LA-UR- 08-4427

Approved for public release; distribution is unlimited.

Title: OPTIMAL INTERDICTION OF REACTIVE MARKOVIAN EVADERS

Author(s): ALEXANDER GUTFRAIND
ARIC HAGBERG
FENG PAN

Intended for: PROCEEDINGS:
SIAM /ACM CONFERENCE

Los Alamos
NATIONAL LABORATORY
EST. 1943

OPTIMAL INTERDICTION OF UNREACTIVE MARKOVIAN EVADERS

ALEXANDER GUTFRAIND, AG362@CORNELL.EDU ARIC A. HAGBERG, HAGBERG@LANL.GOV
FENG PAN, FPAN@LANL.GOV

ABSTRACT. The network interdiction problem [4] arises in a wide variety of areas including military logistics, infectious disease control, and counter-terrorism [13]. In the classical formulation one is given a weighted network $G(N,E)$ and the task is to find b nodes (or edges) whose removal would maximally increase the least-cost path from a source node s to a target node t. In practical applications, G represents a transportation or activity network; node/edge removal is done by an agent, the "interdictor" against another agent, the "evader" who wants to traverse G from s to t along the least-cost route. Our work is motivated by cases in which both agents have bounded rationality: *e.g.* when the authorities set up road blocks to catch bank robbers, neither party can plot its actions with full information about the other. We introduce a novel model of network interdiction in which the motion of (possibly) several evaders is described by a Markov process on G.We further suppose that the evaders do not respond to interdiction decisions because of time, knowledge or computational constraints. We prove that this interdiction problem is NP-hard, like the classical formulation[1, 2], but unlike the classical problem the objective function is submodular. This implies that the solution could be approximated within $1-1/e$ using a greedy algorithm. Exploiting submodularity again, we demonstrate that a "priority" (or "lazy") evaluation algorithm can improve performance by orders of magnitude. Taken together, the results bring closer realistic solutions to the interdiction problem on global-scale networks.

1. INTRODUCTION

Mathematical modeling of network interdiction was introduced originally in the study of military supply chains and interdiction of transportation networks [10, 5]. But in general, the network interdiction problem applies to wide variety of areas including ballistic missile defense, infectious disease control, and disruption of terrorist networks. Recent interest in the problem has been revived due to the threat of smuggling of nuclear materials [13]. In this context interdiction of edges corresponds to the placement of special radiation-sensitive detectors along the selected transportation links.

Network interdiction problems have two opposing actors: one or more network "evaders" (smugglers, etc.) and an "interdictor" (leader, border agent, etc.) Each evader attempts to minimize some objective function in the network, *e.g.* the probability of capture while traveling from network location s to location t, while the interdictor attempts to limit success by removing network nodes or edges. Most often the interdictor has limited resources and can thus only remove a very small fraction of the nodes or edges. The case where the interdictor can choose at most k edges to maximize the shortest path is known as the k most vital arcs problem [4] and even when there is just one evader it has been shown to be NP-hard [1, 2] and hard to approximate [3].

This classical shortest-path formulation is not suitable for some interesting interdiction scenarios for a number of reasons. First, in many practical problems there is a thick fog of uncertainty about the underling network properties, such as the cost to the evaders to traverse a network arc in terms of resource consumption or probability of detection. Second, there are also mismatches in the cost and risk computations between the interdictor and the evaders, as well as between different evaders, and all agents have an interest in hiding their actions. Therefore, the interdictor has at best only probabilistic information about the evaders. Thirdly, on the evader side, classical evader models make maximal assumptions about the capacity of the evaders to be informed about the interdictor's strategy, namely, the choice of interdiction set. Practical evaders likely fall far short of the maximum. Therefore, this paper considers the other limit, namely, the case when the evaders do not respond to interdictor's decisions. Section 2 defines this new interdiction problem and section 3 shows that it is NP-hard to solve. Then section 4 introduces an approximate interdiction algorithm that runs with fast polynomial time and has a provable approximation bound.

2. Problem Formulation

The insufficient realism of the classical model represents an opportunity for the development of better interdiction models which would have the additional advantage of being solvable on much larger networks and this is the project of this paper. The formulation can be motivated by the following interdiction situation. Suppose bank robbers want to escape from the bank at node s to their safe haven at node t. The authorities are able to position roadblocks at a few of the roads on the network between s and t. The robbers might not be aware of the interdiction efforts, or believe that they will be outrun the authorities. They certainly do not have the time or the computational resources to identify the max-min solution to the shortest path interdiction problem.

Similar instances of evasion are found for example, when the interdiction is able to remove the edges/nodes clandestinely (*e.g.* place hidden electronic detectors), or if the evader has bounded rationality such as in animal migration. Indeed it may even have no intelligence of any kind and represent a process such as internet packet traffic that the interdictor wants to monitor. Therefore, our fundamental assumption is that the evader does not respond to interdiction decisions. This transforms the interdiction problem from the problem of increasing the evader's cost or distance of travel, as in the classical formulation, into a problem of directly capturing the evader as explicitly defined below. Additionally, the objective function acquires the computationally useful property of submodularity discussed later.

Evaders. In examples discussed above, much of the difficulty in interdiction stems from the unpredictability of evader motion. Suppose, then that each evader can be represented as a Markov process on the network, with transition probabilities M_{ij} between every pair of adjacent nodes i and j. The stochasticity of this process represents, first, the evader's intentional unpredictable movement, and second, the evader's limited information about the network topology and the risks/costs along alternative paths. A simple model of this behavior was recently developed in Ref.*et al.* [7] (see also [11]). Briefly, the model assumes that the probability that an evader at node i would traverse $i \to j$ depends on the probability q_j of (successful) evasion on the shortest path through this edge to the target. Since the higher the probability of evasion, the higher the probability it would be taken, the paper supposes that $M_{ij} \propto \left(\frac{q_j}{q_*}\right)^{\lambda}$ where q_* is the probability of evasion if the shortest path from node i to the target t is followed and $\lambda \geq 0$ is a parameter. Upon reaching the target the evader is removed from the graph. An advantage of the model is that it is possible to significantly speed up the computation of evader motion by assuming that the evader does not backtrack. For a graphical illustration of evader motion, see Fig.2.1. In general, the current problem does not depend on any particular stochastic model, and indeed it need not be Markovian, even thought it will be assumed to be to simplify calculations. For the Markovian case, it is sufficient that M has an absorbing state at t [6].

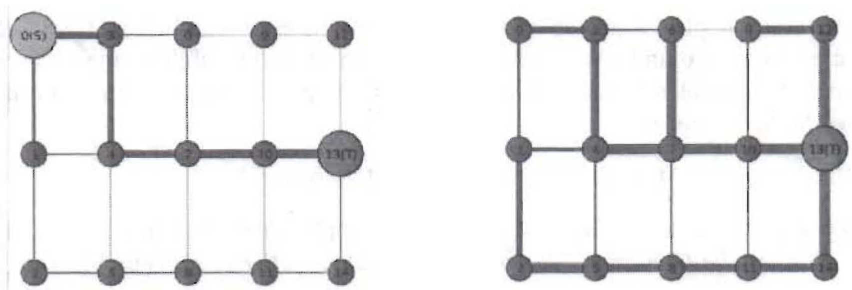

(a) (b)

FIGURE 2.1. Evader motion of an evader on a 3x5 grid graph, starting at node $0(S)$ and terminating at node $13(T)$ based on Ref.[7] with $\lambda = 4$ and no backtracking. (a) probabilities that an evader would pass any one of the edges, as indicated by edge widths. (b) the M matrix: edge width gives the probability that an evader would transit the edge from either of the edge's ends. In both (a) and (b) no edges were interdicted and the asymmetries in probabilities are due to different risks on the paths.

Due to uncertainty about the source site of the evader, node $i \in N$ has probability a_i of being the source site ($\sum_{i \in N} a_i = 1$). Specifically, the expected number of times the evader reaches node i is given by [6]:

$$b_i = \left[a \cdot (I - M)^{-1}\right]_i \tag{2.1}$$

and the expected number of times the evader reaches edge (i, j) is given by:

$$b_i \cdot M_{ij} \tag{2.2}$$

In general, multiple evaders may traverse the network, where evader e_i has likelihood (weight) w^{e_i} and is described by a possibly distinctive source distribution, transition matrix and target node: $a^{e_i}, M^{e_i}, t^{e_i}$. This generalization makes it possible to represent any joint probability distribution $f(s,t)$ of source-target pairs, where each evader is a slice of f at a specific value of t. In this high level view, the evaders collectively represent a stochastic process that connects pairs of nodes on the network, and with which the interdictor attempts to interfer. This has practical applications to problems of monitoring traffic between any set of nodes when there is a limit on the number of "sensors". The underlying network could be *e.g.* transportation, internet, or water supply.

Interdictor. The interdictor in this formulation is similar to the classical one in possessing complete knowledge about the network and all evader parameters parameters a and M. Interdiction is represented by the variables r_{ij}: $r_{ij} = 1$ if edge (i, j) is interdicted, and zero elsewhere (the formulation is similar for node interdiction). Next let d_{ij} be the efficiency of the interdiction at (i, j). Namely, if edge (i, j) is interdicted and then traversed, then the evader is removed from the graph with probability d_{ij}. An equivalent conceptualization that preserves the stochastic properties of M is to say that instead of reaching the next node and eventually the target, the evader is redirected to an absorbing "jail node". In monitoring problems, this redirection does not require the evader to be physically removed from the network because it merely represents a way of talling up traffic that was inspected before reaching the target.

The feasible interdiction strategies R are all subsets $r \subset E$ such that the cost, $c(r)$ of r is at most b. In the simplest formulation $c(r)$ is just the number of interdicted edges. The interdiction problem is to find a set $r \in R$ so as to minimize the probability of the evader reaching target t. Equivalently, the task is to maximize the probability $J(a,M,r,d)$ of the evader going to the jail node:

$$J(a,M,r,d) = 1 - \left[a(I-(M-M\odot r\odot d))^{-1}\right] \tag{2.3}$$

where the symbol "$\odot$" is the element-wise Hadamard multiplication (derived from Eqn.2.1). This could be called the *Unreactive Markovian Evader Interdiction (UMEI)* problem:

$$\max_{r\in R} J(a,M,r,d)$$

In the case of multiple evaders, the objective becomes a weighted sum of Eqn.2.3: $J = \sum_{e_i} w^{e_i} J^{e_i}$, where J^{e_i} is the probability of evader e_i going to jail. Computation of the objective function can be achieved in $\frac{2}{3}|N|^3$ time for each evader, where $|N|$ is the number of nodes, because it is dominated by the cost of Gaussian elimination. If M has special structure, significant speedups can be achieved as in [7].

3. Complexity of Solutions

This section proves a number of technical results about the problem - the equivalence in complexity of node and edge interdiction, and the NP-hardness of the optimization problem.

Lemma. *Edge interdiction with any d_{ij} is polynomially equivalent to node interdiction in complexity.*

Proof. To reduce edge interdiction to node interdiction, replace graph G by G' where G' is constructed by splitting the edges. That is, by adding to the middle of each edge (i,j) in G a node v_{ij} and setting $d_{v_{ij}} = \begin{cases} d_{ij} & (i,j)\in E \\ 0 & (i,j)\notin E \end{cases}$, where E are the edges of G and d_{ij} is the interdiction efficiency of (i,j) in E. Conversely, to reduce node interdiction to edge interdiction, construct from G another graph G' by replacing each node v with efficiency d_v by nodes v^{in}, v^{out} and joining them with an edge such that $d_{ij} = \begin{cases} d_v & (i,j)\in N \\ 0 & (i,j)\notin N \end{cases}$, where N are the nodes of G. Next, change the transition matrix M of each evader such that all transitions into v now move into v^{in} while all departures from v now occur from v^{out}, and $M_{v^{in}v^{out}} = 1$. □

The hardness of the UMEI problem stems from the difficulty interdicting edges while avoiding redundancy: it is wasteful to interdict edges on the same evader path. This resembles the Set Cover problem [8] because including an element in two sets is redundant in a similar way, and this is how reduction is proved.

Theorem. *The UMEI problem is NP-hard even if d_{ij} =constant.*

Proof. First, note that the problem is polynomially equivalent to budget minimization problem:

$$\min_{r\in R} c(r) \text{ s.t. } J(r) \geq \rho$$

where $\rho \in [0,1]$ is the minimum allowed interdiction probability. In turn, there is a decision problem which polynomially reduces to it:

Can we find r such that $c(r) \leq b$ and $J(r) \geq \rho$?

The lemma implies that it would be sufficient to prove the hardness of node interdiction. The core of the proof is the demonstration that the NP-complete Set Cover problem polynomially reduces to the decision problem. In the problem one is given elements $x_i \in X$, and a collection of subsets $C = \{S_j | S_j \in 2^X\}$ and the task is to decide whether there exist b subsets which would contain all of X. For the reduction, construct a graph G with nodes labeled S_j for each $S_j \in C$. For each element x_i introduce an evader x_i. If $x_i \in C_i = \{S_j \in C | x_i \in S_j\}$, then x_i starts at one $S \in C_i$ and transitions once through each of the nodes in C_i. It is clear that evader x_i would be jailed iff at least one of the nodes in its path, that it, in C_i is interdicted. Hence, the interdiction decision problem on H with $\rho = 1$ (all go to jail) is equivalent finding the set cover. □

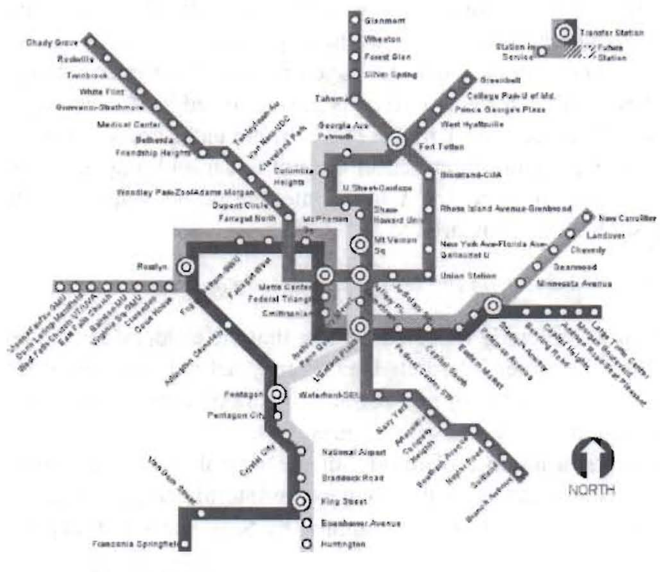

FIGURE 3.1. Washington subway map. Interdiction of any station would catch the evader taking the lines that pass through it. Image ©Colin M.L. Burnett, licensed under GFDL.

The idea in the reduction below is to produce a transportation graph similar to the subway system of Washington, Fig.3.1. For each evader there is a subway line which the evader rides from end to end. "Transfer stations", that is, stations serviced by multiple lines are akin to non-singleton subsets whose interdiction would catch all the evaders taking any of the connecting lines. The set cover problem is exactly problem of finding b stations of any kind which would interdict all the lines.

In general, the cost c_i of interdicting a node or edge i may depend on i. If so, the hardness of the problem is even easier to prove through reduction to the knapsack problem: Given values v_i and costs c_i, construct the UMEI problem on a star graph for a suitable M: Make the central node the target and choose a, M so that probability flows of the evader through edge i give v_i (up to normalization).

4. An Efficient Interdiction Algorithm

The UMEI problem can be efficiently approximated using a greedy algorithm by exploiting submodularity. We prove this property for UMEI, construct a greedy algorithm and prove its performance. We then show that the solution speed can be further improved by exploiting submodularity once again.

A function is called submodular if the rate of increase decreases monotonically, which is akin to concavity. Formally:

Definition. A function $f : S \to \mathbb{R}$ where S is a subset of some space R is *submodular* [12] if for any subsets $S_1 \subseteq S_2 \subset R$ and any $X \subset R$ it satisfies

$$f(S_1 \cup X) - f(S_1) \geq f(S_2 \cup X) - f(S_2)$$

Consider now how the weighted probability $J(r)$ of the evaders being sent to jail changes as the number of interdicted edges increases. This probability is a weighted sum of the number of paths from the source nodes to the target nodes that cross an interdicted edge. Consider now two cases: in case one the edges S_1 have been interdicted and in case two $S_2 \supseteq S_1$ edges have been interdicted. If X is any set of edges (some possibly interdicted already), then with S_1 at least as many but maybe more paths cross X as compared with S_2. This is because with a smaller interdiction set, not fewer and maybe more paths reach X. This means that the gain from adding X to the interdiction set is at least as large but maybe larger when the set is S_1 as compared to S_2. Hence:

Lemma. *$J(r)$ is submodular on the set of interdicted edges, r.*

Note that the proof crucially relies on the fact that the evader does not react to interdiction. If it did, then the larger interdicted set S_2 may actually redirect the evader towards paths that cross X, increasing their number and likelihood more so than S_1. Incidentally, f has a non-decreasing property, but will not need this.

Submodularity has a number of important theoretical and algorithmic consequences. Suppose, as is likely in practice, that the edges are interdicted in steps $l = 1, 2, 3, \ldots$ such that the interdiction set S_l at step l is contained by S_{l+1}, the set at step $l+1$. Moreover, suppose at each step l, the interdiction set S_l is grown by adding the one edge that gives the greatest increase in J giving in the following algorithm:

Algorithm. *Greedy construction of the interdiction S with budget b*

$S \leftarrow \emptyset$
while $b > 0$:
For each element $E_i \in G \setminus S$: *Compute* $\Delta(E_i) := J(S \cup \{E_i\}) - J(S)$
$S \leftarrow S \cup \{\text{argmax}_{E_i} \Delta(E_i)\}$, *resolving ties arbitrarily.*
$b \leftarrow b - 1$

The computational time is $O(b|N|^3|E|)$ for each evader, which is polynomial, but it could be improved further: the $\Delta(E_i)$ for the initial step only could be computed for all edges simultaneously by determining the probability flow f_i through each edge and multiplying by d_i, the interdiction efficiency.

Solution Quality. The quality of the approximation can be bounded as a fraction of the optimal solution by exploiting the submodularity property. Consider the following: if S_b^* is the optimal interdiction set within budget b, then the first edge E_1 found by the greedy algorithm provides a gain that is not less than the average gain for all the edges in S_b^*, that is, $\Delta(E_1) \geq \frac{J(S_b^*)}{|b|}$. A similar idea for all subsequent edges could be use to show that the solution S_b^g with a greedy algorithm satisfies

$$J(S_b^g) \geq \left(1 - \frac{1}{e}\right) \cdot J(S_b^*)$$

where e is Euler's constant [12]. Hence, the greedy algorithm achieves at least 63% of the optimal solution. Intuitively, submodularity implies that most of the gains in interdiction would be attained by edges interdicted in early steps and thus ever a myopic greedy algorithm is able to capture a large fraction of this gain.

Exploiting Submodularity with Priority Evaluation. In addition to its theoretical utility, submodularity can be exploited to improve performance as follows. The basic greedy algorithm recomputes the $\Delta_l(E_i)$ for each edge $E_i \in G \setminus S_l$ at each step l. However, submodularity implies that the gain $\Delta_l(E_i)$ from adding any edge E_i to the interdiction set at step l is not greater than the gain $\Delta_k(E_i)$ computed at any earlier step k ($< l$). Therefore, if for some edge E_j, $\Delta_l(E_j) > \Delta_{k_i}(E_i)$ for all E_i and any past step $k_i \leq l$, then E_j is the optimal edge at step l - there is no need for further computation (as was suggested in a different context by Ref.[9]). As a result, on average it would not be necessary to compute $\Delta(E_i)$ for all edges $E_i \in G \setminus S$ at every iteration. Rather, the computation should prioritize the edges in descending order of Δ. This "priority" or "lazy" evaluation algorithm is easily implemented with a priority queue which stores the gain $\Delta(E_i)$ for each edge, as well as the step k at which it was last calculated (the latter informing whether it needs to be updated.)[1]

The performance gain from this improvement can be dramatic: in many computational experiments, the second best edge from the previous step was the best in the current step. Frequently, only a small fraction of the edges had to be recomputed at each iteration. However, in the worst case the gains would all need to be recomputed, so the speedup would vanish. In order to better gauge the improvement in performance, the algorithm was run on a highly synthetic hard interdiction problem, as follows. The transportation graph was a $10x10$ unweighted grid with the boundary nodes connected to make the graph periodic, and then 10 random ties were added to simulate jumps - a total of 110 edges. There was a single highly-randomizing Markov evader (based on the model of Ref.[7], with $\lambda = 0$) with source nodes uniformly distributed over the graph. The interdictor had a budget of 10 edges, where each interdiction cuts the flow in half ($d_{ij} = 0.5$). It was found that the solution required just 33 evaluations of Δ in the priority evaluation scheme, instead of the $991 = 1 + 9 \cdot 110$ in the simple algorithm: a speedup by a factor of 30 (*cf.* [9]). When the budget was reduced to 5, the speedup was a factor of 25.

5. Further Work

The submodularity property of the UMEI problem developed above provides a rich source for algorithmic improvement. In particular, there is room for and value in more efficient approximation schemes. Simultaneously, it would be interesting to classify the UMEI problem into a known approximability class and to prove its hardness even when there is just one evader. Of more practical interest would be to investigate various trade-offs in

[1]The source code is available from the authors.

the interdiction problem, such as the trade-off between quality and quantity of interdiction devices, as well as to quantify the loss of accuracy in problems where the evader is able to respond to the interdiction decisions, even if suboptimally.

Acknowledgement. AG would like to thank Jon Kleinberg for inspiring lectures and David Shmoys for a helpful discussion. Part of this work was funded by the Department of Energy at Los Alamos National Laboratory under contract DE-AC52-06NA25396 through the Laboratory Directed Research and Development Program.

REFERENCES

[1] Michael O. Ball, Bruce L. Golden, and Rakesh V. Vohra. Finding the most vital arcs in a network. *Oper Res Lett*, 8(2):73–76, 1989.

[2] A. Bar-Noy, S. Khuller, and B. Schieber. The complexity of finding most vital arcs and nodes. Technical report, University of Maryland, College Park, MD, USA, 1995.

[3] E. Boros, K. Borys, and V. Gurevich. Inapproximability bounds for shortest-path network intediction problems. Technical report, Rutgers University, Piscataway, NJ, USA, 2006.

[4] H. W. Corley and D. Y. Sha. Most vital links and nodes in weighted networks. *Operations Research Letters*, 1(4):157 – 160, Sep 1982.

[5] P. M. Ghare, D. C. Montgomery, and W. C. Turner. Optimal interdiction policy for a flow network. *Naval Research Logistics Quarterly*, 18(1):37, 1971.

[6] Charles M. Grinstead and J. Laurie Snell. *Introduction to Probability*. AMS, USA, second revised edition, Jul 1997.

[7] Alexander Gutfraind, Aric A. Hagberg, David Izraelevitz, and Feng Pan. Interdicting a Markovian evader. *arxiv.org*, pages 1–1, 2008.

[8] Richard M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. New York: Plenum, 1972.

[9] Jure Leskovec, Andreas Krause, Carlos Guestrin, Christos Faloutsos, Jeanne VanBriesen, and Natalie Glance. Cost-effective outbreak detection in networks. In *KDD '07: Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 420–429, New York, NY, USA, 2007. ACM.

[10] A. W. McMasters and T. M. Mustin. Optimal interdiction of a supply network. *Naval Research Logistics Quarterly*, 17(3):261, 1970.

[11] D. P. Morton, F. Pan, and K. J. Saeger. Models for nuclear smuggling interdiction. *IIE Transactions*, 39(1):3–14, 2007.

[12] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of the approximations for maximizing submodular set functions-I. *Mathematical Programming*, 14:265–294, 1978.

[13] F. Pan, W. Charlton, and D. P. Morton. Interdicting smuggled nuclear material. In D.L. Woodruff, editor, *Network Interdiction and Stochastic Integer Programming*, pages 1–19. Kluwer Academic Publishers, Boston, 2003.