NPS-OR-23-004

NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

DECEPTION AND RISK-AWARE DYNAMIC ROUTING

by

Jefferson Huang
Ruriko Yoshida

21 October 2023

Approved for public release. Distribution is unlimited.

Prepared for: Naval Surface Warfare Center, Crane Division.
This research is supported by funding from the Naval Postgraduate School, Naval Research Program (PE 0605853N/2098).
NRP Project ID: NPS-23-N059-A

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
10/21/2023	Technical report	10/23/2022	10/21/2023

4. TITLE AND SUBTITLE

Deception and Risk-Aware Dynamic Routing

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
		0605853N/2098
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**
NPS-23-N059-A		

6. AUTHOR(S)

Jefferson Huang, Ruriko Yoshida

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Naval Postgraduate School 1 University Circle Monterey, CA 93943-5000	NPS-OR-23-004

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
Naval Postgraduate School, Naval Research Program; Naval Surface Warfare Center Crane Division	NRP; NSWC Crane	NPS-23-N059-A

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release, distribution is unlimited.

13. SUPPLEMENTARY NOTES

The views expressed in this document are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

14. ABSTRACT

This report addresses the question of how routing algorithms should be modified when they are applied to routing agents, such as unmanned vehicles, in contested environments. Our primary motivation is prior work at the Naval Postgraduate School that showed how optimization-based routing that does not account for the presence of an observing adversary can result in predictable paths from which operational intent can be inferred. We propose the use of a randomized routing strategy based on the solution of a two-player game and evaluate its effectiveness on data derived from a multi-thread experiment. We also propose methods for dynamic routing in contested environments. Our work highlights the trade-off that needs to be made between efficiency and predictability in route planning and can potentially inform the development of new adversary-aware routing algorithms for unmanned systems.

15. SUBJECT TERMS

deception, routing, dynamic, contested

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
U	U	U	UU	39

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Jefferson Huang	(831) 656-2605

THIS PAGE INTENTIONALLY LEFT BLANK

NAVAL POSTGRADUATE SCHOOL
Monterey, California 93943-5000

Ann E. Rondeau
President

Scott Gartner
Provost

Further distribution of all or part of this report is authorized.

This report was prepared by:

HUANG.JEFFERSON.1547945588
Digitally signed by HUANG.JEFFERSON.1547945588
Date: 2023.11.28 09:46:54 -08'00'

Jefferson Huang
Assistant Professor

Digitally signed by Ruriko Yoshida
Date: 2023.11.28 09:51:14 -08'00'

Ruriko Yoshida
Professor

Reviewed by:

CARLYLE.WILLIAM.MATTHEW.1256369440
Digitally signed by CARLYLE.WILLIAM.MATTHEW.1256369440
DN: c=US, o=U.S. Government, ou=DoD, ou=PKI, ou=USN, cn=CARLYLE.WILLIAM.MATTHEW.1256369440
Date: 2023.11.28 10:27:02 -08'00'

W. Matthew Carlyle, Chairman
Operations Research Department

Released by:

SMITH.KEVIN.B.1230026366
Digitally signed by SMITH.KEVIN.B.1230026366
Date: 2023.12.02 06:46:44 -08'00'

Kevin B. Smith
Vice Provost for Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

This report addresses the question of how routing algorithms should be modified when they are applied to routing agents, such as unmanned vehicles, in contested environments. Our primary motivation is prior work at the Naval Postgraduate School that showed how optimization-based routing that does not account for the presence of an observing adversary can result in predictable paths from which operational intent can be inferred. We propose the use of a randomized routing strategy based on the solution of a two-player game and evaluate its effectiveness on data derived from a multi-thread experiment. We also propose methods for dynamic routing in contested environments. Our work highlights the trade-off that needs to be made between efficiency and predictability in route planning and can potentially inform the development of new adversary-aware routing algorithms for unmanned systems.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

DDG	guided missile destroyer
LP	linear programming problem
MTX	multi-thread experiment
NPS	Naval Postgraduate School
NSW	Naval Special Warfare
OR	Operations Research
UAV	unmanned aerial vehicle
USV	unmanned surface vehicle
UxV	unmanned vehicle
USN	U.S. Navy
VRP	vehicle routing problem

THIS PAGE INTENTIONALLY LEFT BLANK

Executive Summary

Background

The purpose of this study was to explore ways in which deception, with the aim of obscuring the operational intent of a Blue agent, could be incorporated into optimization-based algorithms for routing the agent in a potentially contested environment. This is one of the active areas of research being conducted at Naval Surface Warfare Center Crane Division, which is our topic sponsor's organization. One of the assumptions was that a Red agent could observe the movements of the Blue agent and use these observations to attempt to predict Blue's destination. Specifically, Red was assumed to be able to use the prediction algorithm developed in (Wigington 2021), where it was shown using data from a multi-thread experiment (MTX) on San Clemente Island conducted by Naval Postgraduate School (NPS) researchers, that this prediction algorithm could accurately predict the locations of supported ground forces based on the locations of supporting unmanned vehicles (UxVs). Our starting point was a game-theoretic formulation of the interaction between the Blue agent seeking reach a goal before Red is able to guess it, and the Red agent seeking to guess Blue's goal location, that was proposed in Tsitsiklis and Xu (2018). The structure of the game suggested that a particular class of strategy for Blue would be effective. This class of strategies involves the introduction of randomness into Blue's movements, in a way designed to make it appear to Red that every possible destination of Blue is roughly equally likely. In other words, it approximately maximizes Red's uncertainty about Blue's destination. The policy includes a tuning parameter that reflects how much randomness Blue is willing to introduce into their route; in general, Blue must deal with a trade-off between efficiency of the route, and predictability. We implemented and tested this routing strategy for a single UxV, using data derived from the MTX, to assess the degree to which it can defeat Red's prediction capabilities. This work was part of the thesis research of LT Kyle Plunkett, U.S. Navy (USN), who graduated with a Master's degree in Operations Research (OR) in March 2023 (Plunkett 2023). We also studied an alternative formulation of the problem of routing in a contested environment, based on iteratively solving an online shortest path problem that includes a cost term reflecting the current threat landscape. This is part of ongoing thesis research by CPT Yan-ru Lin of the Republic of China Marine Corps, an NPS OR Master's student who

joined the project in March 2023 and is slated to graduate in December 2023.

Findings and Conclusions

Our primary finding is that the randomized routing strategy suggested by the two-player game proposed in Tsitsiklis and Xu (2018) can significantly degrade the predictive capability of the prediction algorithm developed in Wigington (2021). Specifically, it was able to increase the average route prediction error from a maximum of 75 meters, to roughly an average of 300 meters Plunkett (2023). We also assessed the destination prediction error, which was roughly 368 meters on average (Plunkett 2023). We also observed that it was sufficient for one of the UxVs being observed to act as a decoy, by following a randomized path, while the others followed more efficient paths, to significantly degrade the prediction algorithm's performance. This finding provides empirical support for the importance of using deception (e.g., via decoys), and can inform the development of routing algorithms that take an observing adversary into account. In addition to these findings, we also assessed the method based on solving online shortest path problems using a scenario set in the Indo-Pacific, where the goal is to supply several Taiwanese ports in the context of an evolving threat landscape. The application to this notional scenario provides a proof-of-concept of the proposed approach and indicates its potential applicability to routing problems in the Indo-Pacific theater.

Recommendations for Further Research

There are many promising directions for further research. One direction is to try to improve the prediction model in Wigington (2021), to model a more sophisticated adversary and to determine whether the randomized routing strategies evaluated in this research (Plunkett 2023) still effective against it. If not, further research should include studying ways in which the deception mechanism can be improved. In general, the goal would be to continue to better understand how to balance deception with efficiency, and to address the reality that near-peer adversaries will likely have the capability to take advantage of observations of how our autonomous systems behave.

CHAPTER 1:
Introduction & Background

The 2022 U.S. National Defense Strategy notes that the rapid development of new technologies, such as those based on advances in artificial intelligence or autonomous systems, has the potential to significantly change and complicate the nature of armed conflict in the near term (Austin 2022, p. 6). This has, for example, already manifested itself in the Russia-Ukraine war, where both Russian and Ukrainian forces are making extensive use of unmanned vehicles across air, ground, and maritime domains; see, for example, Kallenborn and Plichta (2023). The fact that these platforms are either fully autonomous, or are remotely operated, means that the sensing capabilities that are required for their operation are also, in the words of the former Chief of Naval Operations, Admiral (Ret.) Mike Gilday, "making contested spaces more transparent and more lethal" (Gilday 2022, p. 5).

One one hand, the data collected via these sensing capabilities can be leveraged for better tactical or operational planning, and may even influence strategic planning, via the use of data-driven optimization or machine learning methods. For example, in 2017 the Naval Postgraduate School (NPS) conducted a multi-thread experiment (MTX) involving both researchers across NPS as well as fleet sponsors to demonstrate how a networked team of unmanned vehicles (UxVs) could provide communications support and battlespace awareness for a Naval Special Warfare (NSW) team conducting a notional direct-action mission on San Clemente Island, located off the coast of California (Schehl 2018). In order to provide this support, the UxVs were routed using optimization-based algorithms to both ensure that the UxV platforms could both be well-positioned to surveil locations of interest, as well as to maintain a robust communications capability for the NSW team (Lowry 2020).

The use of optimization-based algorithms in applications such as routing UxVs can, however, present security risks if the optimization objective does not account for the potential presence of a technologically-capable adversary. For example, Wigington (2021) showed that the locations of the supporting UxVs during the aforementioned MTX, which were prescribed by optimization-based algorithms that optimized for sensing value and communications robustness, could be used to accurately predict the actual movements of the NSW

team on the ground. Specifically, using a hierarchical time series model for the NSW team's movements given the movements of the UxVs that were fitted using actual NSW/UxV location data collected during the MTX, the locations of supported NSW units could be predicted with an accuracy on the order of tens of meters (Wigington 2021).

Our work is motivated by the need to account for the presence of an adversary in the design of decision-making algorithms, particularly those for the routing of UxVs. One of the goals is to make it more difficult for an adversary who is observing the prescribed movements/routes to infer operational intent (e.g., future movements, or destinations) based on those routes. Another goal is to account for potentially changing threat landscapes. Both of these are crucial considerations for planning algorithms that are meant to provide prescriptions for agents (e.g., UxVs) operating in contested environments.

1.1 Incorporating Deception

A natural way to reduce the extent to which an adversary can predict one's future movements is to make use of some form of deception. This idea is certainly not new; during World War II, for example, Allied convoys used so-called Zig-Zag Control Clocks to execute coordinated zig-zag maneuvers designed to disguise the convoy's true course from German submarines (Jones 2014). A similar effect could also be achieved via specially designed camouflage, such as the so-called "dazzle camouflage" paint patterns employed by Allied ships during both World Wars (Williams 2001, see Figure 1.1 below).

We are primarily interested in how deception can be incorporated into optimization-based route planning algorithms. In this context, usually a trade-off needs to be made between maximizing the route's efficiency and minimizing the degree to which it is "predictable" in the sense of signaling operational intent. For example, one way to measure a route's efficiency is in the travel time or distance needed to execute the route; here, shorter distances or travel times correspond to higher efficiency. A maximally efficient route between two points would be a shortest path between the points. If an adversary knows that a shortest path will be followed, however, it becomes relatively easy to infer operational intent by ruling out candidate destinations based on where the agent (e.g., UxV) has gone so far (see the left-hand side of Figure 1.2). This can be addressed by introducing movements that, while reducing the efficiency of the route, increase the degree to which the agent's operational

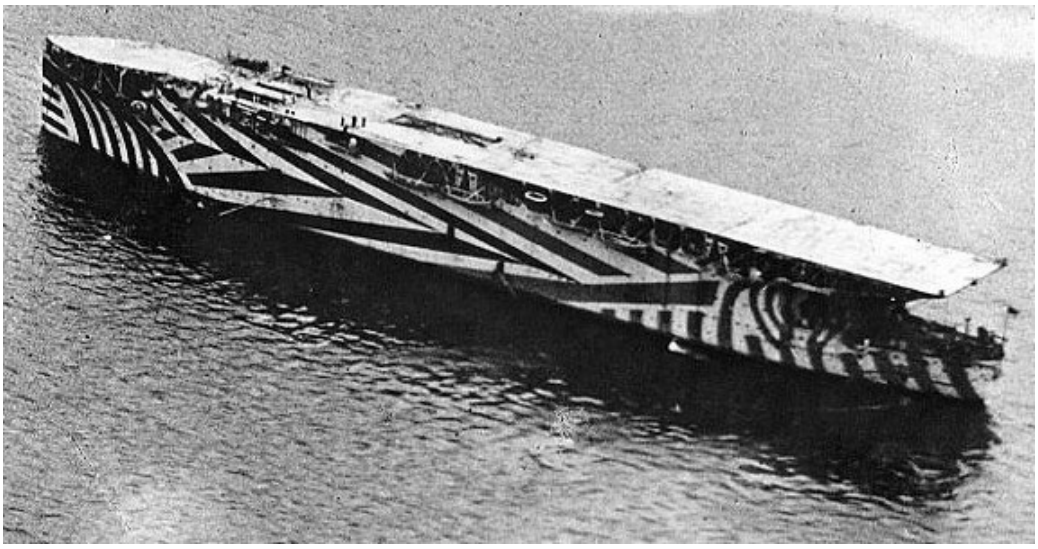

Figure 1.1. The HMS Argus with dazzle camouflage (Williams 2001).

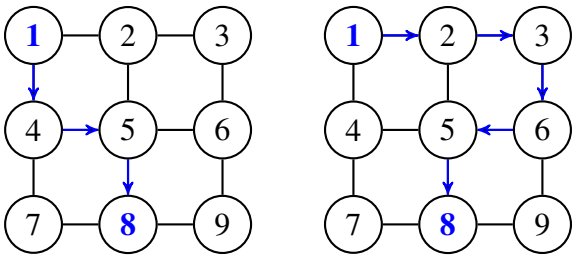

Figure 1.2. An efficient but predictable path (left), and a less efficient but less predictable path (right), from node 1 to node 8 in a grid network.

intent is hidden from an observing adversary who, for example, may be attempting to guess the agent's destination (see the right-hand side of Figure 1.2).

The question then becomes one of how to properly balance efficiency with predictability. Motivated by the so-called "red cell analysis" carried out by Wigington (2021), we consider this question in the context of routing a UxV in a way that reduces the capability of an adversary using the time-series methods in Wigington (2021) to predict the UxV's destination. Specifically, Plunkett (2023) studies the effectiveness of a randomized routing strategy derived from a game-theoretic formulation of the routing problem proposed by Tsitsiklis and Xu (2018) against this particular "red cell".

1.2 Routing in Contested Environments

Algorithms for making routing recommendations in contested environments also need to account for dynamically evolving threat or risk landscapes. Motivated by this, we consider the efficacy of iteratively re-solving the routing problem in light of new information on the current threat landscape as it becomes available. As operating in a contested environment also often entails reduced or denied communications capabilities, we also consider the problem of coordinating multiple agents (e.g., UxVs) in a de-centralized way. These topics are the subject of an ongoing Master's thesis project by NPS student CPT Yan-ru Lin, Republic of China Marine Corps, who is slated to graduate in December 2023.

1.3 Report Outline

The remainder of this report is organized as follows. Chapter 2 includes a formulation (Section 2.1) of a two-player zero-sum game that models a Blue agent's need to reach a destination in a way that minimizes the chance that Red will correctly guess their final location. Section 2.2 provides a description of the randomized routing strategy motivated by analyzing the two-player zero-sum game formulated in Section 2.1, and Section 2.3 describes numerical assessments of the randomized routing strategy on data from the 2017 MTX on San Clemente Island. Chapter 3 considers the problem of routing in contested environments. Section 3.1 presents an approach to risk-aware dynamic routing via online optimization, and Section 3.2 describes a framework for coordinating multiple agents in a communications-denied environment. Section 3.3 describes ongoing work in applying the methods of Sections 3.1 and 3.2 to a notional maritime resupply scenario around Taiwan. Finally, Chapter 4 provides conclusions and recommendations for future work.

CHAPTER 2:
Randomized Routing Strategies for Deception

In this chapter, we describe the game-theoretic model that motivates the class of randomized routing strategies we considered for planning deceptive routes for agents such as UxVs in mobile networked control systems. These strategies can to some extent balance the need for deception with the need for operating efficiency. Section 2.1 formulates the two-player zero-sum game and states some known structural results about this game. Section 2.2 specifies the class of randomized routing strategies we considered, which are parameterized by a "delay budget" for the agent. Finally, Section 2.3 presents numerical results from the application of this strategy to UxV routing using data from the 2017 MTX conducted by NPS researchers on San Clemente Island (Schehl 2018).

2.1 Game-Theoretic Routing Model

We consider the following routing model proposed by (Tsitsiklis and Xu 2018). The agent, which we denote as BLUE, moves among the nodes of an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with node set $\mathcal{V} = \{1, \ldots, n\}$ and edge set $\mathcal{E}$. The agent's starting node is denoted by v_1, and the agent wants to visit a certain goal node $v_* \in \mathcal{V}$. At the same time, an adversary which we denote by RED observes the movements of BLUE. The goal of this adversary, who does not know which of the nodes in $\mathcal{V}$ is BLUE's goal node, is to correctly predict which node is BLUE's goal node before BLUE reaches it for the first time.

The game proceeds in discrete time steps $t = 1, 2, \ldots$ as follows. On each time step t, BLUE is at some node $v_t \in \mathcal{V}$ and needs to decide which node $v_{t+1} \in \mathcal{V}$ to visit next, subject to that node being directly accessible via some edge (i.e., with the constraint that $(v_t, v_{t+1}) \in \mathcal{E}$). Simultaneously, RED must decide to either issue a prediction of which node is BLUE's goal node, or to continue observing BLUE, given the history of locations $v_1, \ldots, v_t$ that BLUE has visited so far.

2.1.1 Prediction Risk

The objective of Blue is to minimize a quantity called the *prediction risk* (Tsitsiklis and Xu 2018), while the objective of Red is to maximize it. To define the prediction risk, define the random variables $\hat{V}_t$ to be equal to 0 if Red elects to not issue a prediction on time step t, and to be equal to Red's prediction of Blue's goal node otherwise. The game ends if Red makes an incorrect prediction or Blue reaches its goal node v_* before $\hat{V}_t = v_*$ (i.e., Red issues the correct prediction), in which case Blue wins; otherwise, Red wins. The *prediction risk* is defined as the probability that Red will win; Blue wants to minimize it, while Red wants to maximize it.

The prediction risk can be viewed as a function of the strategies adopted by Blue and Red. A strategy for Blue is a possibly randomized decision rule that, for each time step t, prescribes the next node v_{t+1} to visit given Blue's current location v_t; we will generically denote such a routing strategy by ψ, and let Ψ denote the set of all routing strategies. The strategy adopted by Blue implies a corresponding (possibly random) trajectory $V_1^\psi, V_2^\psi, \ldots$ of visited nodes, and an associated time at which the goal node v_* is visited for the first time, denoted by

$$T^\psi = \min\left\{t \in \{1, 2, \ldots\} \mid V_t^\psi = v_*\right\}.$$

A strategy for Red is a possibly randomized decision rule under which, for each time step t, Red either sets $\hat{V}_t = 0$ (i.e., declines to make a prediction) or issues a goal node prediction $\hat{V}_t \in \mathcal{V}$ given the locations $V_1^\psi, \ldots, V_t^\psi$ that Blue has visited so far; we will generically denote such a prediction strategy by χ, and let $\mathcal{X}$ denote the set of all prediction strategies. Letting

$$T^{\psi,\chi} = \min\left\{t \in \{1, 2, \ldots\} \mid \hat{V}_t \in \mathcal{V}\right\}$$

be the time at which Red makes a goal node prediction for the first time, the prediction risk (i.e., the probability that Red wins) can be written as

$$c(\psi, \chi) = P\left(\hat{V}_{T^{\psi,\chi}} = v_*,\ T^{\psi,\chi} \leq T^\psi\right).$$

The objective of Blue is to minimize the prediction risk , while the objective of Red is to maximize it. One way to view the overall objective from Blue's perspective is that they want to find a *robust* routing strategy, in the sense that it minimizes the maximum prediction risk achievable by Red. In other words, a robust routing strategy for Blue solves the optimization

problem

$$\underset{\psi \in \Psi}{\text{minimize}} \quad c_*(\psi) = \max_{\chi \in \mathcal{X}} c(\psi, \chi).$$

Here, Blue is effectively taking a worst-case perspective of Red, in that Red will respond to any routing strategy that Blue follows with a maximally effective prediction strategy χ.

It is important to keep in mind, however, that deceiving Red is not Blue's only objective. Blue is trying to reach a goal location v_*, and will typically prefer to reach it sooner rather than later. In other words, all other things being equal, a more "efficient" route (e.g., a shortest path) is preferable. Hence Blue would actually like to solve a bi-criterion optimization problem: find the most efficient route that also minimizes the prediction risk (i.e., the probability that Red will win). Rather than attempt to directly solve is bi-criterion problem, we instead modeled the efficiency consideration as a constraint on the expected value of the time it will take Blue to reach the destination node v_*. Given a maximum allowable expected travel time b, let

$$\Psi(b) = \{\psi \in \Psi \mid E[T^\psi] \leq b\}$$

denote the set of all routing strategies ψ under which the travel time to the goal node v_* is at most b; given a particular "travel time budget" b, $\Psi(b)$ is the associated set of feasible routing strategies. The associated (constrained) optimization that Blue looks to solve in order to obtain a robust routing strategy is then:

$$\begin{aligned} &\text{minimize} \quad c_*(\psi) = \max_{\chi \in \mathcal{X}} c(\psi, \chi) \\ &\text{subject to} \quad \psi \in \Psi(b) \end{aligned} \tag{2.1}$$

2.2 Robust Randomized Routing Strategies

Tsitsiklis and Xu (2018) showed that routing strategies exist for Blue where the worst-case prediction risk (i.e., the prediction risk under any Red prediction strategy) can be bounded above by a constant that is inversely proportional to the travel time budget b. More precisely, letting $d_{\mathcal{G}}$ denote the *diameter* of the graph $\mathcal{G}$, which is defined as the length of the longest shortest path between any pair of nodes in $\mathcal{G}$, it follows from (Tsitsiklis and Xu 2018,

Corollary 1) that if BLUE's starting node v_0 is chosen uniformly at random from the set of all nodes $\mathcal{V}$, and b is BLUE's travel time budget, then there exists a routing strategy $\psi_* \in \Psi(b)$ for BLUE such that, for any prediction strategy χ adopted by RED,

$$c(\psi_*, \chi) \leq \frac{2}{b - d_{\mathcal{G}}}. \tag{2.2}$$

Specifically, Tsitsiklis and Xu (2018) specify a class of routing strategies called *segment-based strategies* that can provide this prediction risk guarantee.

A segment-based strategy consists of BLUE randomly selecting and following a path that is designed so that, even if the adversary knows the locations that BLUE plans to visit, each location will appear to be roughly equally likely to be the goal location v_* to the adversary. First, a set $\mathcal{S}$ of paths (referred to in Tsitsiklis and Xu (2018) as "segments") are generated that have a length of b, where b is the aforementioned travel time budget. More specifically, BLUE generates a total of $|\mathcal{S}| = 2(n-1)$ paths based on a depth-first traversal of the graph $\mathcal{G}$ Tsitsiklis and Xu (2018) that ensures that every location in the graph is contained in at least b of the paths in $\mathcal{S}$. BLUE then selects one of the paths uniformly at random from $\mathcal{S}$, and does the following:

- Navigate from the starting location v_1 to the first location on the selected path, via a shortest path.
- Follow the selected path from $\mathcal{S}$ until the goal node v_* is reached.

Tsitsiklis and Xu (2018) derive an exact formula for the worst-case prediction risk incurred by BLUE when a segment-based strategy is followed. This formula, along with the properties of the set of paths $\mathcal{S}$ described above, is then used to derive the upper bound given in (2.2).

2.3 Application: Routing Unmanned Vehicles

The thesis work of one of the students involved in this research effort (Plunkett 2023) consisted of evaluating the efficacy of segment-based strategies in defeating the predictive capability of the "Red-Cell Analyzer" proposed in Wigington (2021). Plunkett (2023) showed that, even if a single UxV in a team of UxVs follows a segment-based strategy, and the others make use of their originally planned routes, the predictive capability of this Red-Cell Analyzer could be degraded significantly.

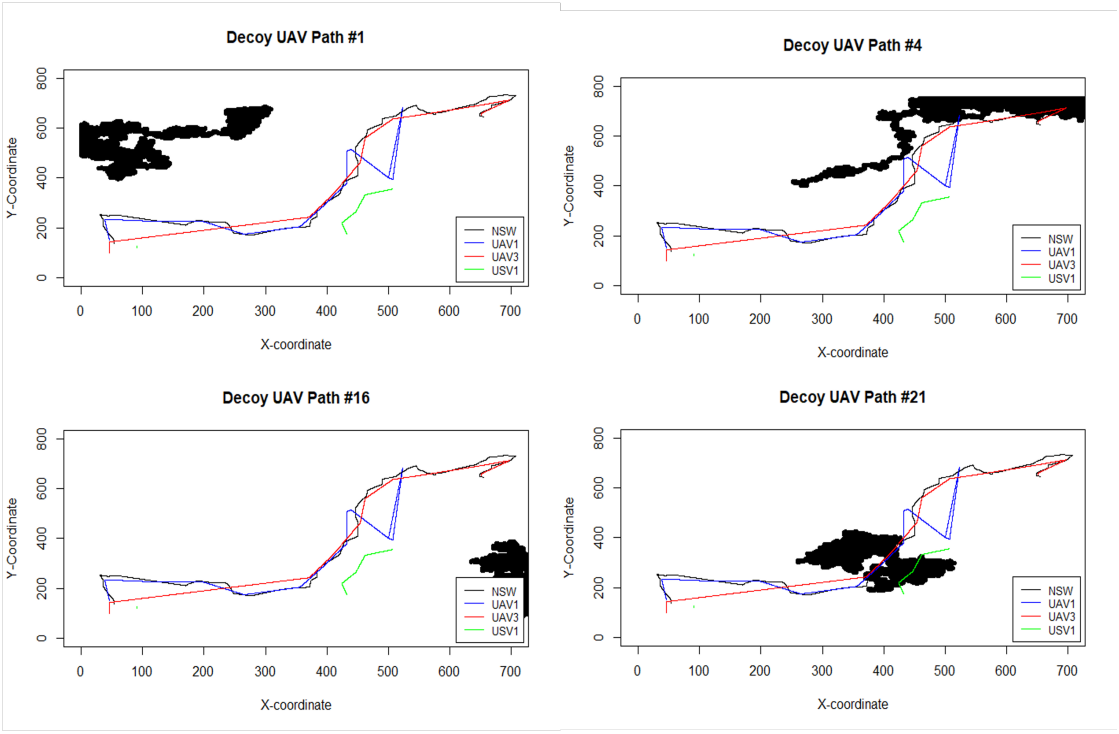

Figure 2.1. Examples of randomly generated paths for the decoy UAV, according to a segment-based strategy (Plunkett 2023).

Plunkett (2023) used a dataset of UxV paths collected from a MTX conducted by NPS researchers on San Clemente Island in 2017 (Schehl 2018). This dataset includes 7,005 observed UxV locations, taken at one-second intervals, of each unit in a team of five UxVs composed of three unmanned aerial vehicles (UAVs) and two unmanned surface vehicles (USVs). During the MTX, this team of UxVs was supporting a NSW team and one guided missile destroyer (DDG).

Since the locations of the UxVs are coordinate over a continuously varying area of operations, the location data had to first be discretized so that the movements of the UxVs could be interpreted as movements over the nodes in a graph (Plunkett 2023, Section 3.3). Then, a segment-based strategy with a travel time budget of $b = 7,005$ and a number of generated paths equal to $|\mathcal{S}| = 25$ was used to generate potential paths that a decoy UAV would then follow (See Figure 2.1).

The prediction capability of the Red-Cell Analyzer (Wigington 2021) was then evaluated

Decoy	Mean Prediction Error
UAV 1	212.5
UAV 2	537.1
UAV 3	356.1

Table 2.1. Mean error in predicted NSW destination, compared to their true destination (Plunkett 2023).

as follows. The routes of each UAV were in turn replaced with a route generated using the segment-based strategy, so that the UAV following this route acted as a decoy. Then, the Red-Cell Analyzer (Wigington 2021) was applied to attempt to predict the path and final destination of the supported NSW team.

This resulted in a striking degradation in the performance of the Red-Cell Analyzer. Table 2.1 shows the mean prediction error, over the 25 randomly generated decoy paths, for the respective cases where the first, second, and third UAV followed the route generated according to the segment-based strategy. For comparison, the *maximum* prediction error of the Red-Cell Analyzer in the absence of a decoy UAV was only 75 meters (Wigington 2021). Plunkett (2023) also provides some additional analyses related varying the travel time budget and the prediction of the entire NSW trajectory; see (Plunkett 2023, Chapter 4) for details.

CHAPTER 3:
Dynamic Risk-Aware Routing in Contested Environments

In this chapter, we outline a methodology for dynamically routing one or more agents (e.g., surface vessels, UxVs) in environments where the agents may not be able to directly coordinate with each other during routing (e.g., they are operating in a communications-denied environment), and where there are "threat hotspots" within the area of operations that may evolve over time. In this setting, it is important to balance the efficiency of the route followed (e.g., the travel time) with the need to account for areas of high threat. Section 3.1 describes the methodology for routing a single agent, and Section 3.2 considers the case of multiple agents. Finally, Section 3.3 describes an application of the dynamic routing methodology to a notional scenario involving the resupply of Taiwanese ports.

3.1 Online Optimization for a Single Agent

Consider a single agent whose goal is to navigate from one location in a network to another. If the agent's goal is to reach their destination as quickly as possible, and costs associated with each link in the network correspond to travel times, then the single agent can solve a shortest-path problem to find an "optimal" path; see Figure 3.1. Specifically, such a path can be obtained by solving a linear programming problem (LP).

To formulate this LP, consider a network represented by a graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$, where $\mathcal{V}$ is the set of locations ("vertices") and A is the set of directed arcs where for $i, j \in \mathcal{V}$, if $(i, j) \in \mathcal{A}$ then j is directly accessible from i without having to visit any of the other locations. Letting c_{ij} be the "cost" incurred when arc (i, j) is used, a cost-optimal path from a given starting location $s \in \mathcal{V}$ to a given destination $t \in \mathcal{V}$ can be computed by solving the following LP:

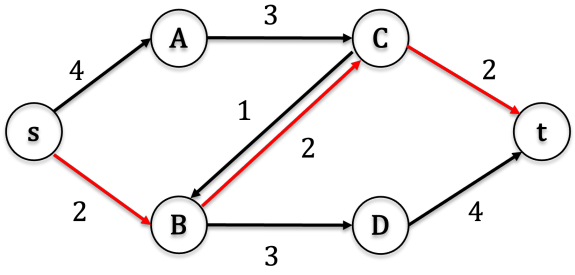

Figure 3.1. A shortest path from the agent's starting location s to the agent's destination t.

$$\begin{aligned}
&\text{minimize} && \sum_{(i,j)\in\mathcal{A}} c_{ij}x_{ij} \\
&\text{subject to} && \sum_{j:(i,j)\in\mathcal{A}} x_{ij} - \sum_{j:(j,i)\in\mathcal{A}} x_{ji} = \begin{cases} 1, & \text{if } i = s \\ -1, & \text{if } i = t \\ 0, & \text{otherwise} \end{cases} \qquad \forall i \in \mathcal{V} \qquad \text{(SP)} \\
& && x_{ij} \geq 0 \qquad \forall (i,j) \in \mathcal{A}
\end{aligned}$$

It is well-known that LPs of the form (SP) have an optimal solution $x^* = (x^*_{ij})_{(i,j)\in\mathcal{A}}$ where all of the x^*_{ij}'s are either 0 or 1; see e.g., (Ahuja et al. 1993, Section 11.12). For such a solution to the LP (SP), x^* can be interpreted as a shortest path from s to t by taking $x^*_{ij} = 1$ as indicating that arc $(i, j) \in \mathcal{A}$ should be used, and $x^*_{ij} = 0$ as indicating that the associated arc should not be used.

In a dynamic environment, the arc costs c_{ij} will change over time; for example, the presence/absence of an adversary near arc (i, j) may increase/decrease the cost of traversing (i, j). Here, we will consider discrete time steps $t = 1, 2, \ldots$ modeling discrete times at which re-routing decisions can be made. For each time step t, let $c_{ij}(t)$ denote the cost of traversing arc $(i, j) \in \mathcal{A}$; this may be interpreted as accounting for all of the intelligence on the "threat landscape" that is available prior to making a potentially updated routing decision at time step t.

The fact that the $c_{ij}(t)$'s differ across time steps t suggests that the LP (SP) should be re-solved at the start of each time step in order to account for information that becomes

1. At $t = 1$, solve (SP_1) and set $i_2 \in \mathcal{V}$ to be a location for which $x^*_{i_1 i_2}(1) = 1$.
2. Increment t by 1.
3. While $i_t \neq t$,
   (a) Solve (SP_t), and set $i_{t+1} \in \mathcal{V}$ to be a location for which $x_{i_t i_{t+1}} = 1$.
   (b) Increment t by 1.

Figure 3.2. Online optimization algorithm for routing a single agent.

available over time, in the form of updated arc costs. This is analogous to the re-calculation of recommended routes in civilian routing applications such as Google Maps[1] or Apple Maps[2]. This idea of re-solving an optimization problem in an "online" fashion has been used in other dynamic routing contexts; see, e.g., Liu (2021), Cone (2022), Liu et al. (2022), Marler (2022), Jockheck (2023), and Marler et al. (2023).

More precisely, for each time step $t = 1, 2, \ldots$, consider the following LP (SP_t) that is parameterized by t:

$$\begin{aligned}
\text{minimize} \quad & \sum_{(i,j)\in\mathcal{A}} c_{ij}(t) \cdot x_{ij} \\
\text{subject to} \quad & \sum_{j:(i,j)\in\mathcal{A}} x_{ij} - \sum_{j:(j,i)\in\mathcal{A}} x_{ji} = \begin{cases} 1, & \text{if } i = s \\ -1, & \text{if } i = t \\ 0, & \text{otherwise} \end{cases} \qquad \forall i \in \mathcal{V} \qquad (\text{SP}_t) \\
& x_{ij} \geq 0 \qquad \forall (i,j) \in \mathcal{A}
\end{aligned}$$

For each time step t, let $x^*(t) = (x^*_{ij}(t))_{(i,j)\in\mathcal{A}}$ denote an optimal solution of (SP_t).

At time step $t = 1$, the agent is at the initial location s. In general, we will denote the location of the agent at time step $t = 1, 2, \ldots$ by i_t, so that $i_1 = s$. The path that the agent ends up following can then be represented by $p = (i_1, \ldots, i_T)$, where T denotes the time step on which the destination location t is reached for the first time (i.e., $i_T = t$). This path p is generated according to the algorithm described in Figure 3.2. An example path for an instance where $s = \text{A}$, $t = \text{D}$, and $p = (\text{A}, \text{B}, \text{C}, \text{D})$, is shown in Figure 3.3.

[1] https://www.google.com/maps/
[2] https://www.apple.com/maps/

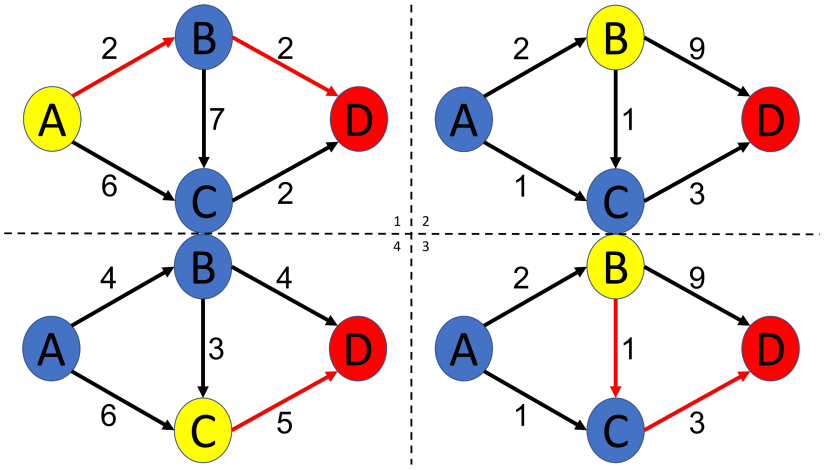

Figure 3.3. Example of a path computed using the online optimization algorithm.

3.2 Distributed Optimization for Multiple Agents

The multi-agent context that we consider is one where there is a set of locations $\mathcal{L} \subseteq \mathcal{V}$ that need to be visited, where for each location it suffices to have one of K available agents visit it. We will assume that the agents may not be able to coordinate with each other after time step $t = 1$. In other words, some initial coordination is possible, but once the agents embark they may have limited or no communication capabilities.

Our approach is as follows. At $t = 1$, a vehicle routing problem (VRP) is solved to determine each agent's itinerary, using the arc costs $c_{ij}(1)$. This can be accomplished by solving the following LP for $t = 1$, where the location d represents the starting location of all the agents[3]. The solution of this VRP partitions the set of locations $\mathcal{L}$ into subsets $\mathcal{L}_k$ of locations assigned to each agent $k = 1, \ldots, K$. More precisely, for each agent k, the solution to the VRP provides an order in which the locations in $\mathcal{L}_k$ should be visited; this order specifies the itinerary that the agent will follow.

[3]It is straightforward to extend this to allow for different starting locations for the agents, but for notational simplicity we will assume that all agents start from a common location/region, e.g., the continental U.S.

$$
\begin{aligned}
\text{minimize} \quad & \sum_{(i,j)\in\mathcal{A}} c_{ij}(1)\cdot x_{ij} \\
\text{subject to} \quad & \sum_{i\in\mathcal{V}} x_{ij} = 1 \qquad \forall j \in \mathcal{V}\setminus\{d\} \\
& \sum_{j\in\mathcal{V}} x_{ij} = 1 \qquad \forall i \in \mathcal{V}\setminus\{d\} \\
& \sum_{i\in\mathcal{V}\setminus\{d\}} x_{id} = K \\
& \sum_{j\in\mathcal{V}\setminus\{d\}} x_{dj} = K \\
& \sum_{i\notin\mathcal{L}} x_{ij} \geq 1 \qquad \forall j \in \mathcal{L}
\end{aligned}
\tag{VRP}
$$

For the remaining time steps $t = 2, 3, \ldots,$ each agent independently visits the locations in $\mathcal{L}_k$ according to their itinerary obtained from the solution of the VRP, where the actual routes are determined according to applying the online routing methodology defined in Section 3.1 to each agent.

3.3 Application: Resupply of Taiwanese Ports

The methodology described in this chapter is the subject of the thesis work of CPT Yan-ru Lin, who is a student in the Operations Research (OR) curriculum at NPS. In her thesis, CPT Lin will apply the methodology to a notional scenario involving the resupply of a number of Taiwanese ports using a team of supply vessels that, once underway, cannot communicate with each other. For details and up-to-date information on the results and analysis of this notional dynamic routing application, please contact Dr. Jefferson Huang at `jefferson.huang@nps.edu`.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4:
Conclusions and Recommendations

This chapter provides some conclusions (Section 4.1 and recommendations regarding potential follow-on research (Section 4.2).

4.1 Conclusions

Our primary finding is that the randomized routing strategy suggested by the two-player game proposed in Tsitsiklis and Xu (2018) can significantly degrade the predictive capability of the prediction algorithm developed in Wigington (2021). Specifically, it was able to increase the average route prediction error from a maximum of 75 meters, to roughly an average of 300 meters (Plunkett 2023). We also assessed the destination prediction error, which was roughly 368 meters on average (Plunkett 2023). We also observed that it was sufficient for one of the UxVs being observed to act as a decoy, by following a randomized path, while the others followed more efficient paths, to significantly degrade the prediction algorithm's performance. This finding provides empirical support for the importance of using deception (e.g., via decoys) and can inform the development of routing algorithms that take an observing adversary into account. In addition to these findings, we also assessed the method based on solving online shortest path problems using a scenario set in the Indo-Pacific, where the goal is to supply several Taiwanese ports in the context of an evolving threat landscape. The application to this notional scenario provides a proof-of-concept of the proposed approach and indicates its potential applicability to routing problems in the Indo-Pacific theater.

4.2 Recommendations

There are many promising directions for further research. One direction is to try to improve the prediction model in Wigington (2021), to model a more sophisticated adversary and to determine whether the randomized routing strategies evaluated in this research (Plunkett 2023) still effective against it. If not, further research should include studying ways in which the deception mechanism can be improved. In general, the goal would be to continue to

better understand how to balance deception with efficiency, and to address the reality that near-peer adversaries will likely have the capability to take advantage of observations of how our autonomous systems behave.

List of References

Ahuja RK, Magnanti TL, Orlin JB (1993) *Network Flows: Theory, Algorithms, and Applications* (Prentice Hall).

Austin LJ (2022) National Defense Strategy of The United States of America.

Cone SW (2022) *Casualty Evacuation Optimization in a Conflicted Environment*. Master's Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA, URL https://hdl.handle.net/10945/71052.

Gilday MM (2022) Chief of Naval Operations Navigation Plan.

Jockheck HW (2023) *Online Optimization to Increase Small Unmanned Aerial Vehicle Surveillance Capacity in Joint Forcible Entry Operations*. Master's Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA, URL https://hdl.handle.net/10945/72196.

Jones E (2014) Zig-Zagging: How to Confuse the Enemy at Sea. *Royal Museums Greenwich Blog* December 9, URL https://www.rmg.co.uk/stories/blog/curatorial/zig-zagging-how-confuse-enemy-sea.

Kallenborn Z, Plichta M (2023) Release the Robot Hounds: Providing Unmanned Ground Vehicles to Ukraine. *Commentary, Center for Strategic & International Studies* April 3, URL https://www.csis.org/analysis/release-robot-hounds-providing-unmanned-ground-vehicles-ukraine.

Liu Y (2021) *Solving Reward-Collecting Problems with UAVs Against the Stochastic Adversary Through Reinforcement Learning and Online Optimization*. Master's Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA, URL https://hdl.handle.net/10945/67148.

Liu Y, Vogiatzis C, Yoshida R, Morman E (2022) Solving reward-collecting problems with UAVs: A comparison of online optimization and Q-Learning. *Journal of Intelligent & Robotic Systems* 104:35.

Lowry B (2020) *Distributed Submodular Optimization for a UxV Networked Control System*. Master's Thesis, Department of Mechanical and Aerospace Engineering, Naval Postgraduate School, Monterey, CA, URL https://hdl.handle.net/10945/64926.

Marler K, Yoshida R, Vogiatzis C (2023) U.S. Marine Corps rapid planning and logistics routing against uncertainty. *Naval Engineers Journal* 135-1:115–125.

Marler KM (2022) *A Decision Process for Surface Medical Evacuation Routing Under Adversary Threat and Uncertain Demand Using Online Optimization*. Master's Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA, URL https://hdl.handle.net/10945/71114.

Plunkett KE (2023) *An Evaluation of Randomized Routing Strategies for Deception in Mobile Networked Control Systems*. Master's Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA, URL https://hdl.handle.net/10945/72041.

Schehl M (2018) NPS Research Team Explores the Boundaries of Unmanned Systems Through MTX. *Press Release, Office of University Communications, Naval Postgraduate School* June 11.

Tsitsiklis JN, Xu K (2018) Delay-predictability trade-offs in reaching a secret goal. *Operations Research* 66.

Wigington LW (2021) *Red Cell Analysis for Mobile Networked Control Systems*. Master's Thesis, Operations Research Department, Naval Postgraduate School, Monterey, CA, URL https://hdl.handle.net/10945/67831.

Williams DL (2001) *Naval Camouflage 1914–1945: A Complete Visual History* (Naval Institute Press).

Initial Distribution List

1. Anthony S. Tai, PhD
   Chief Engineer
   Electromagnetic Warfare S&T Division
   Spectrum Warfare Department
   Naval Surface Warfare Center, Crane Division
   Crane, Indiana

2. Defense Technical Information Center
   Ft. Belvoir, Virginia

3. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California