NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

LEVEL LOADING SURFACE SHIP MAINTENANCE AVAILABILITIES

by

Lauren G. O'Malley

June 2020

Thesis Advisor:	Kyle Y. Lin
Second Reader:	Emily M. Craparo

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE				*Form Approved OMB No. 0704-0188*

1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** June 2020	**3. REPORT TYPE AND DATES COVERED** Master's thesis

4. TITLE AND SUBTITLE LEVEL LOADING SURFACE SHIP MAINTENANCE AVAILABILITIES	**5. FUNDING NUMBERS** RNP5P
6. AUTHOR(S) Lauren G. O'Malley	
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) Naval Surface Force, U.S. Pacific Fleet	**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.

12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release. Distribution is unlimited.	**12b. DISTRIBUTION CODE** A

13. ABSTRACT (maximum 200 words)

Navy surface ships need to undergo regular maintenance pier side to meet fleet operational requirements. These maintenance jobs—commonly known as maintenance availabilities—are often contracted out to private shipyards in a ship's home port. While the Navy needs a maintenance schedule that meets fleet operational requirements, the shipyards prefer the workload to be leveled over time to sustain a trained and skilled workforce. Currently, surface ship maintenance scheduling is planned manually. This thesis develops a mixed integer linear programming model to produce an optimal surface ship maintenance schedule to account for two competing objectives: (1) level the workload over time in a regional port, and (2) minimize the schedule shift from fleet operational requirements. In a case study conducted in the Port of San Diego, the optimization model reduces workload fluctuation substantially over a 5-year period by slightly shifting the original maintenance schedule. The optimal schedule provides private shipyards with a more sustainable and predictable workload, which in turn reduces the risk of maintenance backlogs for the Navy.

14. SUBJECT TERMS ship maintenance, port loading, surface ships, fleet readiness, optimization	**15. NUMBER OF PAGES** 55
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
Unclassified	Unclassified	Unclassified	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.

LEVEL LOADING SURFACE SHIP MAINTENANCE AVAILABILITIES

Lauren G. O'Malley
Ensign, United States Navy
BS, United States Naval Academy, 2019

Submitted in partial fulfillment of the
requirements for the degree of

**MASTER OF SCIENCE IN APPLIED SCIENCE
(OPERATIONS RESEARCH)**

from the

**NAVAL POSTGRADUATE SCHOOL
June 2020**

Approved by: Kyle Y. Lin
Advisor

Emily M. Craparo
Second Reader

W. Matthew Carlyle
Chair, Department of Operations Research

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Navy surface ships need to undergo regular maintenance pier side to meet fleet operational requirements. These maintenance jobs—commonly known as maintenance availabilities—are often contracted out to private shipyards in a ship's home port. While the Navy needs a maintenance schedule that meets fleet operational requirements, the shipyards prefer the workload to be leveled over time to sustain a trained and skilled workforce. Currently, surface ship maintenance scheduling is planned manually. This thesis develops a mixed integer linear programming model to produce an optimal surface ship maintenance schedule to account for two competing objectives: (1) level the workload over time in a regional port, and (2) minimize the schedule shift from fleet operational requirements. In a case study conducted in the Port of San Diego, the optimization model reduces workload fluctuation substantially over a 5-year period by slightly shifting the original maintenance schedule. The optimal schedule provides private shipyards with a more sustainable and predictable workload, which in turn reduces the risk of maintenance backlogs for the Navy.

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

CLF	Combat Logistic Force
CM	continuous maintenance
CNO	Chief of Naval Operations
COMNAVSURFOR	Commander, Naval Surface Force
CR	continuing resolution
csv	comma separated values
FY	Fiscal Year
GAO	Government Accountability Office
MILP	mixed integer linear program
NAVSEA	Naval Sea Systems Command
NNN	Navy the Nation Needs
NPS	Naval Postgraduate School
NWRMC	Northwest Regional Maintenance Center
OFRP	Optimized Fleet Response Plan
PSI	Private Sector Improvement
PSO	Private Shipyard Optimization
RMC	Regional Maintenance Center
SSDSP	Surface Ship Drydock Schedule Planner
SWRMC	Southwest Regional Maintenance Center

THIS PAGE INTENTIONALLY LEFT BLANK

Executive Summary

Surface ships in the Navy need to undergo maintenance in a regional port on a regular basis. These maintenance jobs—commonly known as maintenance availabilities—are often contracted out to private shipyards. While the Navy needs ships to undergo maintenance at specific times to maximize fleet readiness, the private shipyards prefer a workload that has minimal fluctuation over time. Level loading the work allows private shipyards to train and maintain a skilled workforce, which in turn improves their chance to complete maintenance availabilities on time.

This thesis develops a port loading tool to help fleet planners create an optimal maintenance schedule that levels maintenance availabilities as much as possible while meeting the Navy's requirement for fleet readiness. The port loading tool uses mathematical optimization to account for two competing objectives: (1) minimize workload fluctuation over time in a regional port, and (2) minimize schedule shift of availabilities from the current fleet maintenance schedule. The output of the optimization model is a maintenance schedule that optimally balances workload fluctuation and schedule shift.

The port loading model is formulated as a mixed integer linear program. The model is coded in the Python programming language with the Pyomo optimization library. It takes a minute or two on a personal computer to process data, build the model, and derive an optimal schedule. The tool opens in a web browser, and comes with a graphical interface that allows the user to upload spreadsheet files required to run the optimization model. Upon completion of the program, the user receives an optimal maintenance schedule displayed in plots and accessible through spreadsheet files.

The port loading tool provides a platform for fleet planners to deal with the intricate requirements of the Navy's operations while taking into account the need for a leveled workload at private shipyards. There are several potential applications for this tool.

1. Allow fleet planners to generate multiple long-term schedules that balance workload fluctuation and schedule shift for maintenance availabilities.
2. Identify periods of congestion for private shipyards in terms of workload requirements and justify alterations to the current fleet schedule.

3. Explore a variety of scenarios with different restrictions from the perspective of both the Navy and private shipyards.
4. Create an opportunity to improve current contracting strategies.

We explore a case study for the Port of San Diego to demonstrate our port loading tool. San Diego is selected based on the congestion within the port. The current fleet schedule is provided by Naval Sea Systems Command (NAVSEA), and the case study runs from October 1, 2019 to September 30, 2024. The current schedule has a 12.93% workload fluctuation over the 5-year period. The drastic fluctuation makes it difficult for the private shipyards to train and maintain a skilled workforce, which increases risk of schedule delays. By using the port loading tool, the optimal schedule reduces the workload fluctuation from 12.93% in the original schedule to 3.16%. All schedule shifts are within the allowable window required by the fleet, with an average shift of 10.26% for each availability. The tool also allows the user to assign different weights to workload fluctuation and schedule shift to find a schedule that best balances the two competing objectives.

This thesis makes an important contribution in level loading maintenance availabilities in a regional port. The tool provides quantitative justification for minor adjustments to the current fleet schedule to level the workload over time. The fleet planners can use the port loading tool to create a win-win maintenance schedule for both private shipyards and the Navy. A thriving industry base in a regional port is key to ensure the long-term health and readiness of the fleet.

Acknowledgments

This thesis would not have been possible without the support of many people in both my professional and personal life. Thank you to my professors, mentors and family who have continually shown me support throughout my academic career.

I would especially like to thank Dr. Kyle Lin for his expertise and guidance throughout the development and completion of this thesis. I have learned so much in terms of mathematical formulation and application and technical writing that helped advance my skill set in all aspects of life. This thesis would not have been possible without his insight and patience.

I would also like to thank Dr. Emily Craparo for her analytic input. She challenged me to think beyond the bounds of the problem, and had an attention to detail that took this thesis to the next level.

Finally, I would like to thank my family and friends. To my parents, thank you for listening to my math jargon as I explicitly explained my thesis development and application. To Jared, for showing me continual love and support. To my brother and sister, for reminding me of life outside of the academic world, and to my NPS cohort, for challenging me in academics and reminding me we are all in this together. Thank you to all.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:
Introduction

Surface ships routinely undergo maintenance during life cycles to meet the Navy's operational and readiness requirements. The responsibility of completing ship maintenance jobs—commonly known as maintenance availabilities—rests on a number of different entities and organizations. Ship maintenance is classified at three different levels: (1) Organizational maintenance that the ship's crew routinely performs, (2) Intermediate maintenance that requires the assistance of additional personnel from the contractor workforce, or (3) Depot-level maintenance that personnel conduct at publicly or privately owned ship yards (Maurer 2019). Ship schedules embed maintenance availabilities requiring personnel outside of the ship's crew well in advance to avoid overlaps or delays. Additional time spent in an availability period restricts the time spent training, operating, and meeting the requirements for sustaining the Navy the Nation Needs (Naval Sea Systems Command 2020). With the Navy continually growing and modernizing its surface fleet, fleet planners identify the timely completion of maintenance availabilities as a critical element to success. This puts a burden on both the military and private shipyards to maintain the ship schedule parameters, workforce requirements, and maintenance obligations.

One way to mitigate this burden is to develop a port loading tool to level workloads for maintenance availabilities in a port over time. Specifically, port loading is the technique of matching maintenance requirements with the capabilities of shipyards contracted out to complete ship maintenance. In terms of workload requirements, when the fleet schedule match what the workforce at private shipyards can support, they increase the chances that availabilities complete maintenance on time by reducing workload fluctuation. Thus, leveling workload over time results in private shipyards maintaining a healthy and sustainable workforce available to conduct ship maintenance.

Section 1.1 addresses the background of port loading maintenance and how it is currently scheduled. Section 1.2 discusses our contribution to efficiently schedule pier side maintenance. Section 1.3 reviews similar works related to this thesis, and section 1.4 outlines the framework for the remainder of the thesis.

1.1 Background

1.1.1 The Motivation

On the forefront of the minds of Navy leaders today is how the current fleet can transform to the fleet of the future with improvements in maintenance schedules. Historically, the United States Navy has been incomparable across the globe. Under direction of Commander Naval Surface Force (COMNAVSURFOR), the fleet has maintained and managed more ships, personnel, and surface area of the seas than any other nation. The fleet as it stands now has an inventory of 301 vessels, and Navy leaders plan for the fleet to grow to 355 vessels by fiscal year 2034 (Naval Sea Systems Command 2020). According to the plan, the breakdown of the vessels in 2034 is as follows:

1. Nuclear Powered Vessels:
    - Ballistic Missile Submarines (12)
    - Aircraft carriers (12)
    - Attack submarines (66)
2. Conventional Surface Ships:
    - Large, Multi-Mission Surface Combatants (104)
    - Small, Multi-Role, Surface Combatants (52)
    - Amphibious Warfare Ships (38)
    - Combat Logistic Force (CLF) (32)
    - Command and Support (39)

All nuclear powered vessels require maintenance at public shipyards owned by the Navy. Private shipyards are typically responsible for the maintenance of conventional surface ships.

When predicting what the future may hold for the Naval Fleet, it is important to recognize what an operational fleet of this size demands of fleet support systems. COMNAVSURFOR is responsible for manning, training, and equipping the Navy's surface fleet. Together with Region Maintenance Centers (RMC), COMNAVSURFOR works to allocate budgets and resources so that the fleet is able to perform operationally while providing support for scheduling ship repair and maintenance periods (Northrup 2015). Naval Sea Systems Command (NAVSEA) is responsible for scheduling ship maintenance availabilities with fleet maintenance officers, and the RMC coordinate intermediate maintenance and execute all contracts with private shipyards (Northrup 2015). Operational readiness and ship maintenance depends on the coordination between all of these entities.

Recent studies focused on fleet readiness determined that maintenance backlogs is one of the biggest factors that reduce operational readiness across the fleet. High-tempo operations, unpredictable workload fluctuation, and delays in maintenance that continue to build upon themselves result in backlogs across the fleet. From Fiscal Years (FY) 2011 to 2014, 72% of surface combatant ships experienced maintenance backlogs that detracted from operational employability (Pendleton 2016). Beginning in FY14, the Navy began implementing the Optimized Fleet Response Plan (OFRP) which "seeks to maximize employability while preserving maintenance and training with continuity in ship leadership and carrier strike group assignments, and restoring operational and personnel tempos back to acceptable levels" (Pendleton 2016, p. 1) Under this new policy, however, the Navy continues to face maintenance delays, which reduces the Navy's ability to perform training and operations around the world. In order to grow in size and effectively remain highly operational as the world's superpower, the Navy needs to make up for previous maintenance backlogs and optimize schedules in the future.

Operations depend on mission capable ships that rely on routine and punctual ship maintenance. If a ship does not complete its maintenance on time, then other ships are burdened with the responsibility of carrying out additional operational duty. Most of the ships tasked out to operational roles are already serving out at sea, which leads to added stress on both the physical systems and the crew not accounted for in initial ship schedule. The additional stress can cause unanticipated issues down the road, which can potentially be another source of maintenance delays that perpetuates the problem of backlogs further. Backlogs hinder training and operations and are extremely costly to the fleet as a whole. A successful, operational fleet is not possible without the prompt completion of ship maintenance.

Surface ships are contracted out to private shipyards within a specified region when maintenance tasks are beyond the crew's capabilities. Thus, private shipyards play a critical role in maintenance on conventional vessels in the fleet. These vessels include Cruisers, Destroyers, Amphibious Assault ships, and Military Sealift Command ships (Pendleton 2016), which total to 220 ships currently, and these vessels will be referred to as surface ships for the remainder of this thesis. In conducting maintenance for the fleet of surface ships, private shipyards are often constrained by a limited supply of resources including personnel, facilities, and equipment available. Typically, the schedule of maintenance availabilities is made without considering these constraints, which leads to unrealistic expectations on private shipyards to adhere to the Navy's projected schedule.

An unrealistic maintenance schedule may lead to a number of constraints for private shipyards. Examples of these constraints include unanticipated maintenance requirements, workforce inexperience, and workload fluctuation (Pendleton 2016). The workload required to complete maintenance availabilities fluctuates at an uncontrollable and unpredictable rate. Laying off and rehiring workers in the private shipyard industry is common with peaks and valleys existing in the demand of workers. Industry experts stated, "workers not recalled within approximately 90 days would migrate to other areas or other industries and seldom return" (Martin et al. 2017, p. 53). While there is usually a number of workers to hire without the necessary training needed to complete maintenance, the labor pool of experienced workers to consider for rehiring is often left depleted. According to industry officials from the Government Accountability Office (GAO), "These cycles result in unsustainable lows followed by potentially unmanageable highs in workload that they expect will eventually erode the ship repair industrial base's skilled workforce" (Pendleton 2016, p. 29). To top that, about 32% of shipyard employees have less than five years of experience, which further undercuts the potential to complete ship maintenance on time (Pendleton 2016). Subject matter experts from private shipyards hypothesize that the combination of factors, such as inadequate planning for resources, unpredictable work, quantity of overtime, and quantity of work stoppages, contribute to the perpetuating cycle of maintenance backlogs (Caprio 2012).

The Navy recognizes the importance of finding a solution to maintenance backlogs in order to meet both current and future operational requirements. A solution will enable training and operations to return to a manageable schedule for surface ships, and it will also save the fleet's budget by completing maintenance regularly and on time.

1.1.2 The Challenges

The Navy is looking into a number of changes that could help relieve the problem of maintenance delays. Some of these efforts include (Maurer 2019):

- Revising the size of ship crews: The Navy is working to add more personnel to its surface ships to ensure vessels have enough personnel on board to properly conduct operations and maintenance routinely.
- Hiring additional workers at shipyards: The number of workers at public shipyards have increased from 30,600 in 2014 to 37,400 in 2019. The additional workers are helpful, but these numbers only reflect what the Navy is able to do at public shipyards, and does not provide any resolution to the workforce at private shipyards.
- Performance to Plan: A study group formed to focus on better understanding maintenance challenges and the capacity needs of the future. This analytic effort is looking at improving the accuracy of forecast models currently used.

These efforts will help surface ships effectively complete organizational maintenance, and provide insights on potential solutions to the challenges faced in scheduling availabilities for intermediate and depot level maintenance.

The Navy is also taking specific steps towards change with regards to the maintenance delays at private shipyards. As aforementioned, Regional Maintenance Centers (RMC) are the primary actors in managing, overseeing, and carrying out contracts with private shipyards (Naval Sea Systems Command 2020). In the FY20 Long-Range Plan for Maintenance and Modernization of Naval Vessels report to Congress, it was stated that fleet maintenance schedulers, "continuously balance operational commitments against engineered maintenance periodicity and industrial base constraints to develop an executable maintenance and modernization schedule" (Naval Sea Systems Command 2020, p. 2). Surface ship maintenance availabilities are currently scheduled manually to meet the Navy's fleet readiness requirements.

Scheduling these availabilities manually takes hours, if not days, of tedious work. The process consists of copying and pasting multiple values to and from different Excel spreadsheets for each availability. The process combines a series of three tools programmed in Excel. The first tool fits a rough estimate of a labor curve with a positively skewed distribution for any availability length. The output of this tool is then copied into a different Excel sheet to obtain values of monthly man-days that meet the specified target total of man-days over the time horizon. These values of monthly man-days are then copied into a forecasting tool that combines the information from all availabilities scheduled to plot the workload estimate for private shipyards.

While this process functions to forecast work loads, the time taken to complete this process grows excessively for multiple availabilities over a longer time horizon. Additionally, the manual process is unadaptable and error prone highlighting the need for a better solution.

Beyond the solutions aforementioned, the Navy is also developing Private Shipyard Optimization (PSO) and Private Sector Improvement (PSI) programs that are anticipated to identify and eliminate challenges in the private shipyard industry. The end goal in developing these programs is to ensure affordable and on time delivery of surface ships from maintenance periods. The PSO program takes a broad viewpoint on investigating the optimal placement of facilities and major equipment necessary for maintenance, while the PSI program addresses workload stability, governance, contracting and process optimization (Naval Sea Systems Command 2020).

In order to see these plans through, maintenance and modernization requirements for the fleet of the future must be fully funded and efficiently executed. In a statement to the Armed Services Committee, the Honorable James Geurts and Vice Admiral Thomas Moore said, "Beginning FY20 under a continuing resolution (CR) introduces uncertainty, as the Navy attempts to execute work planned for the current year with funding based on last year's budget" (Geurts and Moore 2019, p. 2). The lack of certainty in the budget hinders the Navy's ability to schedule availabilities with private shipyards. This in turn limits the shipyard's ability to forecast and plan for future workloads. Without proper planning, peaks and valleys continue to occur in the workload, which creates instability in the skilled workforce of a particular port. As the CR continues further into the year, the effects become more pronounced, which significantly impacts planned work on both coasts (Geurts and Moore 2019).

All of these efforts are moving in the right direction of solving the problem of maintenance backlogs. However, delays continue to persist while several mitigation strategies are estimated to require years of continued management and significant investments before major changes are seen (Maurer 2019). To help contribute to these efforts efficiently and effectively, this thesis aims to create a port loading tool to be used by fleet planners. The tool is based on a mixed integer linear programming (MILP) to optimally schedule ship maintenance availabilities within the constraints of the current fleet schedule and those faced by private shipyards. Tactically, the tool can be used to better plan for individual surface ships maintenance schedules, which allows ships to properly train and perform without interference of maintenance delays. Operationally, the tool can help sustain a level loaded workforce at private shipyards in home ports across the nation. Strategically, when world events cause notable changes to fleet activity, the current fleet plan can be reassessed instantaneously to find a new optimal solution.

1.2 Our Contribution

Optimizing port loading is one way to mitigate the risk when resources required to complete maintenance availabilities exceed the capacities in a specified home port. Port loading is the technique of matching maintenance requirements with the capabilities of shipyards contracted out to complete ship maintenance. The contribution of this thesis is to develop a port loading tool that the Navy can use to schedule availabilities while level loading the work required of private shipyards in a port over time.

The scheduling tool will account for two competing objectives. The first objective is to minimize any deviation from the level workload across the planning horizon, which allows private shipyards to better sustain a skilled and constant workforce. The second objective is to minimize the schedule shift from fleet operational requirements, which allows ship schedules to remain highly operational and consistent with the needs of the Navy. Together, these competing objectives mitigate the risk of maintenance delays, which in turn increase operational readiness across the fleet.

The tool is developed as a mixed integer linear program (MILP), and it is implemented using the Python Programming Language (van Rossum 1991) and uses Pyomo (Hart et al. 2008) as an optimization tool. The input data reflects the current fleet schedule to include

- Availability identification
- Availability start date
- Availability end date
- Availability's projected labor

With this input data, the optimization tool adjusts current availability start dates in its allowable window in order to level load the projected workloads with minimal schedule shift. The output of the tool is an ideal, executable schedule for each availability in a specified port.

The port loading optimization tool has several applications:

- Produces an executable schedule considering both fleet requirements and private shipyard constraints.
- Allows the Navy to identify peaks and valleys in workloads for private shipyards well in advance to better plan and forecast for the future.
- Justifies any need for alterations in ship schedules.
- Provides a tool to explore a variety of scenarios implementing different restrictions from both the perspective of the Navy and private shipyards.
- Creates opportunity to improve the current contracting strategies.

The port loading optimization tool comes with a simple user interface, which allows an end user to run the model with a few clicks. The user interface is based on an application development library in Python called Dash (Parmer 2016). This interface allows users to upload the spreadsheet files required to run the optimization tool and receive output files from the program. The output files include a comma separated values (csv) sheet with optimal start dates for each availability and a stack area chart showing the distribution of labor required throughout the planning horizon. This tool is anticipated to be paired with the Surface Ship Drydock Schedule Planner (SSDSP) tool, which is a quantitative drydock loading model (Hilliard 2019). Together, these tools will allow fleet planners to evaluate capacities and capabilities of private shipyards and commercial drydocks when making decisions on the schedules of surface ships.

1.3 Related Works

Optimally scheduling ship maintenance has been a field of study in a number of government organizations. In the early 1990s, Brown (1992) created the Naval Shipyard Optimal Drydock Loading and Capacity Utilization Model. The goal of Brown's work was to create a tool that could study drydock capacity utilization in various scenarios. The idea was prompted by the Assistant Secretary of the Navy in a time when the fleet was shrinking in size and the

Navy's budget declining. Working groups were formed to develop alternatives under this new lens of the Navy to satisfy a smaller fleet's maintenance requirements, minimize excess workload capacity for maintenance workforce, and determine the breakpoints in drydock requirements. With the current process of scheduling taking up to two weeks of tedious manual work, the need of an efficient optimization tool became apparent as decision makers were curious in approximately twenty-five drydock loading plans. Brown came to the rescue by utilizing a mixed-integer linear program to maximize drydock utilization by finding the proper ship-to-dock assignments that met the constraints of ship schedules, drydock capabilities, drydock current loads, and drydock preventative maintenance. Brown's work set the foundation for future researchers to investigate optimization tools that can quickly expose what can be done to save time, money, and resources in coordinating fleet wide maintenance availabilities.

Recently, Schaefer (2017) focused on the port loading aspect of maintenance availability scheduling in private shipyards. His thesis aimed to estimate maintenance availability completion time with the projected labor executed during the planning horizon for each trade skill required. The model inputs include the shipyard's labor capacity, maintenance availability start dates and duration, and the total labor required for a particular trade skill over the course of the availability. The implementation of this model allowed decision makers to have a realistic estimation of when ship maintenance could be completed, what labor capacity was feasible for a particular trade skill in a port, and identified risks in execution before bidding and schedule a contract for maintenance availabilities.

Hilliard (2019) created a Surface Ship Drydock Schedule Planner (SSDSP) to optimally use commercial drydocks for maintenance availabilities. The tool formulates a mixed integer linear program to convert a surface ship maintenance schedule that may or may not be feasible due to commercial drydock constraints to an executable schedule that maximizes drydock utilization with minimal schedule shift to the existing schedule. This tool provides decision makers with the insight of potential inefficiencies, and allows decision makers to test different mitigation strategies to minimize delays in ship maintenance.

The work done by all of these past researchers provided great insights and solutions to the challenges in scheduling surface ships' maintenance availabilities. My thesis aims to build off of the work done in this field to expand the potential of using optimal scheduling across all aspects of maintenance in the fleet.

1.4 Thesis Outline

The rest of this thesis proceeds as follows. Chapter 2 introduces the methodology of the level loading optimization model. Chapter 3 conducts a case study of the home port of San Diego as this port is the most congested with surface ships. Chapter 4 provides concluding thoughts and a discussion on the way ahead in surface ship maintenance scheduling.

CHAPTER 2:
Level Loading Optimization

Surface ships in the Navy need to undergo regular maintenance pier side in order to meet fleet operational requirements. These maintenance jobs—referred to as maintenance availabilities—are contracted out to private shipyards in a ship's home port. Typically, the workload required to complete pier side maintenance fluctuates over time. This fluctuation causes a problem for private shipyards to maintain a trained and skillful workforce, which ultimately can lead to ship maintenance delays.

This chapter presents an optimization model that produces an executable schedule in order to meet fleet operational requirements while minimizing fluctuation from a leveled workload within a home port. The leveled workload at a home port is calculated as the average workload over a specified planning horizon. By maintaining a leveled workload, the risk in maintenance delays is reduced.

Section 2.1 introduces the optimization model, which assigns start dates to availabilities in a specified home port. Section 2.3 introduces additional considerations to make a more realistic optimization model, which fulfills the criteria of the two competing objectives.

2.1 Optimal Level Loading

This section introduces a mixed integer linear programming model to achieve level loading at a homeport. Each availability in the planning horizon is allowed to start during a predetermined time window. The objective is to select a starting date for each availability so that the projected workload required in each time period (such as a month) are as close to each other as possible.

Indices and Sets

$a \in A$	Maintenance availabilities.
$s \in S_a$	Discrete dates in the planning horizon that will refer to the potential starting dates for each availability a.

$t \in T$ — Monthly time periods over the planning horizon. The length of a period is determined by the calendar month and year.

Data

LVL_t — The ideal leveled workload in period t. To achieve level loading, a common choice is to distribute the total workload evenly throughout the time horizon.

TGT_t — The target workload in period t, which is LVL_t minus the workload pre-allocated to period t that is not subject to change, such as continuous maintenance and availabilities that are already in progress.

CAP_t — The labor capacity in the port at time t.

$\mathrm{LBR}_{a,s,t}$ — The workload projected for availability a in period t if the availability starts on date s. This quantity can be calculated based on the labor curve for the type and duration of the availability.

Decision Variables

$x_{a,s}$ — Binary, 1 if availability a is scheduled to start on date s.

y_t — Numeric value indicating the workload projected in period t above the target workload TGT_t.

z_t — Numeric value indicating the workload projected in period t below the target workload TGT_t.

u_t — Numeric value indicating the workload projected in period t above the labor capacity CAP_t.

Formulation

Throughout this thesis, any time there is a summation over an index, it is over the entire set unless otherwise noted.

$$
\begin{aligned}
\min \quad & \sum_t (y_t + z_t) + C \sum_t u_t && (2.1)\\
\text{s.t.} \quad & \sum_{s \in S_a} x_{a,s} = 1, \qquad \forall a \in A && (2.2)\\
& y_t \geq \sum_{a,s} x_{a,s}\ \mathrm{LBR}_{a,s,t} - \mathrm{TGT}_t, \qquad \forall t \in T && (2.3)\\
& z_t \geq \mathrm{TGT}_t - \sum_{a,s} x_{a,s}\ \mathrm{LBR}_{a,s,t}, \qquad \forall t \in T && (2.4)\\
& u_t \geq \sum_{a,s} x_{a,s}\ \mathrm{LBR}_{a,s} - \mathrm{CAP}_t, \qquad \forall t \in T && (2.5)\\
& x_{a,s} \in \{0,1\}, \qquad \forall a \in A,\ s \in S_a && (2.6)\\
& y_t \geq 0, \qquad \forall t \in T && (2.7)\\
& z_t \geq 0, \qquad \forall t \in T && (2.8)\\
& u_t \geq 0, \qquad \forall t \in T && (2.9)
\end{aligned}
$$

In the objective function (2.1), the first summation consists of two terms, y_t and z_t. For each time period t, the first term y_t refers to the projected workload above the target workload, while the second term z_t refers to the projected workload below it. The sum of these two terms is therefore the deviation of the projected workload from the target workload in a time period. The second summation has one term, u_t, which is the amount of workload projected in period t above the labor capacity, CAP_t. The objective function minimizes the total deviation of projected workload from the target workload over the entire planning horizon and makes it undesirable to exceed the labor capacity for all time periods in T.

Constraint (2.2) ensures each availability will start on one of its feasible starting dates. Constraints (2.3) and (2.7) together enforce

$$
y_t = \max \left\{ 0, \sum_{a,s} x_{a,s} \mathrm{LBR}_{a,s,t} - \mathrm{TGT}_t \right\},
$$

which is the amount of workload projected in period t above the target workload TGT_t.

Constraints (2.4) and (2.8) together enforce

$$z_t = \max\left\{0, \mathrm{TGT}_t - \sum_{a,s} x_{a,s}\mathrm{LBR}_{a,s,t}\right\},$$

which is the amount of workload projected in period t below the target workload TGT_t. For $t \in T$, at most one variable between y_t and z_t can be positive, and $y_t + z_t$ in the objective function represents the deviation of the workload required in period t from the target workload TGT_t.

The straightforward manner to ensure that the workload projection in a period does not exceed its labor capacity is to impose a new constraint

$$\sum_{a,s} x_{a,s}\ \mathrm{LBR}_{a,s} \le \mathrm{CAP}_t, \qquad \forall t \in T$$

The disadvantage of this constraint, however, is that the resulting MILP may become infeasible. To ensure that the MILP always has a feasible solution, we create an elastic variable u_t to denote the workload project above the labor capacity and heavily penalize it in the objective function. To do so, we add the following constraint

$$u_t \ge \sum_{a,s} x_{a,s}\ \mathrm{LBR}_{a,s} - \mathrm{CAP}_t, \qquad \forall t \in T$$

and change the objective function to

$$\sum_{t\in T}(y_t + z_t) + C\sum_{t\in T} u_t \tag{2.10}$$

where $C \gg 1$ is a large constant to heavily penalize any workload that exceeds the labor capacity, CAP_t. Consequently, the solver is strongly encouraged to return a feasible solution with $\sum_{t\in T} u_t = 0$ if there is one. In numerical testing, setting $C = 100$ and setting $C = 1000$ essentially produce the same result; we recommend setting $C = 1000$.

2.2 Running the Model

The size of the MILP grows with the number of availabilities, the planning horizon, and the flexibility of each availability's starting date. For a typical application, the planning horizon ranges between 2 to 10 years, while the number of availabilities may range between a few dozen to a few hundred. The starting date of an availability may be allowed to shift left or right for up to 30 days if it is scheduled to start within two years, or up to 45 days if it is scheduled to start after two years. While the MILP solves quickly when the problem size is small, the time needed to compute an optimal solution may grow substantially as the problem size grows.

One remedy to compute the optimal solution more efficiently is to run the MILP twice. In the first run, we let S_a include only a few selected dates evenly spread out in the feasible window for each availability $a \in A$ to substantially cut down the problem size. The solution obtained from this first run would produce a rough estimate of the optimal schedule. In the second run, we let S_a include the several dates surrounding the optimal solution for availability a produced from the first run, for $a \in A$. For example, if availability a is allowed to start between 3/15 and 5/15, then in the first run we can assign

$$S_a = \{3/16, 3/22, 3/28, 4/3, \ldots, 5/3, 5/9, 5/15\}.$$

Suppose the optimal starting date for availability a from the first run is 3/28, then in the second run we assign

$$S_a = \{3/23, 3/24, 3/25, \ldots, 3/31, 4/1, 4/2\}.$$

The rationale of this approach is that the labor curve is usually smooth and unimodal, so that the objective function in the MILP is also smooth and unimodal in each availability's starting date. Therefore, recovering an optimal solution in these two stages performs well.

If we decide to run the MILP model twice to reduce overall computation time, then the overall computation time will typically be the smallest if each run of the MILP model has roughly the same size. Let n denote the length of availability a's feasible time window, and m the number of dates selected in the first run, so the interval between consecutive selected dates is about n/m. Because in the second run we need S_a to cover twice the interval, or

$2n/m$, it is most efficient to choose m such that

$$m = \frac{2n}{m}, \qquad \text{or} \qquad m = \sqrt{2n}.$$

For example, if availability a is allowed to shift left or right for up to 30 days, and we decide to run the model for 2 times, then $n = 30 + 1 + 30 = 61$, so $m = \sqrt{2 \times 61} \approx 11$ dates evenly spread out to cover the entire window of 61 days. In other words, the selected dates are about $61/11 \approx 6$ days apart.

This methodology can be expanded to three or more runs of the MILP model in a similar manner if so desired. The marginal benefit of saving computation time, however, is most pronounced when going from 1 run to 2 runs.

2.3 Trade off Between Workload Fluctuation and Schedule Shift

In the formulation in Section 2.1, each availability is allowed to start in a certain time window and any starting date is as good as any other in the window. This section extends the objective function in the MILP model to incorporate schedule shift. To do so, we introduce the following new data:

Data

$\mathrm{SFT}_{a,s}$	The penalty incurred due to schedule shift if availability a starts on date s.
L_a	Length of availability a.

To encourage an optimal starting date close to the ideal starting date, we write $\mathrm{SFT}_{a,s}$ as the number of days availability a is shifted from its ideal starting date, if starts on date s instead. The penalty incurred from shifting the start date of availability a can be computed as

$$\sum_{s \in S_a} x_{a,s}\ \mathrm{SFT}_{a,s} \tag{2.11}$$

There are two competing objectives. The one in (2.10) tries to minimize workload fluctuation over time and is measured in man-days, while the one in (2.11) tries to minimize schedule shift and is measured in days. To reconcile these two objectives, we first write

$$\frac{y_t + z_t + Cu_t}{\mathrm{LVL}_t}$$

as the workload deviation in period t as a ratio to the ideal workload. The weighted average of this quantity over the entire planning horizon is thus

$$\sum_t \left(\frac{y_t + z_t + Cu_t}{\mathrm{LVL}_t} \right) \left(\frac{\mathrm{LVL}_t}{\sum_t \mathrm{LVL}_t} \right) = \frac{\sum_{t \in T} (y_t + z_t + Cu_t)}{\sum_t \mathrm{LVL}_t} \tag{2.12}$$

where the weight in each period t is its ideal workload LVL_t. This quantity measures *workload fluctuation*.

To convert the objective function in (2.11) in a similar manner, write L_a for the length of availability a, so the shift of availability can be measured as its ratio to its length, namely $\sum_{s \in S_a} x_{a,s}\ \mathrm{SFT}_{a,s}/L_a$. The weighted average of this ratio over all availabilities—with the weight being each availability's length—is

$$\sum_a \left(\frac{\sum_{s \in S_a} x_{a,s}\, \mathrm{SFT}_{a,s}}{L_a} \right) \left(\frac{L_a}{\sum_a L_a} \right) = \frac{\sum_a \sum_{s \in S_a} x_{a,s}\ \mathrm{SFT}_{a,s}}{\sum_a L_a} \tag{2.13}$$

This quantity measures *schedule shift*.

Because the quantity in (2.12) measures the workload fluctuation in ratio, and the quantity in (2.13) measures the schedule shift also in ratio, we can consolidate the two terms into a single objective function

$$\alpha\ \frac{\sum_{t \in T} (y_t + z_t + Cu_t)}{\sum_t \mathrm{LVL}_t} + (1 - \alpha) \frac{\sum_a \sum_{s \in S_a} x_{a,s}\ \mathrm{SFT}_{a,s}}{\sum_a L_a}, \tag{2.14}$$

where α is a number between 0 and 1 to trade off the two competing objectives. Setting $\alpha = 1$ minimizes workload fluctuation, which reduces to the model discussed in Section 2.1. Setting $\alpha = 0$, on the other hand, would minimize schedule shift, which would simply produce the current incumbent start date for each availability.

The revised objective function in (2.14) accounts for the two competing objectives: (1) level load workload fluctuation over time in a regional port, and (2) minimize the schedule shift from fleet operational requirements.

CHAPTER 3: Case Study in the Port of San Diego

This section presents a case study in the Port of San Diego to demonstrate the port loading tool developed in this thesis. The Port of San Diego is a bottleneck for ship maintenance due to a large number of ships stationed there and limited resources at the port. As stated in an article from U.S. Naval Institute News, “the San Diego waterfront has the potential to create a lot of headaches for ships in repair and the fleet operators planning to deploy them” (Eckstein 2017).

The case study demonstrates how the port loading tool optimizes the maintenance schedule for the period of 10/01/2019–09/30/2024 based on data from Naval Sea Systems Command (NAVSEA). Recall that the optimization model has two competing objectives: (1) level the workload over time in a regional port, and (2) minimize the schedule shift from fleet operational requirements. The port loading tool generates an optimal schedule by shifting availability start dates—within the allowable window—to minimize peaks and valleys of the workload. In doing so, private shipyards achieve a more sustainable and predictable workload, which in turn reduces the risk of maintenance backlogs for the Navy.

The rest of this chapter is organized as follows. Section 3.1 discusses the data. Section 3.2 explains how to convert the data into various parameters in the model. Section 3.3 presents the results of this case study.

3.1 Data

The data used to conduct this case study are obtained from NAVSEA in two Excel files. The first Excel sheet contains information on the number of production days per month and estimated labor curves based on a given availability duration. The second Excel sheet is the current maintenance schedule from Southwest Regional Maintenance Center (SWRMC) and Northwest Regional Maintenance Center (NWRMC) that reflects current time lines, types, and assignments for maintenance conducted across the West Coast. We next describe these data.

3.1.1 Production Days

Production days are the number of days in a month that private shipyards are open for business to work on ship maintenance. The number of production days in a month varies from month to month depending on the number of weekdays, weekends, and holidays in that month. Table 3.1 shows the number of production days for each month in 2020.

Table 3.1. The number of production days in 2020.

Jan	Feb	Mar	Apr	May	Jun
21	19	22	22	20	22
Jul	Aug	Sep	Oct	Nov	Dec
22	21	21	21	19	22

3.1.2 Labor Curve of an Availability

The labor curve defines how the total maintenance work is distributed throughout the duration of an availability. The data given by NAVSEA specify the fraction of labor required to be executed each month based on the length of an availability, up to 48 months. Table 3.2 presents a sample of the monthly distribution of labor for availability durations up to six months.

Table 3.2. The monthly distribution of workload for an availability based on its duration.

Availability	Monthly Distribution					
Duration	1	2	3	4	5	6
1 month	1	–	–	–	–	–
2 months	0.4	0.6	–	–	–	–
3 months	0.3	0.5	0.2	–	–	–
4 months	0.25	0.35	0.25	0.15	–	–
5 months	0.2	0.35	0.25	0.15	0.05	–
6 months	0.15	0.25	0.25	0.2	0.1	0.05

3.1.3 Continuous Maintenance

Continuous maintenance (CM) jobs are annual work completed at private shipyards that require personnel from the same labor pool as maintenance availabilities. These jobs continue throughout the year, and a certain fraction is allocated to each month, as seen in Table 3.3. A fraction of 0.089 is allocated to months with 31 days, and 0.077 is allocated to months with 30 days, while 0.067 is allocated to February.

Table 3.3. The workload distribution for continuous maintenance over a year.

Jan	Feb	Mar	Apr	May	Jun
0.089	0.067	0.089	0.077	0.089	0.077
Jul	Aug	Sep	Oct	Nov	Dec
0.089	0.089	0.079	0.089	0.077	0.089

3.1.4 Maintenance Schedule

The NAVSEA data provide the fleet's current maintenance schedule for surface ship maintenance availabilities scheduled in the Port of San Diego from FY16 to FY27. Table 3.4 displays a sample of data provided. Hull indicates the ship class and hull number for each availability. Start Date and End Date represent the current fleet schedule for availabilities by fleet planners. Duration represents the total number of calendar days to complete the maintenance. Total Labor represents the number of man-days required by private shipyards to complete the maintenance.

If the availabilities are executed according to their start dates in the current maintenance schedule, then the aggregate workload will fluctuate over time. The objective of our port loading model is to move the start dates of some availabilities—within the allowable windows—to smooth out the aggregate workload so as to achieve level loading.

3.2 Building Model Parameters

To run the port loading model presented in Section 2.1, we need to convert the data described in Section 3.1 to the model parameters. These parameters include each availability's early start date and late start date, the ideal workload in each period LVL_t, the regional labor

Table 3.4. Maintenance schedule from SWRMC and NWRMC.

Hull	Start Date	End Date	Duration	Total Labor
LPD-23	07/08/2019	10/16/2020	467	102941
CG-57	09/20/2019	06/03/2020	248	80899
DDG-106	11/18/2019	06/29/2020	225	32919
CG-53	02/03/2020	11/20/2020	292	49253
CVN-72	09/30/2020	03/31/2021	183	75133
CG-70	11/16/2020	02/18/2022	460	133383
LCS-10	03/15/2021	09/10/2021	180	81118
DDG-113	05/17/2021	10/15/2021	152	25982
CVN-70	06/15/2022	12/15/2022	184	71483

capacity in each period CAP_t, the target workload in each period TGT_t, and the workload distribution of an availability based on its actual start date $\text{LBR}_{a,s,t}$.

The early start date and late start date of an availability can be calculated based on its start date in the fleet maintenance schedule and the planning date—usually today's date. In our case study, the planning date is 10/01/2019, which is the start of FY20, and the planning horizon is 5 years. If the availability is currently scheduled to start on a day prior to the planning date, such as LPD-23 and CG-57 in Table 3.4, then the availability is in progress as of the planning date so the early and late start dates are both identical to the start date in the fleet schedule. In other words, the start date of an availability that is already in progress as of the planning date is not allowed to change in the optimization model.

If the availability start date is within 2 years of the planning date, then it can shift in either direction for up to 30 days, with the caveat that it cannot shift early to within 250 days of the start date, because 250 days is the lead time the Navy needs to set up the maintenance contract. Consequently, if an availability's current start date is between 250 and 280 days from the planning date, then its early start date is 250 days from the planning date. If an availability's start date is between 280 days and 2 years from the planning start date, then both the early and late start date can shift up to 30 days. Finally, if an availability's start date is more than 2 years away from the planning date, such as CVN-70 in Table 3.4, then the start date can shift up to 45 days in either direction.

Table 3.5 displays the same sample of data provided by NAVSEA as in Table 3.4, but includes the early and late start dates for each availability.

Table 3.5. Maintenance schedule from SWRMC and NWRMC with early and late start dates.

Hull	Start Date	End Date	Duration	Total Labor	Early Start Date	Late Start Date
LPD-23	07/08/2019	10/16/2020	467	102941	07/08/2019	07/08/2019
CG-57	09/20/2019	06/03/2020	248	80899	09/30/2019	09/30/2019
DDG-106	11/18/2019	06/29/2020	225	32919	11/18/2019	11/18/2019
CG-53	02/03/2020	11/20/2020	292	49253	02/03/2020	03/04/2020
CVN-72	09/30/2020	03/31/2021	183	75133	08/31/2020	10/30/2020
CG-70	11/16/2020	02/18/2022	460	133383	10/17/2020	12/16/2020
LCS-10	03/15/2021	09/10/2021	180	81118	02/13/2021	04/14/2021
DDG-113	05/17/2021	10/15/2021	152	25982	04/17/2021	06/16/2021
CVN-70	06/15/2022	12/15/2022	184	71483	05/01/2022	07/30/2022

The parameter LVL_t represents the ideal workload in time period t, which typically reflects the labor pool in the region. In order to calculate LVL_t, we first calculate the *total labor* required to complete all maintenance work during the planning horizon according to the fleet maintenance schedule, including both ship availabilities and continuous maintenance. Next, divide the *total labor* by the total number of production days in the planning horizon, and then multiply this value by the number of production days in each time period t to obtain LVL_t. In other words, LVL_t is the workload in period t assuming that the total labor is distributed evenly over each production day in the planning horizon.

The parameter CAP_t represents the regional labor capacity in time period t. In practice, this value can be estimated by private shipyards or by experts in Navy maintenance community. Or it can be a subjective value, which the Navy does not want the monthly workload to exceed. In our case study, we assume that the regional labor capacity is 10% more than the leveled workload, and set $\mathrm{CAP}_t = 1.1\mathrm{LVL}_t$.

The parameter TGT_t refers to the target workload in period t that can be allocated to availabilities in the optimization model. For each period t, the difference between TGT_t and LVL_t is the amount of labor in period t that has been committed to continuous maintenance and ongoing availabilities. Because continuous maintenance and ongoing availabilities cannot be moved, they are treated as constant in the model. The parameter TGT_t therefore refers to the labor that is *available* to assign to upcoming availabilities whose start dates we wish to optimize. To calculate TGT_t, we start with LVL_t and subtract from it the labor already committed to ongoing availabilities and continuous maintenance, according to Table 3.2 and Table 3.3, respectively.

The parameter $\mathrm{LBR}_{a,s,t}$ represents the value of labor required during period t for availability a, if it starts on date s. The backbone that supports this calculation is the labor curve discussed in Section 3.1.2. The labor curve based on months, as seen in Table 3.2, however, cannot be used directly to calculate $\mathrm{LBR}_{a,s,t}$ for two reasons. First, each month may have a different number of production days. Second, the length of an availability may not be multiple of whole months. To make the conversion, we first use 21 as the nominal number of production days in a typical month to construct a labor curve based on production days, by assuming each production day in the same month requires the same amount of labor. To calculate $\mathrm{LBR}_{a,s,t}$, we use the actual number of production days in each month to calculate the workload required in period t, and use linear interpolation when the length of an availability, measured in production days, is not multiple of 21. This process is replicated for all feasible start dates s that availability a can start on, and the result is a parameter with a proper representation of how much labor availability a requires during time period t if starting on date s.

Running the model also requires setting the value for $C \gg 1$, which heavily penalizes any workload that exceeds the labor capacity, as seen in (2.14). Setting $C \gg 1$ essentially asks the solver to first minimize any workload that exceeds the labor capacity, before trying to minimize workload fluctuation. In other words, the solver is strongly motivated to return a solution with the workload below the labor capacity in every period, if there is one. In numerical testing, setting $C = 100$ and setting $C = 1000$ essentially produce the same result. We set $C = 1000$ in this case study.

3.3 Results

3.3.1 Maintenance Schedule Before Optimization

The case study concerns maintenance schedule in the Port of San Diego between October 1, 2019 and September 30, 2024, for a planning horizon of 5 years. Before any optimization, the original maintenance schedule has a substantial amount of workload fluctuation, as seen in Figure 3.1.

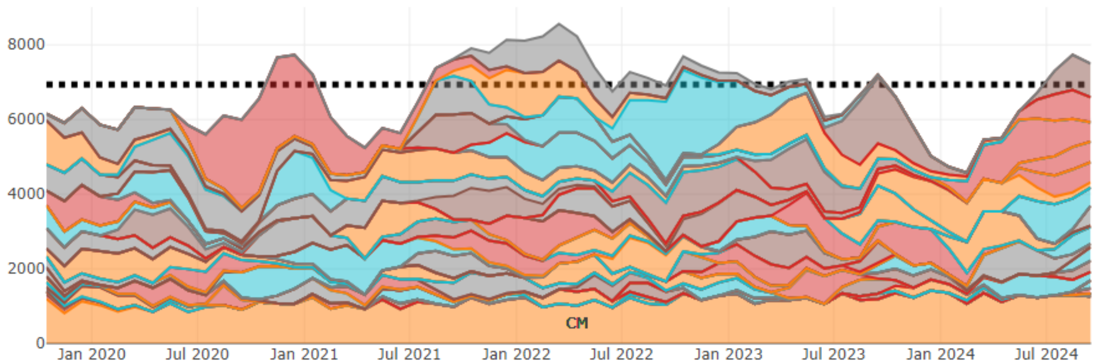

Figure 3.1. Workload with original schedule.

The y-axis in Figure 3.1 corresponds to the number of man-days per day required to complete surface ship maintenance for each month. The dashed line represents LVL_t = 6936 (man-days per day) over the planning horizon, which would be the ideal workload distribution for private shipyards. The bottom curve annotated by CM on the graph shows the portion of workload dedicated to continuous maintenance that occurs throughout the planning horizon, and each colored layer on the graph corresponds to one availability.

As seen in Figure 3.1, the fleet's maintenance schedule fluctuates substantially through the 5 years, with valleys in the first half of 2020 and the beginning of 2024, and peaks in late 2022 to early 2023.

The fleet's maintenance schedule results in 12.93% fluctuation from the level workload over the planning horizon. The percent deviation includes both times when workloads are above the level workload and times when workloads are below the level workload shown as peaks and valleys in Figure 3.1. The drastic fluctuation is what hurts the private shipyard industry in maintaining a skilled workforce to conduct ship maintenance promptly.

3.3.2 Maintenance Schedule After Optimization

We next run the port loading model by setting $\alpha = 1$ in (2.14). By doing this, each availability's start date can move freely in its respective allowable window without an explicit penalty, so setting $\alpha = 1$ maximizes level loading to the best possible extent. As seen in Figure 3.2, the workload becomes more consistent over time, compared with that in Figure 3.1. The optimal schedule reduced the workload fluctuation from 12.93% in the original schedule to 3.16%. The average schedule shift for each availability is 10.26% from the original schedule, but all shifts are within the allowable window required by the fleet. The port is underloaded in the early part of 2020 because the availabilities are not allowed to shift early into the first 250 days of the planning horizon, which is needed for the Navy to set up a contract with private shipyards.

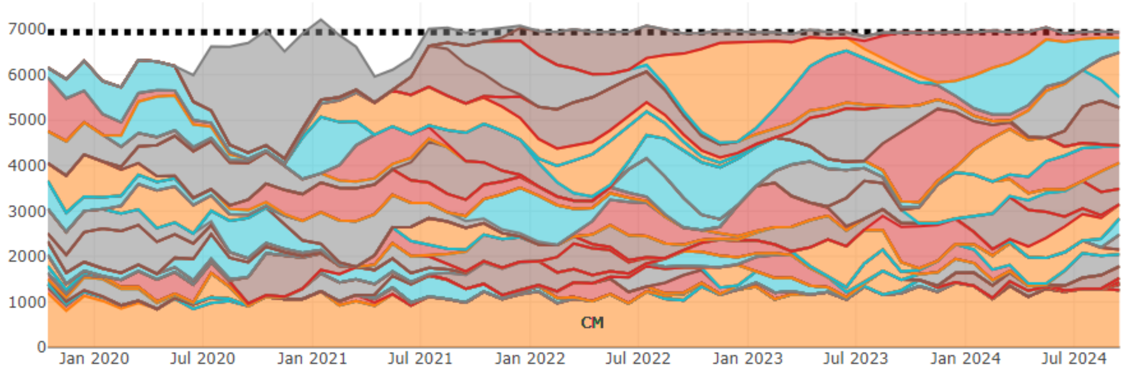

Figure 3.2. Workload with optimal schedule.

Overall, the optimal schedule shows great improvement compared with the original schedule in Figure 3.1. This schedule has far fewer deviations from the leveled workload than the original fleet schedule. By maintaining a nearly leveled workload throughout the 5-year period, the optimal schedule makes it much easier for private shipyards to complete maintenance availabilities on time.

While the deviation from the level workload is not eliminated entirely, the port loading tool demonstrates that there is room for improvement in the current maintenance schedule.

The improvement comes at the cost of shifting the current schedule, but produces a more manageable schedule for both the Navy and private shipyards. In addition to plots shown in Figure 3.1 and Figure 3.2, the tool also outputs a comma separated value (csv) file with the optimal availability start date, end date, and days shifted for each availability, which allows fleet planners to readily have access to the optimal schedule.

3.3.3 Verification of the Two-Stage Optimization

Recall from the discussion in Section 2.2, we design a two-stage approach to run the port loading model to reduce the run time of the optimization model. Stage one allows each availability to start on a few selected dates to obtain an estimate of the optimal solution, and stage two finds the exact optimal solution by improving upon the solution from stage one. Recall from Section 2.2 that in stage one we define a step size for the resolution of the model in how it iterates through potential start dates over the time horizon. For example, when step is equal to 5, the feasible start dates of an availability iterate from the early start date to the late start date with a 5 day difference between each date. If step equals 1, then there is no difference in the results from stage one to stage two of the model, but the computational time to solve this problem is significantly larger than choosing to use a larger step size. Running the MILP model twice speeds up computational time. Table 3.6 shows the results for different step sizes from the two-stage optimization model with $\alpha = 0.5$ in equation (2.14). The optimization time is the time taken to build model parameters and run the MILP optimization model for the case study from 10/01/2019–09/30/2024. Stage 1 optimization defines 327 constraints and 1799 variables of which 1587 are binary, and Stage 2 defines 328 constraints and 1313 variables of which 1100 are binary. The tool uses the CBC solver (Forrest 2000) on personal computer with a 2.60 GHz Intel Core i7 CPU with 8 GB RAM.

From Table 3.6, it is apparent that with a step size equal to 6, we achieve the fastest running model. The solutions obtained from different step sizes are all very close, although not identical, because the MILP solver returns a solution when it finds one within 0.001% of the optimal solution.

Table 3.6. Model output statistics for different step sizes.

	Stage 1		Stage 2		
Step	Schedule Shift	Workload Fluctuation	Schedule Shift	Workload Fluctuation	Optimization Time (sec)
1	4.32%	5.90%	4.32%	5.90%	163
2	4.39%	5.73%	4.26%	5.76%	96
3	4.17%	5.89%	4.19%	5.79%	89
4	4.42%	5.75%	4.16%	5.73%	95
5	4.27%	5.72%	4.19%	5.71%	96
6	5.11%	5.14%	4.66%	5.26%	84
7	4.99%	5.40%	4.56%	5.32%	86
8	4.92%	5.45%	4.43%	5.47%	93
9	4.35%	5.76%	4.11%	5.79%	93
10	5.23%	5.11%	4.55%	5.37%	92

3.3.4 Efficient Frontier

Recall that the MILP model accounts for two competing objectives as seen in equation (2.14). The first objective is to minimize the schedule shift from fleet operational requirements, which allows ship schedules to remain highly operational and consistent with the needs of the Navy. The second objective is to minimize any deviation from the level workload across the planning horizon, which allows private shipyards to better sustain a skilled and constant workforce. Together, these competing objectives mitigate the risk of maintenance delays, which in turn increase operational readiness across the fleet.

The results presented in Table 3.6 are based on $\alpha = 0.5$, or a 50-50 weight between the two competing objectives. However, if the user determines that one objective is more desirable than the other objective, then different weights can be implemented into the model by changing the value of α. As seen in Table 3.7, as the weight assigned to workload deviation increases, the optimal solution has a smaller workload fluctuation at the cost of a larger schedule shift.

Figure 3.3 plots the trade off between the percentages of workload fluctuation and schedule shift when different weights are assigned to the two competing objective functions. The left-most, blue point corresponds to the schedule before optimization with $\alpha = 0$, and the

Table 3.7. Trade off between two objectives when α, the weight assigned to workload fluctuation, varies.

α	Schedule Shift	Workload Fluctuation
0.0	0%	12.93%
0.1	0.35%	12.63%
0.2	0.51%	11.88%
0.3	0.99%	10.34%
0.4	2.04%	8.40%
0.5	4.66%	5.26%
0.6	6.37%	4.05%
0.7	7.22%	3.47%
0.8	7.92%	3.32%
0.9	8.49%	3.17%
1.0	10.26%	3.16%

right-most, red point corresponds to the schedule with $\alpha = 1$. When one objective function increases, the other decreases. The plot shows the efficient frontier of the two competing objective functions, so a decision maker can use it to choose the solution that meets the fleet's need the best.

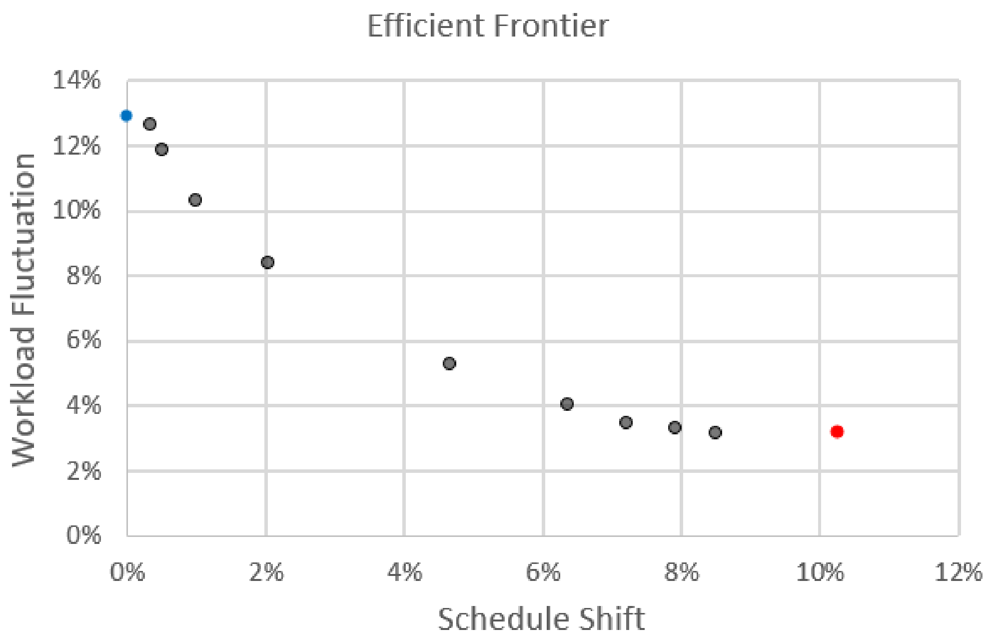

Figure 3.3. Efficient frontier for workload fluctuation and schedule shift.

CHAPTER 4: Conclusion and Future Research

This thesis presents a port loading tool that the Navy can use to schedule maintenance availabilities while level loading the work required of private shipyards in a port over time. The model is based on a mixed integer linear program (MILP) that generates an optimal schedule by shifting availability start dates within an allowable time frame to minimize workload fluctuation over time. The inputs to the model include the current fleet maintenance schedule, a planning date, and planning horizon. With these data, the model seeks a solution that optimally balances two competing objectives: (1) level the workload over time in a regional port, and (2) minimize the schedule shift from fleet operational requirements. The output of the model is a win-win schedule for both the Navy and private shipyards. The new schedule reduces workload fluctuation experienced by private shipyards, which in turn reduces the risk of increased maintenance delays for the fleet.

The port loading tool marks a significant step in using data and optimization to automate fleet planning. While the Navy leader's will always need to provide direction and insight, mathematical models, such as the port loading tool created here, assist in automating time consuming, error prone methods. The port loading tool has several applications:

1. Allow fleet planners to generate multiple long-term schedules that balance workload fluctuation and schedule shift for maintenance availabilities.
2. Identify periods of congestion for private shipyards in terms of workload requirements and justify alterations to the current fleet schedule.
3. Explore a variety of scenarios with different restrictions from the perspective of both the Navy and private shipyards.
4. Create an opportunity to improve current contracting strategies.

The port loading tool performs well within the scope and limitations accounted for by the model. In a case study in the Port of San Diego, the tool reduces workload fluctuation substantially over a 5-year period. The port loading tool comes with a friendly graphical user interface that allows an end user to run a variety of scenarios with different dates and ports.

It is important to have accurate data to produce meaningful results. Further applications of the port loading tool require up-to-date data. For example, if a ship crashes and needs immediate pier-side maintenance, then the tool can be used to determine how to best adjust future maintenance schedule to accommodate the emergency maintenance. In order to overcome quick changes in schedule, the tool can be run on a regular basis to provide real-time assessments of the current maintenance schedule and provide input to adjustments.

The work in this thesis is part of an ongoing effort to use data to improve the scheduling of maintenance availabilities. Hilliard (2019) developed the Surface Ship Drydock Schedule Planner (SSDSP), which optimizes the schedule of availabilities that require a drydock. The output of the model can be used as the input of the port loading tool to *lock* the dates of drydocking availabilities, so that the port loading tool can optimize the dates of the other availabilities.

A few potential research areas can expand on the work done in this thesis. First, the port loading tool can be extended from planning for workload to planning for resources and facilities to further reduce the risk of maintenance delays. Additionally, the port loading tool is currently designed to only level workload at one port, and it does not account for interaction between ports. An area of future research would be to expand the model to simultaneously optimize level loading at two or more ports if availabilities can be contracted out to different ports. Also, the model can be reformatted to provide an initial feasible solution from Stage 1 to Stage 2 of the optimization. This helps the solver fathom more partitions of the feasible region earlier in execution, which can potentially speed up computational time. Finally, the techniques presented in this thesis have the potential to open new opportunities to automate scheduling for other entities in the Navy—such as personnel assignment, fleet management, and aircraft maintenance. An automated scheduling process minimizes the chance of human errors and can improve efficiency within the Navy's logistics.

List of References

Brown RA (1992) *A naval shipyard optimal drydock loading and capacity utilization model*. Master's thesis, Operations Research Department, Naval Postgraduate School, Monterey. http://hdl.handle.net/10945/38569.

Caprio LD (2012) *A preliminary investigation into CNO availability schedule overruns*. Master's thesis, Operations Research Department, Naval Postgraduate School, Monterey.

Eckstein M (2017) Navy facing drydock capacity issue in surface ship repair; testing out new maintenance contract to address shortfall, create efficiency. USNI News. Accessed February 5, 2019, https://news.usni.org/2017/10/12/navy-facing-drydock-capacity-issue-surface-ship-repair-testing-new-maintenance-contract-address-shortfall-create-efficiency.

Forrest J (2000) CBC, version 2.9.9. Accessed March 13, 2019, https://projects.coin-or.org/Cbc.

Geurts JF, Moore T (2019) Statement before the senate armed services committe on ship and submarine maintenance.

Hart W, Watson JP, Woodruff D (2008) Pyomo, version 5.6.1. Accessed March 13, 2019, https://www.pyomo.org.

Hilliard A (2019) *A drydock loading model for surface ship maintenance availabilities*. Master's thesis, Operations Research Department, Naval Postgraduate School, Monterey.

Martin B, McMahon ME, Riposo J, Kallimani JG, Bohman A, Ramos A, Schendt A (2017) A strategic assessment of the future of us navy ship maintenance. *Challenges and Opportunities. RAND Corporation* 53.

Maurer D (2019) Navy maintenance: Persistent and substantial ship and submarine maintenance delays hinder efforts to rebuild readiness, statement of diana c. maurer, director, defense capabilities and management, testimony before the subcommittees on seapower and readiness and management support, committee on armed services, us senate. Technical Report GAO-20-257T, GAO.

Naval Sea Systems Command (2020) Report to Congress on the Long-Range Plan for Maintenance and Modernization of Naval Vessels for Fiscal Year 2020.

Northrup D (2015) *A business process analysis of the surface Navy's depot maintenance program*. Master's thesis, Operations Research Department, Naval Postgraduate School, Monterey.

Parmer C (2016) Dash, version 0.39.0. Accessed March 13, 2019, https://plot.ly/products/dash/.

Pendleton JHM (2016) Progress and challenges in implementing the navy's optimized fleet response plan. Technical report, GAO-16-466R (Washington, DC: US GAO, 2 May 2016).

Schaefer M (2017) *A Port Loading Model For Surface Ship Maintenance Availabilities*. Master's thesis, Operations Research Department, Naval Postgraduate School, Monterey.

van Rossum G (1991) Python, version 3.7.1. Accessed March 13, 2019, https://www.python.org.

Initial Distribution List

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California