NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

UNDERSEA LOGISTICS: ROUTING OPTIMIZATION IN GRAY ZONE ENVIRONMENTS

by

Thuan D. Chu

September 2023

Co-Advisors:	Shelley P. Gallup Ruriko Yoshida
Second Reader:	Douglas J. MacKinnon

This project was funded in part by the NPS Naval Research Program.

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704-0188*
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC, 20503.			
1. AGENCY USE ONLY ***(Leave blank)***	**2. REPORT DATE** September 2023	**3. REPORT TYPE AND DATES COVERED** Master's thesis	
4. TITLE AND SUBTITLE UNDERSEA LOGISTICS: ROUTING OPTIMIZATION IN GRAY ZONE ENVIRONMENTS			**5. FUNDING NUMBERS** NPS-22-N355-A
6. AUTHOR(S) Thuan D. Chu			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000			**8. PERFORMING ORGANIZATION REPORT NUMBER**
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A			**10. SPONSORING / MONITORING AGENCY REPORT NUMBER**
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. This project was funded in part by the NPS Naval Research Program.			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.			**12b. DISTRIBUTION CODE** A

13. ABSTRACT (maximum 200 words)

Unmanned undersea vehicles (UUVs) are being built to meet the growing demand for undersea operational awareness and payload delivery, but there are no tools for operational planners to handle logistics regarding UUVs, such as routing logistic vessels to UUVs of differing capabilities, unique requirements, and time constraints. The feasibility of UUV employment in this manner would result in lower risk and broaden the scope of mission sets for our limited arsenal of manned submarines. Current models seek to optimize a single logistic vessel's pathing to a group of surface vessels. However, there are no models that seek to provide a global solution with multiple logistic vessels to multi-echelon UUVs. We generalize a prior model for medical evacuation with multiple types of unmanned vehicles using time-synchronized systems via a multiple-path orienteering problem. We extend that model to an undersea logistics problem in a contested environment, with heuristics to approximate global solutions. A successful model provides feasible solutions to a simple scenario for proof of concept and possible real-life scenarios.

14. SUBJECT TERMS			**15. NUMBER OF PAGES**
unmanned undersea vehicle, UUV, undersea logistics, vehicle routing, optimization, processor scheduling problem, multiple processor scheduling problem, MPSP, PSP			57
			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	**18. SECURITY CLASSIFICATION OF THIS PAGE** Unclassified	**19. SECURITY CLASSIFICATION OF ABSTRACT** Unclassified	**20. LIMITATION OF ABSTRACT** UU

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

**UNDERSEA LOGISTICS:
ROUTING OPTIMIZATION IN GRAY ZONE ENVIRONMENTS**

Thuan D. Chu
Lieutenant, United States Navy
BS, California State University, Long Beach, 2015

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN OPERATIONS RESEARCH

from the

**NAVAL POSTGRADUATE SCHOOL
September 2023**

Approved by: Shelley P. Gallup
Co-Advisor

Ruriko Yoshida
Co-Advisor

Douglas J. MacKinnon
Second Reader

W. Matthew Carlyle
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Unmanned undersea vehicles (UUVs) are being built to meet the growing demand for undersea operational awareness and payload delivery, but there are no tools for operational planners to handle logistics regarding UUVs, such as routing logistic vessels to UUVs of differing capabilities, unique requirements, and time constraints. The feasibility of UUV employment in this manner would result in lower risk and broaden the scope of mission sets for our limited arsenal of manned submarines. Current models seek to optimize a single logistic vessel's pathing to a group of surface vessels. However, there are no models that seek to provide a global solution with multiple logistic vessels to multi-echelon UUVs. We generalize a prior model for medical evacuation with multiple types of unmanned vehicles using time-synchronized systems via a multiple-path orienteering problem. We extend that model to an undersea logistics problem in a contested environment, with heuristics to approximate global solutions. A successful model provides feasible solutions to a simple scenario for proof of concept and possible real-life scenarios.

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

CA contested area

CVRP Capacitated Vehicle Routing Problem

LMACC Lightly Manned Autonomous Combat Capability

LV logistic vessel

MILP mixed-integer linear program

MPSP Multiple Processor Scheduling Problem

mTSP Multiple Traveling Salesman Problem

NM nautical mile

PSP Processor Scheduling Problem

SPSP Single Processor Scheduling Problem

SVRP Synchronized Vehicle Routing Problem

TSP Traveling Salesman Problem

UUV unmanned undersea vehicle

VRP Vehicle Routing Problem

VRPTW Vehicle Routing Problem with Time Windows

VSP Vehicle Scheduling Problem

THIS PAGE INTENTIONALLY LEFT BLANK

Executive Summary

Advancements in unmanned autonomous vehicles have enabled the U.S. Navy to significantly expand operations in the undersea warfare domain. Undersea operations often take place in gray zone environments, a challenging and ambiguous maritime domain characterized by unconventional activities. In such an environment, these unmanned vehicles play a vital role as they allow for higher flexibility and lower risk compared to operations with manned submarines. As this technology continues to mature, these unmanned undersea vehicles (UUVs) are expected to handle a larger range of operations. Despite this, the logistical aspects concerning unmanned vehicles have lagged behind and have not received the same level of attention (Hill 2022).

UUVs have unique capabilities and service requirements that necessitate human involvement. While UUVs are quickly built and cost-effective to acquire, manned logistic vessels (LVs) are few in number. This thesis proposes a model designed to optimize the routing and scheduling of a limited number of LVs in order to effectively service a network of UUVs. Taking elements from both vehicle and processor scheduling optimization problems, the model takes the form of a mixed-integer linear program (MILP) referred to as a Multiple Processor Scheduling Problem (MPSP) with synchronization constraints.

Among the various objective functions within the processor scheduling domain, this model incorporates two specific criteria: mean tardiness and mean lateness. By minimizing these metrics, the model successfully generates comprehensive schedule recommendations and visualizations for LVs to efficiently service UUVs within the network in a timely manner. To accommodate operations in gray zone environments such as contested areas (CAs), the model optimizes route and schedule recommendations according to the acceptable level of risk specified by the decision-maker.

This work paves the way for extended exploration and in-depth development within the realm of optimization concerning UUV logistics. Subsequent work could involve the fine-tuning of input parameters to align with practical operational conditions or delving into alternative model objectives to broaden the applicability and versatility of the model.

References

Hill R (2022) NPS-22-N355: Distributed maritime logistics for theater undersea warfare. NRP Topic Portal. Accessed April 4, 2023, https://nps.edu/web/naval-research-program/topics/-/topics/lLTOSW5Vv9hI/20304.

Acknowledgments

I would like to thank my advisors, Dr. Yoshida and Dr. Gallup, for the opportunity to conduct this work and for your continuous guidance. Thank you to my second reader, Dr. MacKinnon, for your adeptness in posing insightful questions during our weekly meetings. Dr. Chrysafis Vogiatzis, thank you for your guidance in the early stages of model formulation.

To my cohort, especially the NPS Honor Roll, thank you for sharing this experience with me.

Finally, thank you to my wife, Ailene, for your never-ending support and for selecting NPS over Japan.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:
Introduction

Militaries have and will continue to leverage unmanned autonomous vehicles to enhance their operational capabilities. In today's complex and rapidly evolving geopolitical landscape, these operations often extend into challenging and ambiguous environments known as gray zones. Encompassing a wide range of activities that occur between peace and war, these efforts include everything "from nefarious economic activities, influence operations, and cyberattacks to mercenary operations, assassinations, and disinformation campaigns" (Starling 2022).

Within the maritime domain, gray zones such as contested areas (CAs) blur the distinction between peace and conflict, presenting unique challenges for the armed forces. The presence of CAs further complicate underdeveloped logistics surrounding unmanned undersea vehicles (UUVs). These vehicles play a crucial role in diverse military operations, ranging from intelligence gathering to surveillance and reconnaissance. Therefore, ensuring the efficient deployment and sustained operation of these vehicles becomes imperative for upholding maritime dominance and securing strategic advantages within the intricate contexts of both present and future operating environments.

This thesis aims to address the challenges of optimizing logistics for servicing UUVs within gray zone environments. By leveraging optimization techniques, we strive to develop a tool that planners and decision-makers can employ to efficiently plan and schedule routes for servicing these vehicles. This will enhance the overall efficiency and effectiveness of UUV operations within these complex and dynamic settings.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 2: Background

In this chapter, we aim to establish the context for the problem addressed in this thesis, along with a concise overview of pertinent optimization methods and prior research efforts.

The problem was derived from Naval Research Program topic *NPS-22-N355: Distributed Maritime Logistics for Theater Undersea Warfare*, where Hill (2022) inquires on "how to perform logistics resupply methods for a small, distributed, agile undersea force in Phase 0-2 environments." In the topic details, Hill states that "the undersea warfare community owns, and will continue to acquire, platforms and capabilities beyond those of the previous submarine community" and that "given the dwindling of surface resupply capabilities and the advances in communications/denial of communications/distributed lethality, over the decades since the last mature submarine resupply capability existed, how, in future fights, will the undersea warfare community be resupplied, and how will they assist in resupplying other members of the joint theater force?"

2.1 Motivation

The undersea warfare environment is in a state of evolution towards unmanned and autonomous vehicles. "Challenges to logistics are complex and include the development of how and when needs arise for different platforms doing myriad duties in undersea and seabed warfare" (Gallup 2022). There are a plethora of unmanned autonomous platforms in development for undersea warfare by private and government institutions, each with their own set of needs and requirements for sustainment. In example, a Knifefish UUV can conduct mine countermeasure operations for 16 hours while the Riptide UUV can conduct signals intelligence operations for 36 hours (Baesystems 2023; Naval Technology 2022).

Figure 2.1 depicts a hypothetical scenario where a network of UUVs have been deployed to an operating area south of Natuna Besar, operating in a gray zone environment. These UUVs are assigned to monitor a 120 x 120 NM operating area—making the area transparent. A limited number of logistic vessels (LVs) are assigned to service all UUV needs. On a given day, certain UUVs are required to be serviced, each with differing requirements and priorities. Depending on the needs of a UUV (battery recharge, maintenance requirements, etc.), each UUV will have a different service time, time window for service and may require more than one LV for service. UUVs are programmed to rendezvous at a certain location during their time window for service, which coincides with remaining battery capacity. A UUV not serviced by the end of its time window will naturally stop operating and rise to the surface by design. The risk of losing a UUV is proportional the time it spends floating on the surface. With the limited number of LVs available at disposal, it may be infeasible to service all UUVs within their respective time windows. What is the best order for servicing these UUVs to minimize this risk? What if there are CAs within the operating area?

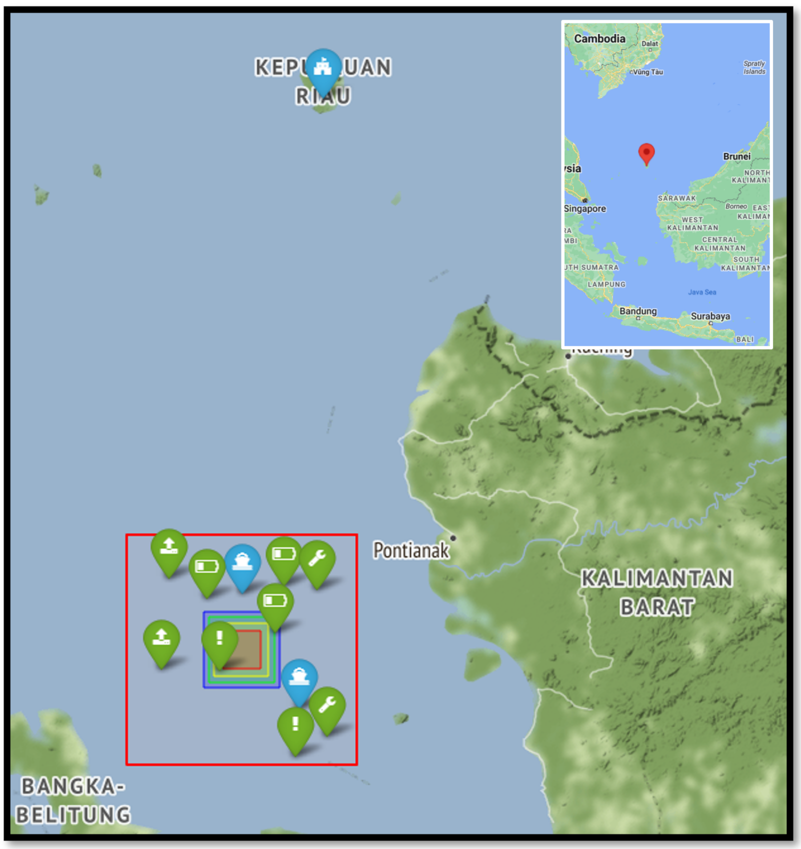

Figure 2.1. Map snippet of operating area generated using Folium. Shown here is a 120x120NM operating area (red) with two UUVs (blue), nine UUVs with differing service requirements (green), and one CA in the center.

2.1.1 Logistic Vessel

As part of our mission to address the logistical support challenge for UUVs, we consider a new type of ship called the Lightly Manned Autonomous Combat Capability (LMACC), which is a lightweight, fast, agile, and well-armed warship with minimal crew (Gallup 2023). The LMACC integrates autonomy within the vessel, allowing humans to focus on perceiving and navigating complex situations. "These vessels will work in concert with unmanned surface vessels" (Gallup and Mun 2020). The LMACC's ability to collaborate with unmanned surface vessels makes it an excellent choice for carrying out logistic support missions. In the analysis conducted for this thesis, the LMACCs will serve as the designated logistic vessel.

2.2 Optimizing Logistics

There are a couple of ways that we can approach the problem presented. For example, we could determine the optimal order for LVs to service UUVs such that the least fuel (measured by total distance traveled) is expended by LVs or that all UUVs are serviced within the allotted time windows. We will briefly discuss a few optimization models that went into consideration for our final model.

2.2.1 The Traveling Salesman Problem

The Traveling Salesman Problem (TSP) is a fundamental routing problem in optimization where the goal is to route a salesperson from a starting location to all customers within a network only once and return back to origin in the shortest distance (Ahuja et al. 1993). The Multiple Traveling Salesman Problem (mTSP) is a generalization of the TSP where all salesmen depart from origin and travels to a subset of customers before returning to origin.

2.2.2 The Vehicle Routing Problem

Similar to the mTSP, the Vehicle Routing Problem (VRP) is a generalization of the TSP where a fleet of vehicles travels to customers within a network (Ahuja et al. 1993). Ahuja states further that additional constraints result in variations of the VRP such as

- Capacitated Vehicle Routing Problem (CVRP) where each vehicle has a capacity limit such as space, fuel, operation time, etc.

- Vehicle Routing Problem with Time Windows (VRPTW) where customers can only be visited and/or serviced within a certain time window.
- Synchronized Vehicle Routing Problem (SVRP) where multiple vehicles are required simultaneously to service a single customer.

2.2.3 The Vehicle Scheduling Problem

Like the VRP, the Vehicle Scheduling Problem (VSP) is related to optimizing the movement of vehicles to serve a set of locations or customers. The VSP focuses specifically on the scheduling of tasks performed by vehicles and aims to minimize the makespan, improve resource utilization, or optimize other scheduling criteria (Bodin et al. 1981). It involves assigning tasks to vehicles, determining their execution sequence, and allocating time slots for task completion.

2.2.4 The Processor Scheduling Problem

Rardin (2017) states that the Processor Scheduling Problem (PSP) is a basic optimal job scheduling problem which "seeks an optimal sequence in which to complete a given collection of jobs with varying processing times by a single processor that can only accommodate one job at a time." Each job has a start time for when it is released for processing and due date. The Multiple Processor Scheduling Problem (MPSP) is a generalization of the Single Processor Scheduling Problem (SPSP) where multiple machines are tasked to complete the collection of jobs. Processor scheduling problems have a wide variety of objective functions such as makespan (maximum completion time), mean lateness, mean tardiness, etc.

2.3 Previous Work

Cone (2022) describes the synchronization of two echelons of vehicles to effect patient transfer for extraction in a VRP. He characterizes this problem as a "capacitated team orienteering problem with time windows and synchronization." Cone's thesis provides an optimization method for routing various medical evacuation assets to casualties with various severities in a contested environment. Like traditional VRPs, his model provides a minimum number of transport vessels required to meet time-window and other customer requirements.

Cone's work introduces the concept of vehicle synchronization which will be applied to our model.

In Sorenson (2020), surface supply ships are routed within a network of ports operating within a contested environment. Sorenson achieves this by conceptualizing each vessel type as a distinct layer within the network, with each layer characterized by a specific risk level. The presence of CAs introduces a threat region, within which the survival probability of a supply ship traveling on an edge that intersects the threat region is confined between zero and one. Furthermore, every vessel type is associated with a minimum survival probability, and routing decisions are made based on edges with survival probabilities surpassing this threshold. The primary objective of this model is to maximize the amount of supplies delivered to customers within a designated timeframe. Although this is a different approach to accounting for contested areas, the underlying concept of treating various vessel types as individual network layers is applicable within our own model.

2.4 Our Model

Given the unique characteristics of our problem, employing a traditional VRP/VSP model would often be impractical due to the constrained number of available LVs within the described scenario, which could prevent the fulfillment of all time window requirements. Therefore, we combine elements of SVRP, MPSP and previous works to formulate the model discussed in the following chapter.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 3:
Methodology and Formulation

The goal of this thesis is to produce a tool that would allow decision-makers to not only plan and allocate resources to solve the problem set defined previously, but to reveal potential capability and capacity gaps of UUVs and LVs. In this chapter, we describe our process and methodology for routing and scheduling LVs around CAs.

3.1 Scenario Assumptions

Due to the ongoing advancements in technology, it is challenging to provide a precise set of inputs for the model. Therefore, for the proof of concept, we have provided an approximate range for inputs such as UUV recovery time and service time. On a similar note, model inputs for CA and LV locations are also randomized. Examples of model input parameters are provided in Table 3.1.

Table 3.1. UUV parameters

Latitude	Longitude	Priority	Start Window	End Window	Service Time	Sync
-1.17	107.30	4	333	484	90	0
-0.85	107.79	1	1215	1440	117	1
-0.44	107.87	1	1172	1353	73	0
-1.16	106.80	3	316	503	71	0
-0.47	108.15	2	578	743	104	0
-1.92	107.96	4	194	358	82	0
-0.55	107.20	1	246	403	70	0
-1.75	108.24	2	349	556	118	0
-0.38	106.87	3	138	283	84	0

Data input for the model is randomly generated with some constraints to confine the scenario. All parameters are randomly generated with a uniform distribution. The constraint for latitude and longitudes is within the 120 nautical mile (NM) operating area defined in the scenario aforementioned in Section 2.1. UUV priorities are established based on the service required (1: Battery Replacement, 2: Routine Maintenance, 3: Storage Replacement, 4: Critical Maintenance), meaning tasks with higher priorities are penalized higher for being late or tardy. Service times can vary between one to two hours, and time windows are generated such that the duration encompasses service time plus an additional random time interval up to two hours. Lastly, the binary variable Sync determines whether the UUV requires more than one LV in order to be serviced.

3.1.1 Contested Areas

In this model, we made a general assumption that takes into account the gray zone environment we operate in. We consider contested areas as regions posing a higher level of risk compared to non-contested ones. However, for the sake of abstraction, we do not delve into specific details regarding the contents of these areas. This model allows the decision-maker to state the level of risk associated with the CAs, which we defined as how close to the center of the CA the LVs are allowed to get. The model breaks down risk into five areas:

- 0: No risk. Stay out of the CA.
- 1: Low risk. Stay out of 75% of the CA.
- 2: Medium risk. Stay out of 50% of the CA.
- 3: High risk. Stay out of 25% of the CA.
- 4: Max risk. Ignore the CA.

For example, if the decision-maker says that the CA has a level *2* risk associated with it, then LVs must remain outside fifty percent of the area, measured from center outward. If there are UUVs inside a CA that is not reachable with the level of accepted risk, then the UUV will be excluded from the model's analysis for routing and scheduling.

Table 3.2. CA parameters

Latitude	Longitude	Size (NM)
-1.06	107.38	38
-1.73	107.12	29
-0.49	107.9	35

Table 3.3. LV parameters

Latitude	Longitude	Speed (Kt)
0.2	107.0	10
-2.0	108.0	12
-0.5	108.4	10

3.2 Determining Shortest Path and Process Time

In contrast to the traditional VRP, our scenario involves routing vessels with a unique complexity: there are no obstructions between any two points in our operating area, resulting in a fully connected network. This feature provides direct connections between all locations, leading to an increased number of possible routes and combinations for the vessels to traverse. A fully connected network is required since we are modeling each service vehicle as a processor, UUVs as jobs requiring process, and ensuring that each processor is capable of servicing any job in any order. Introducing CAs to our problem presented an additional challenge. We needed to ensure that the network's full connectivity was maintained while adhering to the constraints imposed by the CAs. Finding a feasible solution required efficiently navigating around these contested regions without compromising the overall network connectivity. To accomplish this, we introduced intermediary nodes into the network.

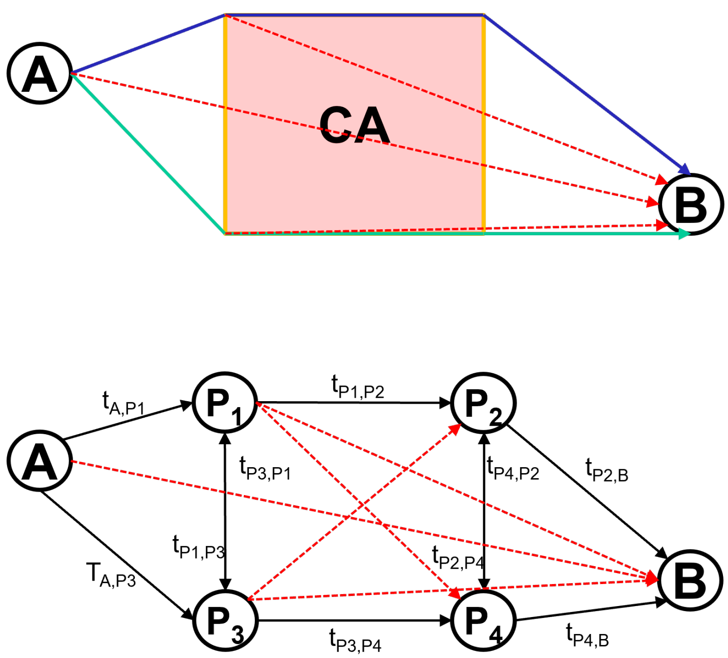

Figure 3.1. A visualization of valid paths between Point A and Point B when separated by a contested area. Point A represents the location of a UUV or LV, and Point B represents another UUV location. P1, P2, P3, and P4 represent the corners of the CA. The time it takes to travel from two points is denoted by t. For example, $t_{A,P1}$ is the time it takes to get from Point A to Point P1. Red arrows represent invalid paths from any point to Point B, as they intersect the CA.

Each intermediary node represents a corner of the CA and is determined based on the acceptable level of risk. The model calculates the minimum time required to process a job by breaking it down into smaller steps. Firstly, valid edges between Point A and Point B are determined, representing the travel time between each point based on the vehicle's speed, excluding those that intersect the contested area. Next, the shortest path is determined using Dijkstra's algorithm which efficiently calculates shortest path in a weighted graph (Ahuja et al. 1993). By combining the shortest path with the predetermined service time for the UUV at Point B, the total service time required for that UUV by the specific LV is obtained. This process is repeated for all pairs of nodes in the network, providing the input for our multi-processor scheduling model.

3.3 Complete Formulation

Sets and Indices:

$i, j \in N = \{1, 2, \ldots, n\}$		Unmanned Undersea Vehicle (UUV)
$k, l \in V = \{1, 2, \ldots, v\}$		Logistic Vehicle (LV)
$N^c \subset N$	$\forall c \in C$	UUVs requiring synchronization

Data:

t_{ij}^k	travel time of vehicle k from UUV i to UUV j [minute]
s_i	service time required for UUV i [minute]
a_i	lower bound time window for UUV i [minute]
b_i	upper bound time window for UUV i [minute]
p_i	priority of UUV i
d_i	whether UUV i requires synchronizing between LVs [0,1]

Decision Variables:

$x_{ik} =$ start of service time for UUV i by vehicle k

$$y_{ijk} = \begin{cases} 1 & \text{if job } i \text{ starts before job} j \text{ by vehicle } k \\ 0 & \text{otherwise} \end{cases}$$

$$z_{ik} = \begin{cases} 1 & \text{if job } i \text{ is processed by vehicle } k \\ 0 & \text{otherwise.} \end{cases}$$

Formulation:

$$\min \frac{1}{n} \sum_{i \in N} \sum_{k \in V} (x_{ik} - b_i) z_{ik} p_i \tag{3.1}$$

s.t.

$$\max\{0, a_i\} \leq x_{ik} \qquad \forall i \in N, \forall k \in V \tag{3.2}$$

$$x_{ik} + s_i + t_{ij}^k z_{jk} \leq x_{jk} + M(1 - y_{ijk}) \qquad \forall i \in N, \forall j \in N, \forall k \in V \tag{3.3}$$

$$x_{jk} + s_j + t_{ji}^k z_{ik} \leq x_{ik} + M y_{ijk} \qquad \forall i \in N, \forall j \in N, \forall k \in V \tag{3.4}$$

$$z_{0k} = 1 \qquad \forall k \in V \tag{3.5}$$

$$x_{0k} \leq b_0 \qquad \forall k \in V \tag{3.6}$$

$$\sum_{k \in V} z_{ik} = 1 + d_i \qquad \forall i \in N^c \tag{3.7}$$

$$(x_{ik} - x_{il}) z_{ik} z_{il} = 0 \qquad \forall k \in V, \forall l \in V, \forall i \in N^c \tag{3.8}$$

$$x_{ik} \geq 0 \qquad \forall i \in N, \forall k \in V \tag{3.9}$$

$$z_{ik} \in \{0, 1\} \qquad \forall i \in N, \forall k \in V \tag{3.10}$$

$$y_{ijk} \in \{0, 1\} \qquad \forall i \in N, \forall j \in N, \forall k \in V. \tag{3.11}$$

3.3.1 Set and Indices

A graph G(V,E), where V denotes the vertices, and E represents the edges. Within this graph, we have two distinct sets—one for UUVs or jobs requiring service, and the other for LVs or processors responsible for performing the service. Additionally, a third set is defined, consisting of a subset of UUVs that require two LVs to service them synchronously.

3.3.2 Data and Parameters

The main data required for a processor scheduling problem includes the service time requirements and time windows, which provide constraints on when a job is available for processing, how long it takes to process, and when it is due. The travel time of each vehicle between two nodes contributes to the total process time required for each node. While we could combine this travel time with the service time beforehand, we chose to present them separately in our model to highlight the distinction between our approach and a standard processor scheduling optimization model. In addition to these parameters, the inclusion of priority and synchronization further increases the complexity and deviation from the standard model.

3.3.3 Objective Function

There are many different types of objective functions for processor scheduling optimization models to include makespan (maximum completion time), flow time, lateness, and tardiness. "Optimization over each of these is equivalent to optimization over the corresponding mean measure because the total is constant n times the mean" (Rardin 2017). For our objective (3.1), we chose to minimize the mean lateness for servicing each UUV, scaled by the priority of each job. The model establishes a correlation between four priorities and the service requirements for UUVs. These priorities are chosen arbitrarily to illustrate a straightforward implementation of prioritization within the model. Notably, higher priority assigned to a job corresponds to a more substantial reward (or penalty) when servicing the specific UUV associated with that job. To demonstrate the flexibility of our model, we also implemented the mean tardiness objective:

$$\min \frac{1}{n} \sum_{i \in N} \sum_{k \in V} max\{x_{ik} - b_i, 0\} z_{ik} p_i \tag{3.1a}$$

The motivation for examining these two objective functions was to assess whether it is more advantageous to incentivize timely servicing of the UUV, rather than solely penalizing lateness in providing the service.

3.3.4 Constraints

We provide details on our model constraints as follows:

1. Constraint 3.2 requires each job to start after present time (time = 0) and the time at which the start time window opens.
2. Constraints 3.3 and 3.4 are "disjunctive constraint pairs to ensure that only one job is in progress on any processor at any time" (Rardin 2017). Rardin also notes that "*M* is a large positive constant to enforce this conditional constraint pair."

3. Constraints 3.5 and 3.6 ensure that processing must start with the current location of LVs.
4. Constraint 3.7 ensure that all jobs require at least one processor, and those that require synchronization has two processors assigned.
5. Constraint 3.8 ensure jobs that require synchronization must start at the same time.
6. Constraints 3.9, 3.10 and 3.11 apply variable-type restrictions.

3.3.5 Computational Complexity

Since our problem is a generalization of the VRP and the PSP, its complexity is known to be at least NP-hard (Paulo and Vigo 2002), and approaches NP-complete in certain cases (Garey and Johnson 1979). We have observed that for small-sized problems such as the scenario presented in Figure 2.1, global optimal solutions are obtained in approximately one second. In the next section, we will present the results from our model and discuss the optimality gap analysis for medium-sized networks.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 4:
Results and Analysis

In this chapter, we present the outcomes of our model's application to the straightforward scenario illustrated in Figure 2.1 and model performance for more complex networks.

4.1 Model Outputs

Our objective was to develop a user-friendly tool that empowers decision-makers to explore various scenarios and visualize potential courses of action (COA). To achieve this, we designed and implemented an intuitive template for easy data input. Users can conveniently provide data via an Excel sheet, following the format demonstrated in Tables 3.1, 3.2 and 3.3.

Figure 4.1 presents an Excel file containing the schedule for each LV in chronological order. This schedule offers a detailed overview of the expected arrival and departure times of each LV is expected to arrive at and depart each node, at different nodes, providing essential information for decision-makers to plan activities and operations effectively. The time interval between arrival and departure represents the estimated travel time from the previous node combined with the required service time.

In Figure 4.2, all information from the Excel file is available in an interactive file, allowing users to visualize and analyze the situation more comprehensively. By using this visual representation, we can easily discern the path each vehicle will follow for completing jobs in a particular order. It also helps identify which jobs are expected to be completed on time, which ones require synchronization, and those that are unreachable due to the accepted level of risk.

Vehicle	Node	Arrival Time	Departure Time
0	0	0.1	67.41
0	5	194	276
0	7	372.8	490.8
0	4	876.97	1056.9
0	1	1215	-
1	0	0.1	0.1
1	8	193.59	277.59
1	3	513.68	584.68
1	6	803.71	1010.05
1	1	1215	1332
1	2	1456.59	-

Node 2:	103.59 mins late
Node 3:	10.68 mins late
Node 4:	133.97 mins late
Node 6:	400.71 mins late

Figure 4.1. Schedule Excel file

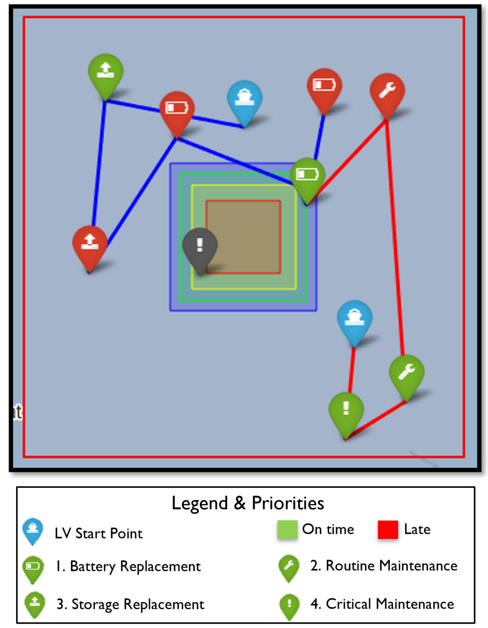

Figure 4.2. Interactive map HTML file

The legend accompanying the visual in Figure 4.2 provides details on color scheming and prioritization of UUVs. Notably, in this scenario, the accepted risk level prevents us from servicing the UUV node requiring critical maintenance, symbolized by a grey icon with an exclamation mark. The other UUV service requirements include battery replacement, storage replacement, and routine maintenance. The numbers adjacent to the icons correspond to a priority level for these service requirements. However, it is important to mention that these service requirements and priorities are arbitrary and not exhaustive, as they represent a subset of all potential service needs.

4.2 Analyzing Risk

To explore scheduling and routing variations, we conduct additional model runs, considering distinct levels of acceptable risk as defined in Section 3.1.1. Furthermore, we analyze the model's performance between two objective functions, mean lateness and mean tardiness. In the following discussion, note that node numbers referenced are specific to the sub scenarios with different levels of risk.

Figures 4.3(a) and 4.3(b) present the optimal solutions in the absence of accepted risk. In Figure 4.3(a), a prohibition sign overlays the entire CA to illustrate invalid paths crossing into the CA, rendering the UUV nodes inside unreachable for servicing. The objective value of -101 minutes indicates that, on average, the LVs initiate service on UUVs 101 minutes before their End Time Window (with priority factored in). Alternatively, a mean tardiness time of 100 minutes in 4.3(b) shows that, of the three late jobs, their lateness averaged to 100 minutes. Despite the disparity in objective functions, we observe from Figures 4.3(c) and 4.3(d) that the optimal schedule and paths for LVs remain identical in both cases.

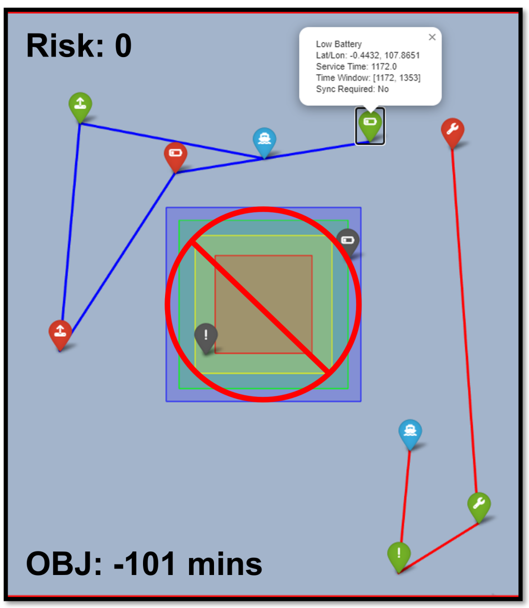

(a) Mean lateness map

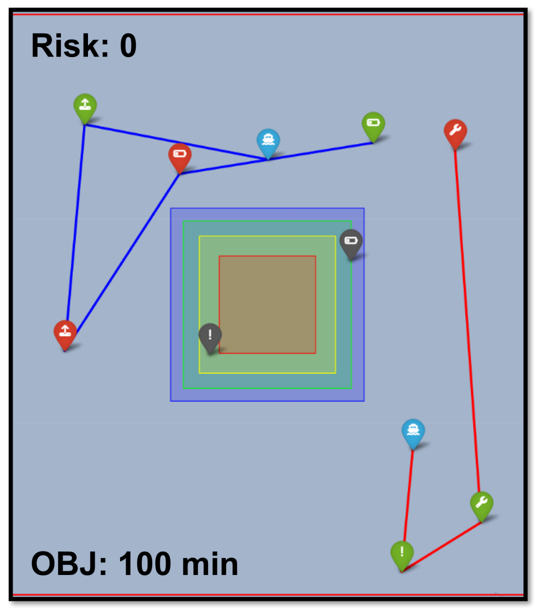

(b) Mean tardiness map

Vehicle	Node	Arrival Time	Departure Time
0	0	0.1	67.41
0	4	194	276
0	6	372.8	490.8
0	3	876.97	-
1	0	0.1	0.1
1	7	193.59	277.59
1	2	513.68	584.68
1	5	803.71	968.2
1	1	1172	-

(c) Mean lateness Excel

Vehicle	Node	Arrival Time	Departure Time
0	0	1	67.41
0	4	194	276
0	6	372.8	490.8
0	3	876.97	-
1	0	0.1	0.1
1	7	193.59	277.59
1	2	513.68	584.68
1	5	803.71	968.2
1	1	1172	-

(d) Mean tardiness Excel

Figure 4.3. Risk 0 model outputs

Our observations reveal minimal differences between `Risk 0` (Figure 4.3) and `Risk 1` (Figure 4.4) in the provided scenario, likely attributed to the scenario's simplicity designed to validate the model's functionality. While most nodes share identical arrival times, departure times, and objective values, it is noteworthy that `Risk 1` results in a swap of service providers for `Node 1`, with `Vehicle 0` taking over from `Vehicle 1`. This occurrence is possible due to the absence of restrictions on the total distance traveled by LVs, as our primary concern is timely service, as reflected in the objective function.

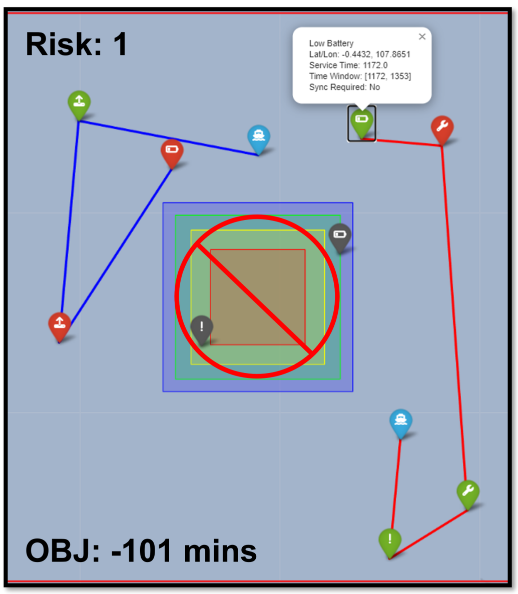

(a) Mean lateness map

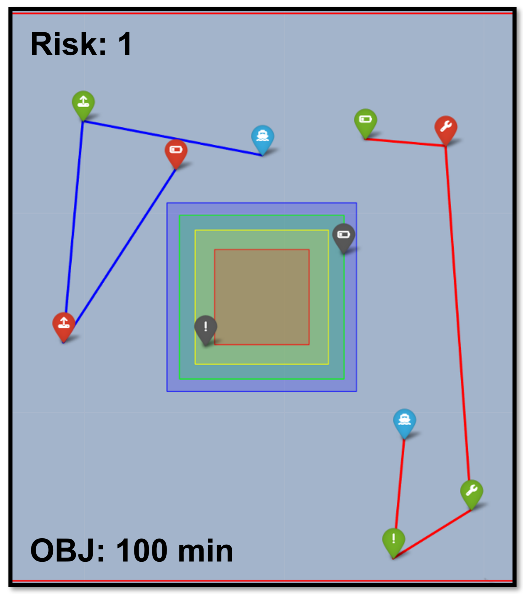

(b) Mean tardiness map

Vehicle	Node	Arrival Time	Departure Time
0	0	0.1	67.41
0	4	194	276
0	6	372.8	490.8
0	3	876.97	1086.61
0	1	1172	-
1	0	0.1	0.1
1	7	193.59	277.59
1	2	513.68	584.68
1	5	803.71	-

(c) Mean lateness Excel

Vehicle	Node	Arrival Time	Departure Time
0	0	1	67.41
0	4	194	276
0	6	372.8	490.8
0	3	876.97	1086.61
0	1	1172	-
1	0	0.1	0.1
1	7	193.59	277.59
1	2	513.68	584.68
1	5	803.71	-

(d) Mean tardiness Excel

Figure 4.4. Risk 1 model outputs

With the acceptance of `Risk 2` as shown in Figure 4.5, the LVs are now successfully routed to UUV `Node 1`, located within the lime-green zone of the CA, requiring a battery replacement. This service achievement was made possible through precise synchronization between both vehicles.

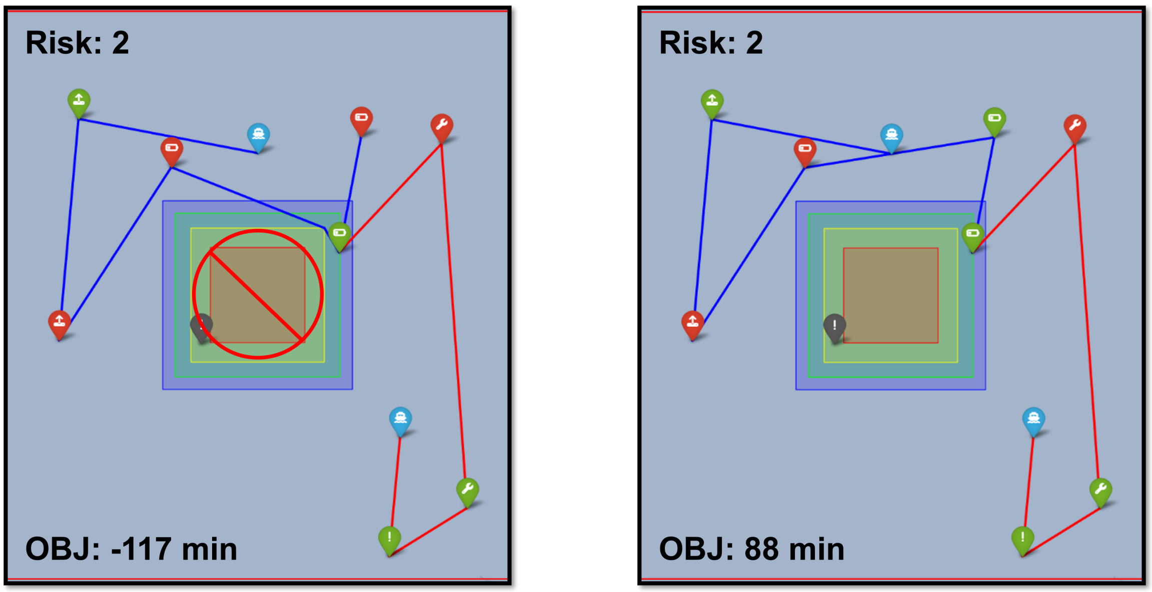

(a) Mean lateness map

(b) Mean tardiness map

Vehicle	Node	Arrival Time	Departure Time
0	0	0.1	67.41
0	5	194	276
0	7	372.8	490.8
0	4	876.97	1056.9
0	1	1215	-
1	0	0.1	0.1
1	8	193.59	277.59
1	3	513.68	584.68
1	6	803.71	1010.05
1	1	1215	1332
1	2	1456.59	-

(c) Mean lateness Excel

Vehicle	Node	Arrival Time	Departure Time
0	0	1	67.41
0	5	194	276
0	7	372.8	490.8
0	4	876.97	1222.5
0	1	1380.59	-
1	0	0.1	0.1
1	8	193.59	277.59
1	3	513.68	584.68
1	6	803.71	968.2
1	2	1172	1256
1	1	1380.59	-

(d) Mean tardiness Excel

Figure 4.5. Risk 2 model outputs

Figures 4.5(c) and 4.5(d) showcase the arrival time at `Node 1` under two objective functions: mean lateness indicates an arrival at the beginning of the time window, while mean tardiness results in an arrival near the end. Notably, leveraging the mean tardiness objective, `Vehicle 1` effectively serviced `Node 2` before `Node 1`, ensuring timely service for both nodes. This demonstrates how diverse objective functions empower decision-makers to explore trade-offs and make well-informed choices.

Figure 4.6 shows that by accepting `Risk 3`, all UUVs within the scenario come under consideration during the optimization of routing schedule. We observe that the scheduling of LVs remains consistent across both objective functions, except for the service time assigned to the final UUV, `Node 2`. In the context of mean lateness, the servicing of `Node 2` was efficiently executed at approximately 1369 minutes (Figure 4.6(c)), whereas the execution of `Node 2`'s service for mean tardiness was deferred until the 1440-minute mark (Figure 4.6(d)). This difference occurs because there are no benefits in trying to arrive more quickly when calculating mean tardiness, which ultimately suggests that there isn't just one optimal solution.

When examining the scenario under `Risk 4`, no discernible variations in the results were found. Therefore, we chose not to include it in our discussion.

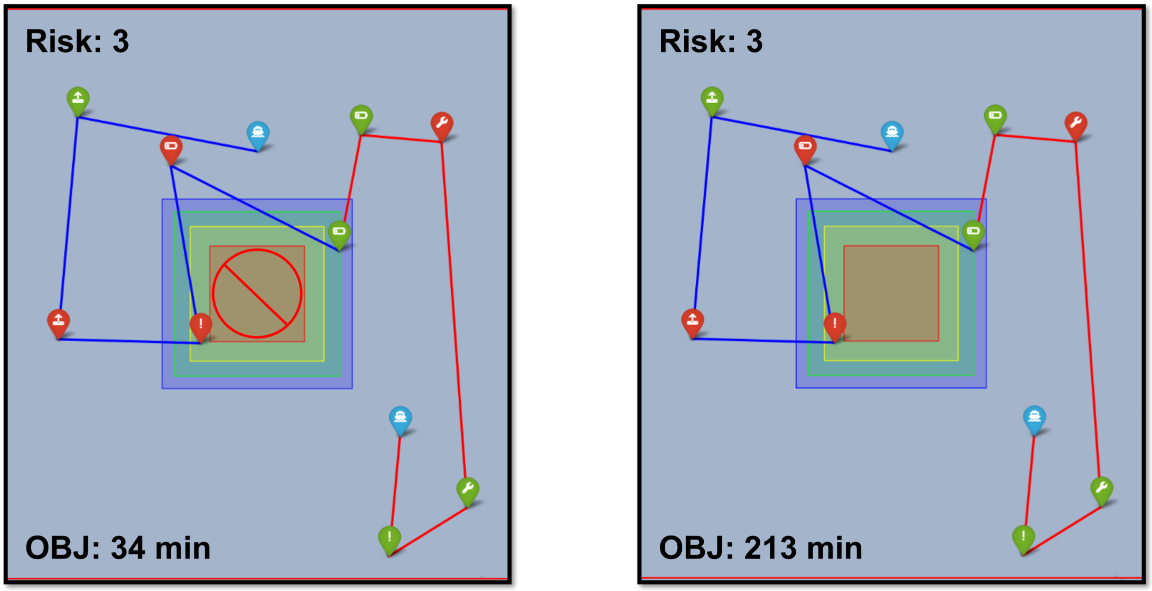

(a) Mean lateness map (b) Mean tardiness map

Vehicle	Node	Arrival Time	Departure Time
0	0	1	67.41
0	6	194	276
0	8	372.8	490.8
0	5	876.97	1086.61
0	3	1172	1245
0	2	1369.59	-
1	0	0.1	0.1
1	9	193.59	277.59
1	4	513.68	584.68
1	1	735.08	825.08
1	7	1015.62	1169.95
1	2	1369.59	-

(c) Mean lateness Excel

Vehicle	Node	Arrival Time	Departure Time
0	0	0.1	67.41
0	6	194	276
0	8	372.8	490.8
0	5	876.97	1086.61
0	3	1172	1315.41
0	2	1440	-
1	0	0.1	0.1
1	9	193.59	277.59
1	4	513.68	584.68
1	1	735.08	825.08
1	7	1015.62	1240.36
1	2	1440	-

(d), Mean tardiness Excel

Figure 4.6. Risk 3 model outputs

4.3 Optimality Gap Analysis

In this section, our aim is to evaluate the effectiveness and reliability of the solution derived from our model. This evaluation enables us to gauge the proximity of the obtained solution to the ideal, theoretical solution often referred to as the "optimal solution."

Considering that our previous scenario was somewhat simplistic, we decided to challenge our model by subjecting it to more intricate scenarios—some of which may pose challenges to timely resolution. To achieve this, we maintained fixed positions for four Launch LVs and three CAs, while introducing randomized parameters for fifteen UUVs within the predefined operating area, the randomization of the accepted level of risk. For the rest of this discussion, we define a scenario as a specific model run with randomized UUV parameters and level of accepted risk. The parameters attributed to LVs and CAs for this experimental setup are outlined in Tables 4.1 and 4.2.

Table 4.1. CA parameters (OGA)

Latitude	Longitude	CA Size (NM)
-1.06	107.38	38
-1.73	107.12	29
-0.49	107.9	35

Table 4.2. LV parameters (OGA)

Latitude	Longitude	Speed (Kt)
0.2	107.0	10
-2.0	108.0	12
-0.5	108.4	10
-0.68	106.8	12

4.3.1 Additional Constraints and Parameters

Throughout these simulations, we incorporated additional constraints into the model, prompting it to halt if the gap failed to narrow adequately within a 30-minute window, not to exceed two hours. Additionally, we introduced a straightforward heuristic instructing the model to initiate by assigning UUVs to the closest LV.

Within the Gurobi solver framework, we made adjustments to two model parameters with the aim of enhancing performance. These changes were prompted by the observation that while the model excelled in generating feasible solutions, it encountered challenges in expeditiously narrowing the gap in certain scenarios as the problem's scale increased.

The `MIPFocus` parameter allows you to direct the solver's focus towards specific aspects of the MILP problem, potentially enhancing solution quality or solution time. By default, this setting maintains a balance between discovering new solutions and refining the existing incumbent solution. In our case, this parameter was assigned a value of `3`, signaling the model to concentrate on enhancing the best objective-bound.

The second parameter modified was `Cuts`, the global cut aggressiveness setting. This parameter was also set to 3 for very aggressive cut generation. Additional parameter details are provided in the Gurobi Reference Manual (Gurobi Optimization, LLC 2023). With these settings in place, we anticipate that the model will efficiently close the optimality gap compared to its default configurations.

4.3.2 Simulation Results

Figures 4.7 and 4.8 are illustrative outputs of two such scenarios for the mean tardiness objectives. In Figure 4.7, the model swiftly attained a global optimal solution within just under five minutes (271 seconds). This can be attributed to the `Risk = 2` constraint, which prevented the servicing of four out of the 15 UUVs.

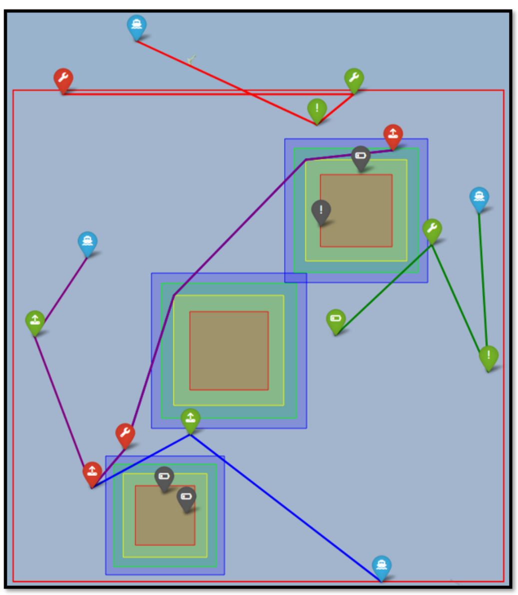

Figure 4.7. Optimal solution found

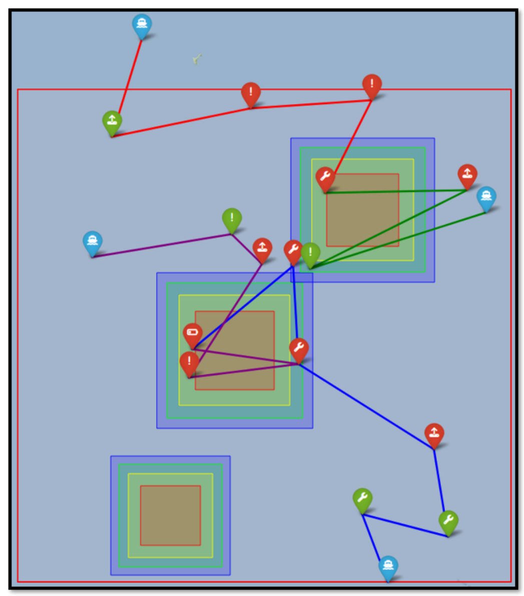

Figure 4.8. Optimal solution not found

In contrast, Figure 4.8 presents a scenario with `Risk = 4`, where the model encountered challenges in promptly narrowing the optimality gap. Over approximately one hour (3477 seconds), the gap was reduced to 58%. This discrepancy is due to the inclusion of all 15 UUVs in the optimization of vehicle schedules, increasing complexity. As is typical in scaling MILP problems, prolonged gap closure times align with a larger count of UUVs necessitating service, regardless of the associated risk level.

Figure 4.9 presents two plots depicting the optimality gap over time when the objective function is mean lateness. The upper plot consists of line graphs, each color representing the gap's progression over time for a specific scenario. In the lower plot, depicted as a box-plot, the average gap across all 100 scenarios is illustrated for each time point. The y-axis denotes the optimality gap as a percentage, while the x-axis represents time in seconds.

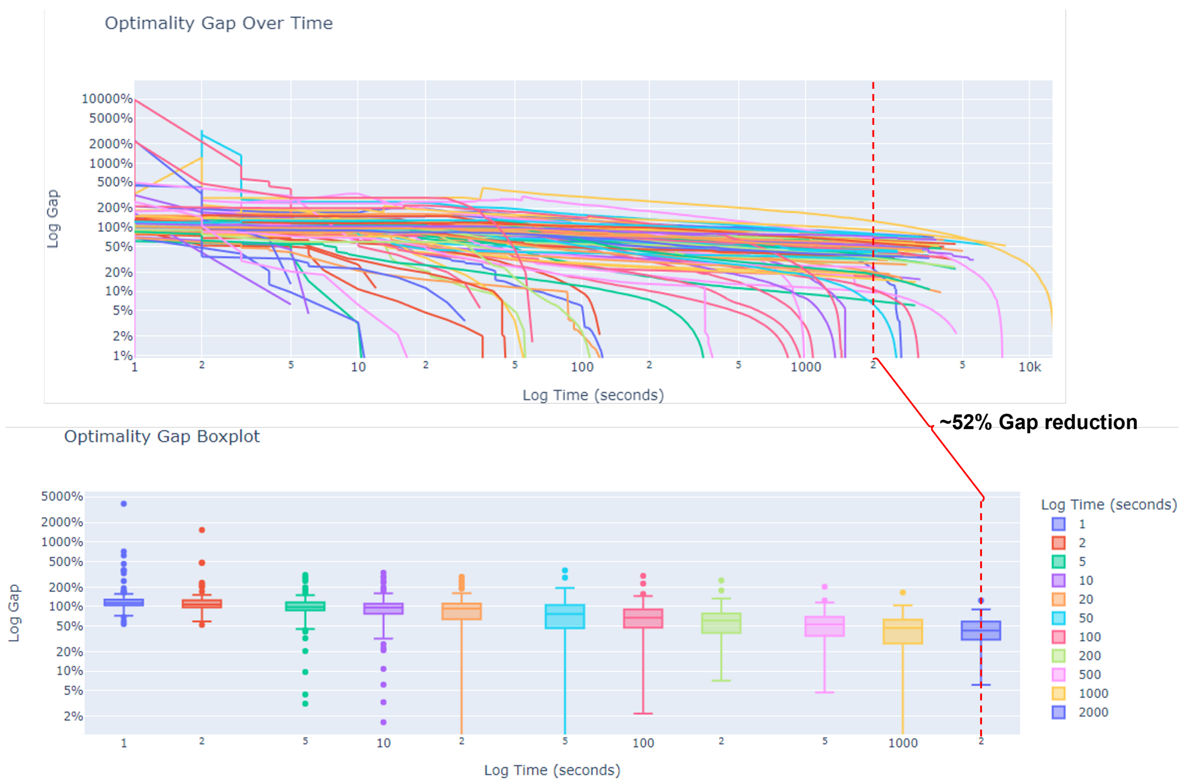

Figure 4.9. This figure depicts two plots illustrating the optimality gap over time for the mean lateness objective. The axes of the plots have been logarithmically scaled to condense the represented values. The x-axes represent time in seconds, while the y-axes represent gap as a percentage. The upper plot shows the optimality gap over time for each scenario, while the lower plot displays the five-number summary (minimum, first quartile, median, third quartile, maximum) for all scenarios at those time intervals.

Commencing from time zero, it is apparent that, in certain scenarios, the optimality gap exceeds 100%. Mean lateness, being a measure of how late tasks are completed relative to their due times, can take on negative values when tasks are completed early. This could lead to a situation where the optimality gap exceeds 100% if the algorithm initially struggles to find feasible solutions within the problem's constraints.

As time advances, observations from the box-plot demonstrate a consistent trend of gap reduction across all scenarios. Remarkably, in approximately one-third of the scenarios, a global optimal solution was identified within the model's designated runtime. Around the half-hour milestone (2000 seconds), indicated by the dashed red line, the average gap for the remaining scenarios settled at 48%. This indicates that on average, the model achieved a solution roughly 48% away from the theoretical best outcome within the stipulated timeframe. This insight holds valuable implications for decision-makers engaged in trade-off analysis between solution quality and computational time.

The same procedure applied to the mean tardiness objective yielded analogous plots, as illustrated in Figure 4.10. Two noteworthy distinctions between Figure 4.9 and Figure 4.10 came to light.

Firstly, the peak optimality gap did not surpass 100% in any scenario. As mentioned earlier, this can be attributed to the nature of the objective functions. Mean tardiness, measuring lateness positively, tends not to exhibit the same phenomenon of the optimality gap exceeding 100%, even amidst challenges faced by the algorithm in finding viable solutions.

Secondly, the model displayed a swifter gap closure rate, averaging at 33% around the half-hour mark for the remaining scenarios (Figure 4.10 box-plot). This potentially indicates that the model was more adept at seeking feasible solutions for the mean tardiness objective, resulting in a more favorable optimization process.

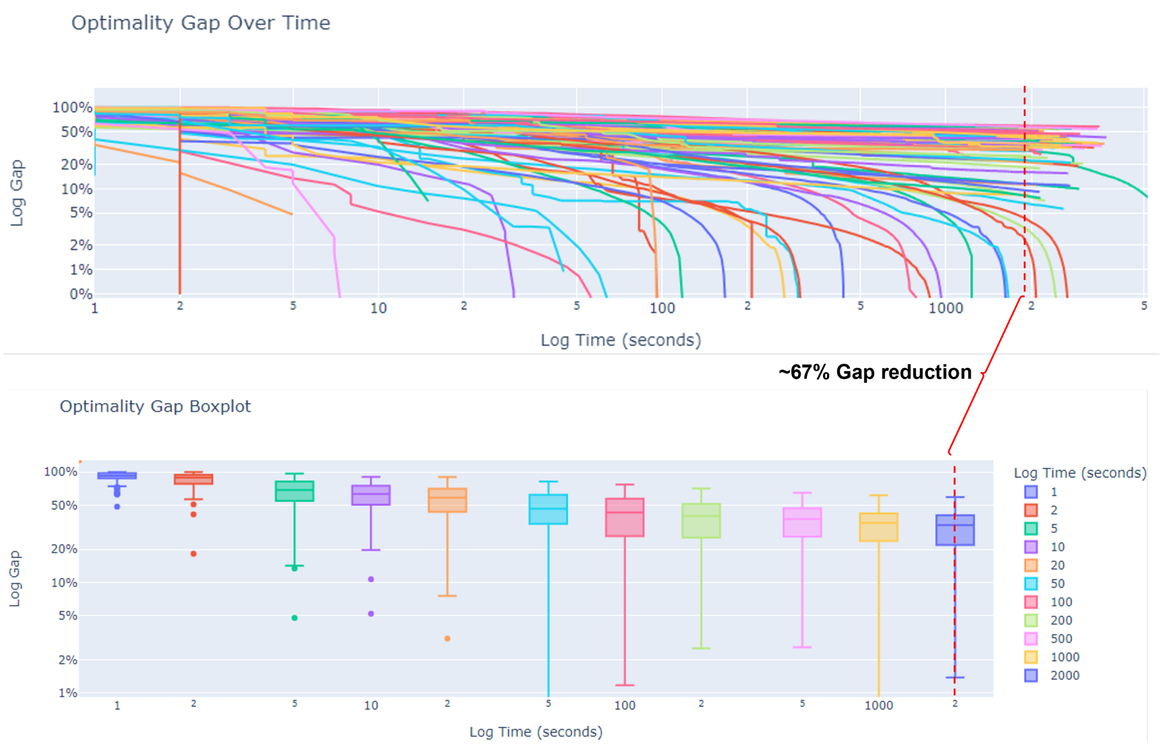

Figure 4.10. This figure depicts two plots illustrating the optimality gap over time for the mean tardiness objective. The axes of the plots have been logarithmically scaled to condense the represented values. The x-axes represent time in seconds, while the y-axes represent gap as a percentage. The upper plot shows the optimality gap over time for each scenario, while the lower plot displays the five-number summary (minimum, first quartile, median, third quartile, maximum) for all scenarios at those time intervals.

Observing the model's improved performance for the mean tardiness objective, we proceeded to explore the impact of altering certain parameters to assess any potential differences. Specifically, we set the `MIPFocus` parameter to `2`, thus directing the solver's attention toward achieving feasible solutions. For this experiment, we introduced 50 random scenarios, allowing the model to operate for eight hours within each scenario. The extension of runtime to eight hours was motivated by our suspicion that the gap closure might be less efficient than under the `MIPFocus` parameter set to `3`; however, this conjecture is not always the case.

Figure 4.11 displays two plots, akin to the ones mentioned earlier, presenting the outcomes of this particular experiment. When comparing the average gap reduction at approximately the half-hour point (2000 seconds), the average gap had decreased to 43%, which is 10% higher than when `MIPFocus` was set to `3`. For scenarios unresolved by the eight hour threshold, the average gap had actually increased to 46%. A noteworthy insight is that post the half-hour mark, unsolved scenarios exhibited wider gaps and a significant reduction in closure rate. In light of the problem's characteristics, opting for `MIPFocus = 3` emerges as a superior choice to expedite gap closure efficiently.

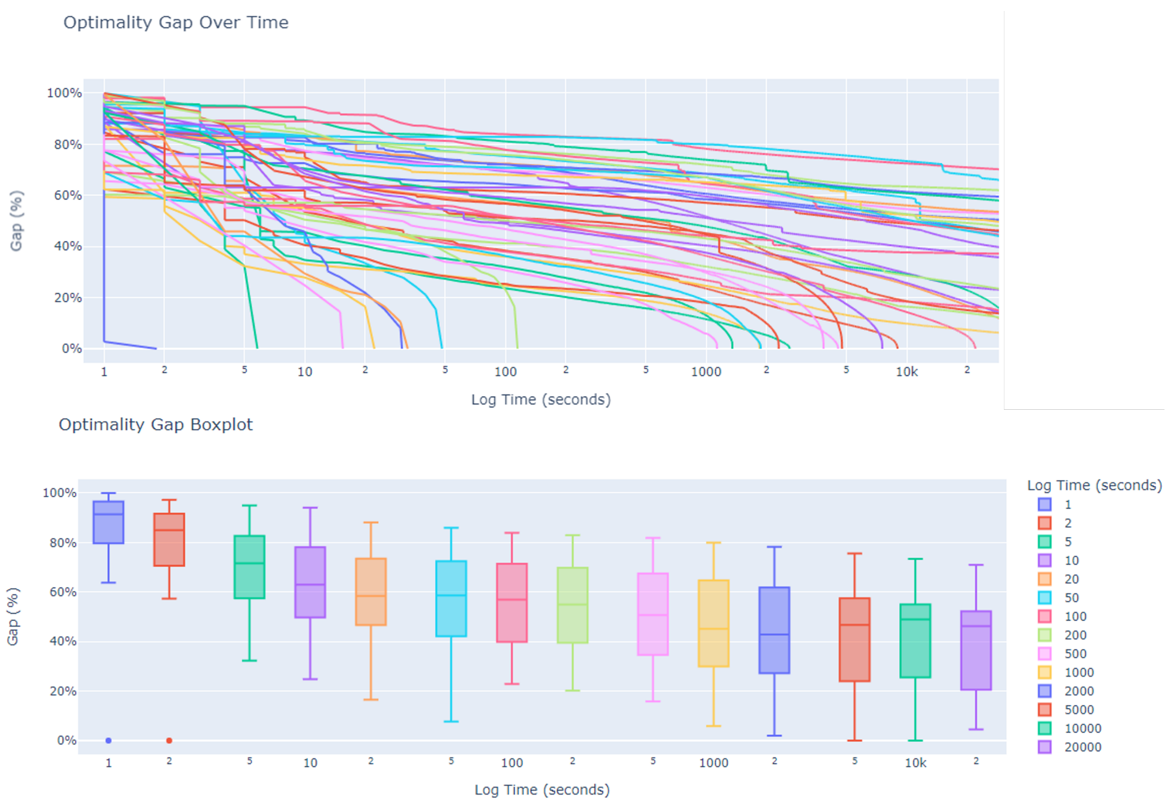

Figure 4.11. This figure displays two plots that illustrate the optimality gap over time for the mean tardiness objective, considering an `MIPFocus = 2` and 8-hour runtime. The x-axes (time in seconds) of the plots has been logarithmically scaled to compactly represent the values. The y-axes represents gap as a percentage. The upper plot shows the optimality gap over time for each scenario, while the lower plot displays the five-number summary (minimum, first quartile, median, third quartile, maximum) for all scenarios at those time intervals.

4.3.3 Summary

In this section, we evaluated our model's reliability and effectiveness by introducing randomized UUV parameters and risk levels. We incorporated additional constraints and heuristics while tuning Gurobi solver parameters to gauge performance changes. Notably, for the mean tardiness objective, significant improvements in performance have prompted us to further explore adjustments to the `MIPFocus` parameter. While revealing its potential for improvement, there are several unexplored parameters that show promise for further optimizing the model. To comprehensively explore the model's potential, it may be beneficial to fix a single scenario while tuning parameters.

CHAPTER 5:
Conclusion

5.1 Conclusion

The undersea warfare environment, much like other operational domains, is steadily progressing towards the integration of unmanned and autonomous vehicles. To keep up with evolving technology and address the logistics challenges inherent in these vehicles, it is necessary to identify and develop ways to support them in an optimized way. Within this context, this thesis offers an alternative perspective on scheduling and routing logistics vehicles to cater to the requirements of UUVs. By employing a hypothetical scenario, we envision a future operational setting characterized by a network of UUVs with diverse needs, necessitating service from a finite pool of LVs. Through our proof of concept, we not only showcase the applicability of our model in such scenarios but also provide visualizations to facilitate comprehension of the results.

To address the intricacies of operations in gray zone environments, we introduce CAs into the scenarios, proposing that decision-makers establish acceptable risk levels for enabling operations within CAs to a certain extent. Lastly, we analyze our model's performance concerning optimality gaps across different objective functions and provide recommendations for further model refinement.

5.2 Future Work

Several avenues are open for further advancing this work, including input parameter refinement, exploring alternative model objectives, and enhancing the capabilities of the provided tool.

Parameter Refinement

The infrastructure and resources required to facilitate the mission set outlined in this thesis are currently in a state of development, therefore it's important to acknowledge that the parameters assigned to UUVs and LVs might not be highly accurate at this stage. In order to enhance the accuracy and effectiveness of our model and tool, it is imperative that we consider forging partnerships with industry experts and contractors who are actively involved in the creation of this technology. Collaborative efforts with these stakeholders can offer insights, data, and expertise that will be instrumental in refining our approach and ensuring that our model aligns more closely with the practical requirements of future operations of this nature.

Alternative Model Objectives and Design

The objective functions implemented in the model for this thesis were simplistic for the purpose of demonstrating the proof of concept. However, to address more intricate real-world constraints like logistic vessel fuel consumption, a more comprehensive approach could involve implementing multi-objective functions. These functions could encompass minimizing the overall distance traveled alongside considerations of tardiness or lateness, thereby capturing a more nuanced representation of real-world operational constraints and objectives. Furthermore, in place of the linear scaling approach for UUV priorities, the objective function can be revised to more accurately account for the decision-maker's utility in ensuring timely servicing of UUVs with varying levels of priority.

Additionally, the model could be improved by integrating elements of online optimization. Online optimization entails the constant adjustment of decisions in response to real-time data updates. This approach holds particular relevance to our situation, given that operations within gray zone environments are characterized by their dynamic and unpredictable nature.

Finally, the application of decentralized optimization would be highly beneficial, wherein individual entities make decisions autonomously using their local information and interactions with their environment, absent of central authority directive. Within the context of the operational scenario underpinning this thesis, each LV would autonomously calculate its optimal schedule without access to the locations and actions of other LVs. This approach becomes particularly valuable in the context of gray zone environments, where communication might be compromised or entirely absent.

Enhancing Tool Capabilities

Currently, the tools developed through this study include an Excel spreadsheet and an interactive map that are customized for the specific objective function programmed into the model at run-time. As we integrate alternative objectives and configurations, there is a possibility to enhance the tool, allowing users to seamlessly transition between different objectives or even execute a selected set of objective functions. Looking ahead, a prospective initiative could revolve around designing a comprehensive tool that encompasses all the previously mentioned objective functions and designs, eliminating the requirement for intermediaries.

THIS PAGE INTENTIONALLY LEFT BLANK

List of References

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

Baesystems (2023) Riptide unmanned undersea vehicles. Accessed April 4, 2023, https://www.baesystems.com/en-us/product/riptide-uuv.

Bodin L, Golden B, Assad A, Ball M (1981) The state of the art in the routing and scheduling of vehicles and crews. Technical report, University of Maryland, College Park, MD, https://hdl.handle.net/2027/ien.35556021333117.

Cone S (2022) Casualty evacuation optimization in a conflicted environment. Master's thesis, Department of Operations Research, Naval Postgraduate School, Monterey, CA, https://hdl.handle.net/10945/71052.

Gallup S (2022) Distributed maritime logistics for theater undersea warfare. NPS-22-N355 Executive Summary, Naval Postgraduate School, Monterey, CA.

Gallup S (2023) LMACC. Naval Postgraduate School. Accessed July 10, 2023, https://nps.edu/web/lmacc/welcome.

Gallup S, Mun J (2020) The lightly manned autonomous combat capability (LMACC). Presentation, Naval Postgraduate School, Monterey, CA. https://dair.nps.edu/bitstream/123456789/4260/1/SYM-AM-20-123.pdf.

Garey M, Johnson D (1979) *Computers and Intractability: A Guide to the Theory of NP-Completeness* (W.H. Freeman, New York City, NY).

Gurobi Optimization, LLC (2023) *Gurobi Optimizer Reference Manual*. https://www.gurobi.com/wp-content/plugins/hd_documentations/documentation/10.0/refman.pdf.

Hill R (2022) NPS-22-N355: Distributed maritime logistics for theater undersea warfare. NRP Topic Portal. Accessed April 4, 2023, https://nps.edu/web/naval-research-program/topics/-/topics/lLTOSW5Vv9hI/20304.

Naval Technology (2022) Knifefish Unmanned Undersea Vehicle (UUV). Accessed April 4, 2023, https://www.naval-technology.com/projects/knifefish-unmanned-undersea-vehicle-uuv/.

Paulo T, Vigo D (2002) *The Vehicle Routing Problem* (SIAM, Philadelphia, PA), https://epubs.siam.org/doi/book/10.1137/1.9780898718515.

Rardin RL (2017) *Optimization in Operations Research* (Pearson, Hoboken, NJ).

Sorenson C (2020) Logistics in a contested environment: Network analysis and flow optimization. Master's thesis, Department of Applied Mathematics, Naval Postgraduate School, Monterey, CA, https://hdl.handle.net/10945/65444.

Starling C (2022) Today's wars are fought in the 'gray zone.' Here's everything you need to know about it. Atlantic Council. Accessed July 16, 2023, https://www.atlanticcouncil.org/blogs/new-atlanticist/todays-wars-are-fought-in-the-gray-zone-heres-everything-you-need-to-know-about-it/.

Initial Distribution List

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California