Computers and Operations Research 90 (2018) 97–109

Contents lists available at ScienceDirect

Computers and Operations Research

journal homepage: www.elsevier.com/locate/cor

Stochastic maximum flow interdiction problems under heterogeneous risk preferences

Xiao Lei[a], Siqian Shen[b,*], Yongjia Song[c]

[a] University of International Business and Economics, Beijing, China
[b] Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI, 48109, USA
[c] Department of Statistical Sciences and Operations Research, Virginia Commonwealth University, Richmond, VA, USA

ARTICLE INFO

Article history:
Received 14 July 2016
Revised 9 June 2017
Accepted 5 September 2017
Available online 6 September 2017

Keywords:
Maximum flow interdiction
Stochastic programming
Conditional value-at-Risk (CVar)
Risk-averse optimization
Mixed-integer linear programming

ABSTRACT

We consider a generic maximum flow interdiction problem that involves a leader and a follower who take actions in sequence. Given an interdiction budget, the leader destroys a subset of arcs to minimize the follower's maximum flows from a source to a sink node. The effect from an interdiction action taken on each arc is random, following a given success rate of decreasing the arc's capacity to zero. The follower can add additional arc capacities for mitigating flow losses, after knowing the leader's interdiction plan but before realizing the uncertainty. We consider risk-neutral and risk-averse behaviors of the two players and investigate five bi-level/tri-level programming models for different risk-preference combinations. The models incorporate the expectation, left-tail, and right-tail Conditional Value-at-Risk (CVaR) as commonly used convex risk measures for evaluating random maximum flows in the leader's and follower's objectives. We reformulate each model as an equivalent mixed-integer linear program and test them on real-world network instances to demonstrate interactions between the leader and the follower under various risk-preference settings.

1. Introduction

The concept of a network is widely used to describe and optimize real-world problems ranging from supply chains, telecommunication, transportation, disaster relief, to public health. Systems in practice can often be represented by large-scale networks with nodes and arcs representing system entities and their interdependencies, respectively. The maximum flow problem is one of the most fundamental and well-studied network flow problems (Ahuja et al., 1993), where the objective is to generate the maximum amount of flows from a given source node (origin) to a sink node (destination).

Network interdiction models are based on Stackelberg games (Washburn and Wood, 1995) in networks, where one player (leader) acts first, and the other player (follower) will act after observing the leader's interdiction decision such as destroying nodes or arcs. The related studies originate from military logistics (see, e.g., Ghare et al., 1971) and have been generalized for various interdiction-related applications, e.g., interdicting smugglers in Morton et al. (2007). In stochastic network interdiction, the interdiction actions could fail with certain probability and/or network parameters such as arc capacity or arc cost could be random. We refer to Cormican et al. (1998); Woodruff (2002) and Hemmecke et al. (2003) for the representative work on stochastic network interdiction, to Wood (2010) and Dimitrov and Morton (2013) for comprehensive reviews of network interdiction problems, and to Lim and Smith (2007); Shen (2011) and Song and Shen (2016) for integer programming related methods for solving two-stage network interdiction models. Song and Shen (2016) also proposed risk-averse optimization and chance-constrained integer programming models but for shortest path interdiction.

We consider a stochastic maximum flow interdiction problem, where the leader destroys a subset of arcs to minimize the follower's maximum flows that are random due to uncertain arc capacities and interdiction effects. The follower has the ability to mitigate flow losses after observing the leader's interdiction plan and before the uncertainties are realized. We model the problem having three stages: (i) the leader interdicts a subset of arcs within a given budget; (ii) the follower adds additional capacities to a subset of selected arcs before knowing the realizations of uncertainties; and (iii) the follower solves for the maximum flows from the source to the sink given the decisions from the first two stages and realizations of uncertain parameters. To model the problem we use bi-level optimization where one problem is embedded (nested) within another (see, e.g., Bard, 1983) and tri-

* Corresponding author.
E-mail addresses: siqian@umich.edu, siqian.shen@gmail.com (S. Shen).

http://dx.doi.org/10.1016/j.cor.2017.09.004
0305-0548/

level optimization which involves three levels of nested problems. In our problem context the follower's arc recovery and flow maximization problems are nested within the leader's arc interdiction optimization problem. Depending on whether or not we can determine the follower's optimal decisions by solving a single-level model, we formulate the overall leader-follower game as a bi-level or a tri-level program. We refer the interested readers to Colson et al. (2005) and Bard (1998) for comprehensive surveys of bi-level programming models, solution algorithms, and applications, and to Vicente and Calamai (1994) for a bibliography review of bi-level and multi-level programming.

According to the leader's and the follower's risk preferences, we discuss five maximum flow interdiction cases, propose a bi-/tri-level programming model for each case, and reformulate them as single-level Mixed-Integer Linear Programming (MILP) models that can be directly solved by off-the-shelf solvers. Through computational analysis, we show that it is meaningful to consider follower's ability of recovering some arc capacities and to classify risk preferences of the players, since the two factors have significant impacts on the leader's optimal interdiction strategy. We also analyze the performance of the leader's interdiction solutions and effects of the follower's recovery solutions under various combinations of risk behavior of the two players. Via extensive computational experiments, we show that the follower being risk averse improves the follower's worst-case maximum flows with insignificant reductions of the average and the best-case maximum flows. One can choose the most suitable model based on specific goals and interdiction contexts in related applications.

The following papers are the most relevant to our work but the authors discuss different maximum flow interdiction cases, assumptions, and variants. Royset and Wood (2007) studied a bi-objective maximum flow interdiction problem and found Pareto-optimal solutions to minimize the interdiction cost and the follower's maximum flows. They employed Lagrangian relaxation and a branch-and-bound algorithm for identifying solutions. Altner et al. (2010) proposed valid inequalities and analyzed integrality gaps for optimizing a generic maximum flow interdiction problem with binary interdiction decisions. Akgün et al. (2011) and Lim and Smith (2007) considered problems of interdicting multi-terminal maximum flows. They employed integer programming models, partitioning algorithms, and heuristic approaches to optimize diverse problem variants.

In the literature on stochastic network interdiction, researchers have dominantly considered a risk-neutral leader who aims to minimize the follower's expected maximum flow. In such a setting, approaches based on Benders decomposition can be applied for solving the bi-level programs. The follower is usually assumed to only passively "accept" the leader's interdiction and make "wait-and-see" decisions to generate flows from the origin to the destination after observing the outcome of uncertain arc capacities and interdiction effects. Janjarassuk and Linderoth (2008) studied such a generic setting but with the follower being unable to recover arc capacities after observing the leader's interdiction plan. They only considered a risk-neutral leader who minimizes the expected value of the maximum flow generated by a wait-and-see follower. Their model is a special case of one of our models with a risk-neutral leader when the follower's recovery budget is zero.

In many real-world applications, the leader may not completely rely on the expected value of the follower's decision outcome. A risk-averse leader may focus on the follower's maximum flows in the extreme cases (e.g., the best-case or the worst-case flows that the follower is able to send). Also, the follower may not act passively and may possess ability to mitigate flow losses. If the follower can take recovery actions before knowing realizations of the uncertainties, then the leader also needs to take the follower's risk preference into consideration when deciding which arcs to interdict. To the best of our knowledge, we are the first to investigate combinations of both risk-neutral and risk-averse behavior of the leader and the follower in the context of stochastic maximum flow interdiction. A risk-averse follower seeks maximum flows that remain larger than a certain threshold in the worst-case percentile, while a risk-averse leader tries to make sure that the follower's maximum flows are not too significant in the best-case percentile. The leader may also try to minimize the maximum flows in the follower's worst-case percentile and in such a case we have a zero-sum game, which refers to the situation that each player's gain (or loss) of utility is exactly balanced by the losses (or gains) of the utility of the other player (see Washburn and Wood, 1995).

The remainder of the paper is organized as follows. In Section 2 we present a generic maximum flow interdiction model where the follower has the option of recovering arc capacities after knowing the leader's interdiction plan. We extend the model by allowing different risk preferences for the two players and formulate bi-level or tri-level programming models for five cases of representative combinations. In Section 3 we derive MILP reformulations of the five models. In Section 4 we conduct numerical tests of randomly generated instances based on real-world network structures. We provide thorough solution analysis and demonstrate out-of-sample test results. In Section 5 we summarize the paper and propose future research directions.

2. Maximum flow interdiction under diverse risk preferences

We consider maximum flow interdiction in a directed network $G = (\mathcal{V}, \mathcal{A})$ with $\mathcal{V}$ representing a set of nodes and $\mathcal{A} \subseteq \mathcal{V} \times \mathcal{V}$ representing a set of arcs, $\{1, \ldots, |\mathcal{A}|\}$. Each arc $(i, j) \in \mathcal{A}$ has a nonnegative capacity u_{ij} and we designate nodes $s,\ t \in \mathcal{V}$ as the origin and destination in the follower's maximum flow problem, respectively.

We define a binary variable $x_{ij} \in \{0,\ 1\}$ for each arc $(i, j) \in \mathcal{A}$, such that $x_{ij} = 1$ if the leader destroys arc (i, j) and 0 otherwise. Vector $x \in \{0, 1\}^{|\mathcal{A}|}$ denotes the leader's interdiction decision. The interdiction cost is given by parameter $r \in \mathbb{R}_+^{|\mathcal{A}|}$ and $r_0 \in \mathbb{R}_+$ is the total budget for interdicting arcs. The leader's feasible region is defined as

$$X := \{x \in \{0, 1\}^{|\mathcal{A}|} \mid r^\top x \le r_0\}. \tag{1}$$

The leader aims to minimize the follower's maximum flows from s to t by optimizing $x \in X$. In a stochastic setting, we consider random arc capacities represented by a vector $\tilde{u} \in \mathbb{R}_+^{|\mathcal{A}|}$. A successfully interdicted arc $(i, j) \in \mathcal{A}$ will have zero capacity. However the leader's interdiction on each arc may not be successful and a random parameter vector $\tilde{\xi} \in \{0, 1\}^{|\mathcal{A}|}$ indicates whether or not an interdiction is successful on each arc. We employ the Sample Average Approximation (SAA) method, which is a common approach for solving stochastic programs with random parameters by reducing the scenario set to a manageable size (see, e.g., Birge and Louveaux, 2011, Shapiro et al., 2009). In this approach, we use Monte Carlo sampling to generate a finite set of scenarios, $\mathcal{K} = \{1, \ldots, |\mathcal{K}|\}$, following some given distribution of the uncertain parameter $(\tilde{u}, \tilde{\xi})$. We denote (u^k, ξ^k) as the realization of $(\tilde{u}, \tilde{\xi})$ in each sample $k \in \mathcal{K}$, where $u^k = (u_{ij}^k, (i, j) \in \mathcal{A})^\top \in \mathbb{R}_+^{|\mathcal{A}|}$ and $\xi^k = (\xi_{ij}^k, (i, j) \in \mathcal{A})^\top \in \{0, 1\}^{|\mathcal{A}|}$. Let p_k be the probability associated with scenario $k \in \mathcal{K}$ and we have $\sum_{k \in \mathcal{K}} p_k = 1$.

After the leader's interdiction, the follower adds additional capacities to a subset of arcs subject to a recovery budget. We define a continuous variable $h_{ij} \ge 0$ as the added capacity to arc (i, j) and denote h_0 as the total budget for allocating capacities. We use $h \in \mathbb{R}_+^{|\mathcal{A}|}$ to denote the follower's recovery decision, and the fol-