenboom, J.P., Kelly, T.K., 2001. Identifying, understanding, and analyzing critical infrastructure interdependencies. *IEEE Control Systems Magazine* 21 (6), 11–25.
- Savelsbergh, M., Uma, R., Wein, J., 2005. An experimental study of LP-based approximation algorithms for scheduling problems. *INFORMS Journal on Computing* 17 (1), 123–136.
- Schulz, A., 1996. *Polytopes and scheduling*. Unpublished doctoral dissertation, Department of Mathematics, Technische Universität Berlin, Germany.
- Smith, W., 1956. Various optimizers for single stage production. *Naval Research Logistics* 3 (1–2), 59–66.
- Sousa, J., Wolsey, L., 1992. A time indexed formulation of non-preemptive single machine scheduling problems. *Mathematical Programming* 54 (3), 353–367.
- Wallace, W., Mendonca, D., Lee, E., Mitchell, J., Chow, J., 2003. Managing disruptions to critical interdependent infrastructures in the context of the 2001 World Trade Center attack. *Beyond September*, 165–198.
- Waterer, H., Johnson, E., Nobili, P., Savelsbergh, M., 2002. The relation of time indexed formulations of single machine scheduling problems to the node packing problem. *Mathematical Programming* 93 (3), 477–494.
- Xu, N., Guikema, S., Davidson, R., Nozick, L., Cagnan, Z., Vaziri, K., 2007. Optimizing scheduling of post-earthquake electric power restoration tasks. *Earthquake Engineering and Structural Dynamics* 36 (2), 265–284.