NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

SCHEDULING SHIP MAINTENANCE JOBS IN MULTIPLE PORTS TO MINIMIZE WORKLOAD FLUCTUATION

by

Sarah B. Naldo

June 2021

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

Approved for public release. Distribution is unlimited.

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** June 2021	**3. REPORT TYPE AND DATES COVERED** Master's thesis		
4. TITLE AND SUBTITLE SCHEDULING SHIP MAINTENANCE JOBS IN MULTIPLE PORTS TO MINIMIZE WORKLOAD FLUCTUATION				**5. FUNDING NUMBERS** RNP5P
6. AUTHOR(S) Sarah B. Naldo				
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) COMMANDER NAVAL SURFACE FORCES PACIFIC 2841 RENDOVA ROAD SAN DIEGO, CA 92155-5490				**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)

Surface ships in the Navy require planned and unplanned pier-side maintenance. These maintenance jobs, known as availabilities, are contracted out to private shipyards. Ship maintenance schedules must meet the Navy’s operational requirements and stay within the capacity of the contracted shipyards. At the same time, it is important to minimize workload fluctuation in a port to help private shipyards train and maintain a skilled workforce. Building on recent work that schedules availabilities in a single port to minimize workload fluctuation, this thesis develops a port loading model to minimize workload fluctuation for all regional ports in the Area of Responsibility by allowing some ships to receive maintenance work out of their home ports. Scheduling availabilities across multiple ports simultaneously to level the workload in each port has two additional benefits: First, an increase in the number of eligible companies who can bid on the maintenance job will drive down the cost for the Navy. Second, allowing more flexibility to assign availabilities to different ports has the potential to further level the workload at these ports. In a case study on three ports in the West Coast over a six-year period, we demonstrate the effectiveness of the multi-port loading model.

14. SUBJECT TERMS optimization, port loading, scheduling, linear programming, ship maintenance, surface ships, fleet readiness			**15. NUMBER OF PAGES** 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.

SCHEDULING SHIP MAINTENANCE JOBS IN MULTIPLE PORTS TO MINIMIZE WORKLOAD FLUCTUATION

Sarah B. Naldo
Ensign, United States Navy
BS, United States Naval Academy, 2020

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 2021

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

Surface ships in the Navy require planned and unplanned pier-side maintenance. These maintenance jobs, known as availabilities, are contracted out to private shipyards. Ship maintenance schedules must meet the Navy's operational requirements and stay within the capacity of the contracted shipyards. At the same time, it is important to minimize workload fluctuation in a port to help private shipyards train and maintain a skilled workforce. Building on recent work that schedules availabilities in a single port to minimize workload fluctuation, this thesis develops a port loading model to minimize workload fluctuation for all regional ports in the Area of Responsibility by allowing some ships to receive maintenance work out of their home ports. Scheduling availabilities across multiple ports simultaneously to level the workload in each port has two additional benefits: First, an increase in the number of eligible companies who can bid on the maintenance job will drive down the cost for the Navy. Second, allowing more flexibility to assign availabilities to different ports has the potential to further level the workload at these ports. In a case study on three ports in the West Coast over a six-year period, we demonstrate the effectiveness of the multi-port loading model.

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

AOR	area of responsibility
CM	continuous maintenance
CWB	coast-wide bid
FY	fiscal year
GAO	Government Accountability Office
MAC-MO	Multiple Award Contract-Multi Order
MILP	mixed integer linear program
MSMO	Multi-Ship, Multi-Option
NAVSEA	Naval Sea Systems Command
NPS	Naval Postgraduate School
NWRMC	Northwest Regional Maintenance Center
POR	Portland port
RPD	resources per day
SEA	Seattle port
SSD	San Diego port
SSDSP	Surface Ship Drydock Schedule Planner
SWRMC	Southwest Regional Maintenance Center

THIS PAGE INTENTIONALLY LEFT BLANK

Executive Summary

Surface ships in the Navy require planned and unplanned pier-side maintenance. These maintenance jobs, known as maintenance availabilities, are contracted out to private shipyards. Ship maintenance schedules must meet the Navy's operational requirements and stay within the capacity of the contracted shipyards. Maintenance delays impede the Navy's operational requirements and pile up work on the contractors. A recent focus of the Navy is to reduce maintenance delays by leveling the maintenance workload. Minimizing workload fluctuation in a port helps private shipyards train and maintain a skilled workforce.

Recent work has developed a model that schedules maintenance availabilities in a single port to minimize workload fluctuation. In this thesis, we extend that work to minimize workload fluctuation *simultaneously* for all regional ports in an entire Area of Responsibility (AOR) by allowing eligible ships to receive maintenance work from outside their home ports. An availability that is eligible for coast-wide bid (CWB) can be contracted to any shipyard in any port in the AOR. The model developed in this thesis attempts to minimize the workload fluctuation by adjusting each availability's start date for the entire AOR and sending a ship out of its home port when necessary, while keeping the ships at their home port as much as possible. The input of the model is a proposed maintenance schedule for an AOR, including each availability's start date, duration, home port, and eligibility for CWB. The output of the model is a maintenance schedule optimized for minimal workload fluctuation and out-of-homeport maintenance. The multi-port loading model is formulated as a mixed integer linear program, and is coded in Python and Pyomo.

We explore a case study for the West Coast AOR—which consists of the ports San Diego, Seattle, and Portland—for Fiscal Years 2021 through 2026. Naval Sea Systems Command (NAVSEA) provides the proposed maintenance schedule. We compare the workload for four scenarios:

1. The original schedule proposed by NAVSEA.
2. The optimal schedule by assigning each availability to the ship's home port, without considering CWB eligibility, and adjusting its start date.
3. The optimal schedule by assigning each availability to one of its eligible ports,

considering CWB eligibility, and adjusting its start date.

4. The optimal schedule by assigning each availability to one of its eligible ports, considering CWB eligibility, and adjusting its start date, with additional constraints on docking availabilities to ensure an executable docking schedule.

In the case study, the workload fluctuation decreases from scenario 1 to scenario 2, and then further in scenario 3, because CWB-eligible availabilities can move out of the ship's home port. The workload fluctuation in scenario 4 sits between those in scenarios 2 and 3, because in scenario 4, docking availabilities are constrained by an executable docking schedule and cannot be shifted to level the workload. For example, the workload in San Diego for the four scenarios are 19.4%, 11.1%, 5.65%, and 8.95%, respectively. The first two scenarios do not allow CWB, so all availabilities stay in their home ports. When we allow coast-wide bid in scenario 3, ships spend 11.2% of their time in maintenance away from their home ports. With the additional docking constraints in scenario 4, the proportion of time ships spend in maintenance away from their home ports increases to 12.7%. Among the 128 availabilities assigned in the case study, 17 of them are assigned to a port other than the ships' home ports in scenario 3, and 16 in scenario 4.

Scheduling maintenance availabilities simultaneously for all ports in the same AOR can reduce the workload fluctuation in each port by taking advantage of the CWB-eligible availabilities. There are several potential applications how the Navy can use the model to help plan maintenance schedule, such as adjusting the maintenance schedule after an incident that requires unexpected maintenance, assessing how new ships affect the maintenance schedule of the others, and projecting labor demand as the Navy grows its fleet. This model is an important contribution in leveling the workload of maintenance availabilities to benefit both the Navy and the private shipyards.

Acknowledgments

This thesis would not have been possible without the support of many people in my professional and personal life. First, I would like to thank Dr. Kyle Lin for advising me on this project. Your expertise and patience were essential to my understanding of the problem, how we can address it, and how we can communicate it. I'd also like to thank Dr. Emily Craparo for providing a sharp and analytic eye, as well as patience and positivity. Finally, thanks to my friends, family, and mentors for anchoring me to Truth and Reality when the drudgery of distance learning threatened to consume me.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:
Introduction

Navy surface ships undergo frequent maintenance, both planned and unplanned. These maintenance jobs—known as maintenance availabilities—must meet the operational requirements of the Navy as well as stay within the capacity of the civilian contractors who perform the maintenance availabilities. A recent focus of the Navy is to promote efficiency in ship maintenance by minimizing delays. Schedule delays impede the Navy's operational requirements and pile up work on contractors and thus cause more delays.

One major cause for maintenance schedule delay is workload fluctuation in a regional port. Private shipyards within a regional port have a certain labor capacity. Large fluctuations in the labor demand over time make it difficult for shipyards to maintain a trained workforce. The Navy has several initiatives to cooperate with private shipyards and optimize the use of their capabilities (Naval Sea Systems Command 2020).

Previous work has produced optimization models that help schedule ship maintenance availabilities within the capacity of contractors and drydocks. Recently, O'Malley and Lin (2020) designed a model to minimize workload fluctuation at a single port. While the work in O'Malley and Lin (2020) adjusts the starting time of each availability in a port to level the workload over time in that port, it does not consider awarding a contract outside a ship's home port. When the scope of a maintenance availability exceeds a certain threshold, the availability is designated by the Navy as coast-wide bid (CWB) eligible, which helps the Navy reduce the contract cost by involving more private shipyards in the other ports in the same area of responsibility (AOR) to bid on the contract. The focus of this thesis is to extend the port loading model in O'Malley and Lin (2020) to account for CWB eligibility. In other words, each CWB eligible availability will be optimally assigned to a port, and every availability will be optimally assigned a starting time, to minimize the workload fluctuation in each port in the AOR simultaneously.

1.1 Motivation

As of Fiscal Year (FY) 2020, the Navy has 301 vessels that require maintenance (Naval Sea Systems Command 2020). The Navy has four public shipyards that primarily support availabilities for nuclear-powered vessels, which leaves about 200 surface ships that rely primarily on private contracted shipyards for maintenance. Martin et al. (2017) reports that there are 22 private shipyards to perform maintenance on these ships across the US. Figure 1.1 displays Navy home ports and ports with shipyards equipped for maintenance of Naval surface ships.

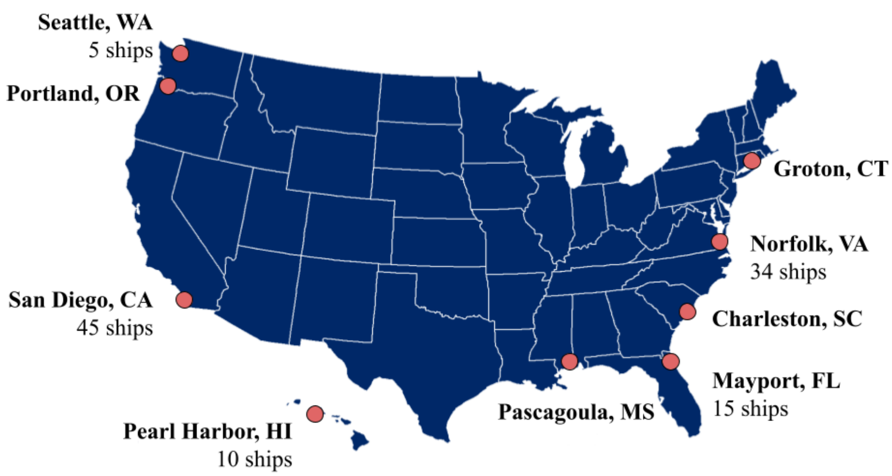

Figure 1.1. Naval surface ship home ports and maintenance ports in the continental United States. Adapted from Naval Sea Systems Command (2020); Martin et al. (2017).

Ideally, the Navy prefers to award a ship's maintenance contract to a shipyard in each ship's home port. However, as seen in Figure 1.1, there is a huge imbalance of ship-to-shipyard ratios among the ports. If a ship in San Diego requires maintenance, but the shipyards in San Diego are nearing capacity, then it may make sense to allow a shipyard in Seattle or Portland to perform the ship maintenance. To allow more shipyards to compete on the same maintenance contract, the Navy designates some availabilities as CWB eligible. Naval Sea Systems Command (NAVSEA), which is responsible for coordinating the maintenance of Navy ships with shipyards, considers availabilities that last longer than six months to be eligible for coast-wide bid (Mackin 2017). The coast-wide bid option helps alleviate some of the capacity shortfalls for busy ports like San Diego (Martin et al. 2017).

The Navy can utilize coast-wide bids when coordinating maintenance schedules with private shipyards. While shipyards certainly require maintenance jobs to stay within their labor capacity, they prefer to have a leveled workload, which promotes a steady, well-trained workforce. The Navy benefits from a well-trained workforce with their private contractors because a well-trained workforce has fewer maintenance delays (Martin et al. 2017).

Not every availability is eligible for coast-wide bid. First, availabilities must last longer than 6 months (Office of the Chief of Naval Operations 2019). This requirement is due to the financial costs of moving a ship out of its home port as well as hardships on the sailors who would have to temporarily move away from their duty station. There must also be "adequate competition," which the Navy defines as two or more qualified bidders (Oakley 2020). If there is adequate competition within the home port, the bidding competition will likely drive the cost of maintenance down enough to outweigh the cost of moving the ship out of its home port. There must also be at least 120 days of lead time to award the contract to the highest bidding shipyard (Oakley 2020).

By 2030, the Navy plans to increase the size of its fleet to 355 ships (Naval Sea Systems Command 2020). As the Navy grows and the ship-to-shipyard ratio increases, efficient maintenance scheduling becomes more important.

1.2 Background

Before 2015, the Navy used a contract-awarding strategy known as Multi-Ship, Multi-Option (MSMO). Under the MSMO strategy, the Navy awarded several availabilities to a shipyard in one contract, limited competition by grouping ship classes, and left much of the planning responsibility to the shipyard. Coast-wide bid was not a priority, and few coast-wide competitions were held for availabilities over six months (Mackin 2017).

In 2015, the Navy transitioned from MSMO to Multiple Award Contract-Multi Order (MAC-MO). With MAC-MO, NAVSEA coordinates planning with a third-party service to better establish cost reimbursement before an availability begins. Additionally, the contracts awarded to private shipyards are more flexible. While availabilities are primarily contracted by ship class within their home port, NAVSEA may award individual contracts for emergent maintenance and *encourages* coast-wide competition for availabilities longer than 6 months (Mackin 2017). In a report by the Government Accountability Office (GAO), Mackin (2017) reports, “The increase in competition opportunities that MAC-MO offers has the potential to help save the taxpayer money, improve contractor performance, and promote accountability for results.”

In May of 2020, the GAO performed another analysis of the MAC-MO contracting approach, and found that 21 of 41 availabilities cost less than initially estimated, but schedule delays still persist (Oakley 2020). Some factors that contribute to maintenance delays include skilled personnel shortages, insufficient shipyard capacity, and adherence to the planning process (Maurer 2019). One way the Navy seeks to address these issues is to optimize port loading, namely, scheduling maintenance work to best utilize the capacity of shipyards.

1.3 Our Contribution

In this thesis, we develop a mixed integer linear program (MILP) that takes advantage of CWB-eligible availabilities to optimize port loading. In the MILP, each availability is assigned a port and a starting date to simultaneously level the workload of all ports in the same AOR.

The MILP accounts for two competing objectives: Primarily, it seeks to minimize deviation from a perfectly leveled workload, which helps private shipyards maintain a steady workforce. In addition, it also seeks to minimize the time a ship spends outside its home port. Although we hope the coast-wide bid option will help significantly level the workload and make the maintenance process more efficient in that way, if moving a ship out of its home port for a certain availability only slightly helps to level the workload, it may not be worth upending the crew and paying the cost of moving a ship out of its home port. Because of these costs, we want to discourage ships from spending unnecessary time away from their home ports.

To run the model, we need to input a baseline schedule—that is, the proposed start dates, duration, required labor, home port, and coast-wide bid eligibility status according to the Navy’s operational needs as conveyed by NAVSEA. With these inputs, we can determine a window of time and a set of ports to which an availability could be assigned. The model shifts availabilities within their windows of time and across their set of ports and returns a schedule with minimal workload fluctuation and discouraged relocation away from the ships' home ports.

We implement the MILP using the Python programming language (van Rossum 1991) and its Pyomo package (Hart et al. 2008). We incorporate the model into an interactive user-friendly interface with a Python application development package called Dash (Parmer 2016).

1.4 Related Works

There exists some earlier work that uses mathematical modeling to improve maintenance scheduling for Navy surface ships.

Brown (1992) introduces the application of MILPs to scheduling Navy ship maintenance. His research was initially a part of a study on the utilization of drydocks, but the product was a model that could maximize overall drydock capacity utilization over a specified time frame. The inputs for Brown's model are the availabilities' start dates and durations, docks eligible to take the availabilities, and preferences for certain docks. The outputs are dock assignments and schedules. Although the circumstances were different—the Navy was reducing the size of its fleet—Brown (1992) introduced a new way to explore maintenance scheduling strategies and utilization.

Schaefer (2017) designs a model that estimates the realistic completion times of availabilities. The model inputs are the shipyard labor capacity, the availability start date, and total labor required. The outputs are projected labor execution and an estimation of each availability's completion time. Although there is no optimization involved in this work, the estimated duration and distribution of labor for availabilities are useful to improve upon the scheduling models that follow.

Hilliard (2019) develops a MILP model to determine the optimal use of commercial dry docks. The model becomes the backbone of the Surface Ship Drydock Schedule Planner (SSDSP) used by the Navy to help plan drydock schedules (Hilliard et al. 2020). The input of the SSDSP is a surface ship maintenance schedule for availabilities that require drydocks along with their start dates and durations, and a set of drydocks, their locations, and their capabilities. The output is an executable schedule that maximize drydock utilization without any docking conflicts. The SSDSP uses MILPs to explore new scheduling strategies, such as double docking or enhancing a given dry dock's capabilities, to mitigate maintenance delays.

The work most closely related to this thesis is O'Malley and Lin (2020), which develops a port loading model to minimize workload fluctuation over time in a single port. The model in O'Malley and Lin (2020) uses a MILP that takes a proposed maintenance schedule of availabilities in the port, their start dates, durations, and labor required. Each availability is allowed to shift within a pre-specified time window. The model returns a new schedule that minimizes workload fluctuation while also minimizing schedule shift.

The focus of this thesis is to extend the model in O'Malley and Lin (2020) to account for CWB-eligible availabilities discussed in Section 1.2. In other words, each CWB-eligible availability will be assigned a port and every availability will be assigned a starting date to simultaneously level the workload of all ports in the same AOR.

1.5 Thesis Outline

The rest of this thesis proceeds as follows. Chapter 2 presents the formulation of the optimization model to level workload over time and across ports. Chapter 3 discusses a case study in the West Coast Area of Responsibility. Chapter 4 concludes.

CHAPTER 2: The Optimization Model

Section 2.1 introduces the optimization model, which assigns optimal start dates and home ports to availabilities in a given AOR. Section 2.2 explains the methods used to run the model. Section 2.3 addresses a competing objective between level loading and a preference to perform maintenance at the ships' home ports.

2.1 Leveling Workload over Time and Ports

This section introduces a mixed integer linear programming model to achieve level loading at multiple ports within an AOR. Each availability in the planning horizon is allowed to start during a predetermined time window. The goal is to select an optimal port and starting date for each availability so that the projected workload required of a single port during one time period, such as a month, is as close as possible to other that port's workload in other time periods.

Indices and Sets

$a \in A$	Maintenance availabilities [unitless].
$s \in S_a$	Potential starting dates for each availability a [date].
$t \in T$	Monthly time periods over the planning horizon. The length of a period is determined by the calendar month and year [unitless].
$p \in P$	All ports in the AOR [unitless].
$p \in P_a$	Ports in the AOR that availability a is eligible to be maintained at [unitless].

Data

$\text{TGT}_{p,t}$	The target workload for availabilities at port p in period t [man hours / day].
$\text{CAP}_{p,t}$	The labor capacity for port p at time t. It is undesirable for the workload to exceed this threshold [man hours / day].

$\mathrm{LBR}_{a,s,t}$ 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 [man hours / day].

Decision Variables

$x_{a,s,p}$ Binary, 1 if availability a is scheduled to start on date $s \in S_a$ at port $p \in P_a$.

$y_{p,t}$ Nonnegative value indicating the workload projected at port p in period t above the target workload $\mathrm{TGT}_{p,t}$ [man hours / day].

$z_{p,t}$ Nonnegative value indicating the workload projected at port p in period t below the target workload $\mathrm{TGT}_{p,t}$ [man hours / day].

$u_{p,t}$ Nonnegative value indicating the workload projected at port p in period t above the labor capacity $\mathrm{CAP}_{p,t}$ [man hours / day].

Formulation

Throughout this thesis, any time there is a summation over an index, it is over the entire set unless otherwise noted. The initial formulation of our model to minimize workload fluctuation is as follows:

$$
\begin{aligned}
\min \quad & \sum_{p,t}(y_{p,t} + z_{p,t}) + C \sum_{p,t} u_{p,t} && (2.1)\\
\text{s.t.} \quad & \sum_{p \in P_a, s \in S_a} x_{a,s,p} = 1, \qquad \forall a \in A && (2.2)\\
& y_{p,t} \geq \sum_{a,s} x_{a,s,p} \, \mathrm{LBR}_{a,s,t} - \mathrm{TGT}_{p,t}, \qquad \forall p \in P, t \in T && (2.3)\\
& z_{p,t} \geq \mathrm{TGT}_{p,t} - \sum_{a,s} x_{a,s,p} \, \mathrm{LBR}_{a,s,t}, \qquad \forall p \in P, t \in T && (2.4)\\
& u_{p,t} \geq \sum_{a,s} x_{a,s,p} \, \mathrm{LBR}_{a,s,t} - \mathrm{CAP}_{p,t}, \qquad \forall p \in P, t \in T && (2.5)\\
& x_{a,s,p} \in \{0, 1\}, \qquad \forall a \in A, \; s \in S_a, p \in P && (2.6)\\
& y_{p,t} \geq 0, \qquad \forall p \in P, t \in T && (2.7)\\
& z_{p,t} \geq 0, \qquad \forall p \in P, t \in T && (2.8)\\
& u_{p,t} \geq 0, \qquad \forall p \in P, t \in T && (2.9)
\end{aligned}
$$

The objective function (2.1) minimizes the deviation between the target workload for each time period and each port and the scheduled or projected workload for each port, measured in man-hours. This deviation could be positive or negative, so we might choose to represent this difference as

$$\sum_{a,s,p,t} |\mathrm{TGT}_{p,t} - \mathrm{LBR}_{a,s,t} x_{a,s,p}|.$$

Instead, we reformulate the problem to obtain a linear model. The first term of the objective function sums $y_{p,t}$ and $z_{p,t}$. For each time period t and port p the term $y_{p,t}$ refers to the positive side of the absolute value expression, or the projected workload above the target workload. Meanwhile, the second term $z_{p,t}$ refers to the negative side, or the projected workload below the target workload. Together, their sum is the difference between the projected workload from the target workload in a time period at any port. We separate this from the second term, $u_{p,t}$, which is the amount of workload projected at port p in period t above the labor capacity, CAP_t. We weight this overage with cost C to discourage going over the labor capacity for all time periods in T and at all ports in P. The objective function in (2.1) thus represents the deviation from the target workload with a weighted penalty for any workload that exceeds the port's capacity.

Constraint (2.2) ensures each availability gets scheduled to start on exactly one of its feasible starting dates at exactly one of the ports for which it is eligible to be scheduled.

We need four constraints in order to reformulate a MILP from the absolute value problem. Constraints (2.3) and (2.7) together enforce

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

which is the amount of workload projected for port p in period t above the target workload $\mathrm{TGT}_{p,t}$. Constraints (2.4) and (2.8) together enforce

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

which is the amount of workload projected at port p in period t below the target workload $\mathrm{TGT}_{p,t}$. These four constraints ensure that for each $p \in P$ and $t \in T$, either $y_{p,t}$ and $z_{p,t}$ can be positive, but not both. The summation of $y_{p,t} + z_{p,t}$ in the objective function represents the deviation of the workload required at port p in period t from the target workload $\mathrm{TGT}_{p,t}$, and since all terms are positive, we have an equivalent to $\sum_{a,s,p,t} |\mathrm{TGT}_{p,t} - \mathrm{LBR}_{a,s,t} x_{a,s,p}|$.

We want to ensure that the projected workload does not exceed the labor capacity of its assigned port. We consider the constraint

$$\sum_{a,s} x_{a,s,p}\ \mathrm{LBR}_{a,s,t} \leq \mathrm{CAP}_{p,t}, \qquad \forall p \in P, t \in T.$$

In theory, this may work, but it may make the problem infeasible. Instead, we add the elastic variable $u_{p,t}$ to the objective function and penalize work that goes over the port's capacity during time t. To define $u_{p,t}$ so, we add the following constraint

$$u_{p,t} \geq \sum_{a,s} x_{a,s,p}\ \mathrm{LBR}_{a,s,t} - \mathrm{CAP}_{p,t}, \qquad \forall p \in P, t \in T$$

and change the objective function to

$$\sum_{p\in P,t\in T}(y_{p,t}+z_{p,t})+C\sum_{p\in P,t\in T}u_{p,t} \tag{2.10}$$

where $C \gg 1$ is a large constant to heavily penalize any workload that exceeds the labor capacity, $\mathrm{CAP}_{p,t}$. The solver is strongly encouraged to return a feasible solution with $\sum_{p\in P,t\in T} u_t = 0$ if there is one.

2.2 Running the Model

The straightforward way to run this model is to let S_a be the set of all feasible dates to begin availability a. However, as the number of availabilities increases, the planning horizon extends farther into the future, or availability starting dates grow more flexible, the size of our MILP grows. This method would take a long time to compute.

Instead, we can run smaller instances of the model multiple times to speed up the process. The first time we run the model, we limit the size of S_a for each availability to a subset of all feasible dates—only every 7 days. This gives us a close enough solution to we assign each availability to an optimal port.

We run the model again, extending S_a to include all feasible dates, but limiting P_a to include only the port assigned during the first run. Section 3.2.1 further explains the process of limiting the size of S_a and P_a.

2.3 Maintenance Away from Home Port

In the formulation in Section 2.1, availabilities that are eligible for CWB may be assigned with equal preference to any of the ports in the AOR. This section extends the objective function in the MILP model to account for moving the ship away from its home port. To do so, we introduce the following new data and variables:

Data

$\mathrm{LVL}_{p,t}$ The ideal leveled workload at port p in period t for all types of labor. These values are obtained from NAVSEA [man hours / day].

α The penalty incurred due to scheduling availabilities away from the ships' home ports [unitless].

D_a The duration of availability a [days].

$h_{a,p}$ Binary, 1 if port p is not availability a's home port.

To encourage an optimal port assignment in the home port of a ship, we need to determine the time ships spend in maintenance away from their home ports. This can be computed as

$$\sum_{a,s,p} D_a x_{a,s,p} h_{a,p}. \tag{2.11}$$

Minimizing this function creates two competing objectives. The objective in (2.10) minimizes workload deviation, with an additional penalty for exceeding the ports' capacities, over time and is measured in man-days, while the one in (2.11) tries to minimize time away from home port and is measured in days. To reconcile these two objectives, we find normalize each function, then add them. First we normalize (2.10) by dividing the penalized workload deviation term by the ideal level workload, $\mathrm{LVL}_{p,t}$. For a single port p during a single time period t, we get

$$\frac{y_{p,t} + z_{p,t} + Cu_{p,t}}{\mathrm{LVL}_{p,t}}.$$

When we weight this quantity at every port over the entire planning horizon, we get

$$\sum_{p,t} \left(\frac{y_{p,t} + z_{p,t} + Cu_{p,t}}{\mathrm{LVL}_{p,t}} \right) \left(\frac{\mathrm{LVL}_{p,t}}{\sum_{p',t'} \mathrm{LVL}_{p',t'}} \right) = \frac{\sum_{p,t}(y_{p,t} + z_{p,t} + Cu_{p,t})}{\sum_{p,t} \mathrm{LVL}_{p,t}}. \tag{2.12}$$

This quantity measures *workload fluctuation*, a unitless ratio.

Now we reformulate the objective function in (2.11) to also be a unitless ratio. For a single availability, which has exactly one home port, this ratio looks like

$$\frac{\sum_{s \in S_a, p \in P_a} D_a x_{a,s,p} \; h_{a,p}}{D_a}.$$

The weighted average of this ratio over all availabilities—with the weight being each availability's length—is

$$\frac{\sum_{a\in A,p\in P_a,s\in S_a} D_a x_{a,s,p} h_{a,p}}{\sum_{a\in A} D_a}. \tag{2.13}$$

This quantity measures *proportion of time away from home port*, and it is also unitless.

Now the quantity in (2.12) measures the penalized workload fluctuation in ratio and the quantity in (2.13) measures the time away from home port in ratio. Because both of the quantities are now unitless, we can combine them into a single objective function

$$\frac{\sum_{p\in P,t\in T}(y_{p,t} + z_{p,t} + Cu_{p,t})}{\sum_{p,t} \mathrm{LVL}_{p,t}} + \alpha \frac{\sum_{a\in A,p\in P,s\in S_a} D_a x_{a,s,p} h_{a,p}}{\sum_a D_a}. \tag{2.14}$$

In this objective, larger values for α put a heavier cost on performing maintenance away from a ship's home port. Setting $\alpha = 0$ minimizes workload fluctuation, and reduces the new objective function (2.14) to the original objective function (2.1) from Section 2.1.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 3:
Case Study

In this chapter, we discuss several applications of the optimization model introduced in Chapter 2. We focus on six fiscal years from FY 2021 to 2026, namely, 10/1/2020–9/30/2026. We use a planning date of 3/22/2021, which means that availabilities that start before 3/22/2021 are treated as work in progress. Any lead time required to precede an availability that starts after 3/22/2021 will be calculated from this date. Our data come from NAVSEA and contains a set schedule of availabilities that have started on or before 3/21/2021 and a proposed schedule of availabilities that start on or after 3/22/2021. We schedule availabilities in the West Coast AOR, which includes the ports of San Diego (SSD), Seattle (SEA), and Portland (POR).

We compare the workload of SSD, SEA, and POR in four scenarios:

1. Original schedule: This is the original proposed schedule from NAVSEA to meet fleet requirement.
2. Level loading over time: In this scenario, we adjust each availability's starting time to level the workload in each port separately; each availability must stay in its home port.
3. Level loading over time and ports: In this scenario, we assign each availability to one of its eligible ports in the AOR and adjust each availability's starting time to level the workload for all ports in the same AOR simultaneously.
4. Level loading over time and ports with drydock constraint: In this scenario, we first use SSDSP to lock the port and starting time of each availability that requires a drydock, and then assign each nondocking availability to one of its eligible ports in the AOR and adjust its starting time to level the workload for all ports in the same AOR simultaneously.

3.1 Data

In this section, we describe the required data for the model. Data for this tool come from NAVSEA in the form of Excel files. The first file is a spreadsheet containing a proposed maintenance schedule from NAVSEA's Southwest and Northwest Regional Maintenance Centers (SWRMC and NWRMC), and the second file contains multiple spreadsheets describing the labor distribution of an availability based on its duration.

3.1.1 Original Schedule

NAVSEA provides data which represent a proposed maintenance schedule and constraints. The original schedule lists every type of maintenance that shares the same labor pool. This includes every availability, including the ship name, home port, eligibility for CWB, the ideal availability start date, the duration of the availability, the estimated labor required to complete the availability in man-days, and the ports in each AOR. It also includes continuous maintenance (CM) for each fiscal year. A sample of these data can be seen in Table 3.1. Although POR is not a home port for US Navy surface vessel, the Navy can send surface ships to POR for maintenance.

Table 3.1. Maintenance schedule from SWRMC and NWRMC.

Hull	Home Port	CWB	Start Date	Duration	Total Labor	AOR
LCS-A	SSD	0	03/01/2021	92	6250	West Coast
CVN-A	SEA	1	05/01/2021	185	70000	West Coast
CVN-B	SSD	1	06/10/2022	565	71483	West Coast
DDG-A	SEA	0	03/23/2023	147	120292	West Coast
LCS-A	SSD	1	09/05/2023	303	86598	West Coast
LCS-B	SSD	1	01/29/2024	123	88103	West Coast
CM	SSD	0	10/01/2025	–	293811	West Coast
CVN-C	SSD	0	03/14/2026	45	15845	West Coast
DDG-B	SEA	0	09/21/2026	90	24176	West Coast

In this study, we represent the workload of a port using a layer-cake chart. Figure 3.1 shows the workload in San Diego according to the original schedule from NAVSEA from FY21–FY26. The plot represents the resources per day (RPD) measured in man-days on the vertical axis and time (in months) on the horizontal axis. Vertical black lines separate each fiscal year, and a dotted vertical black line indicates the planning date. Each layer represents the work required for an availability over time. The layers are added together to show the aggregate workload, namely, all the labor required from the port.

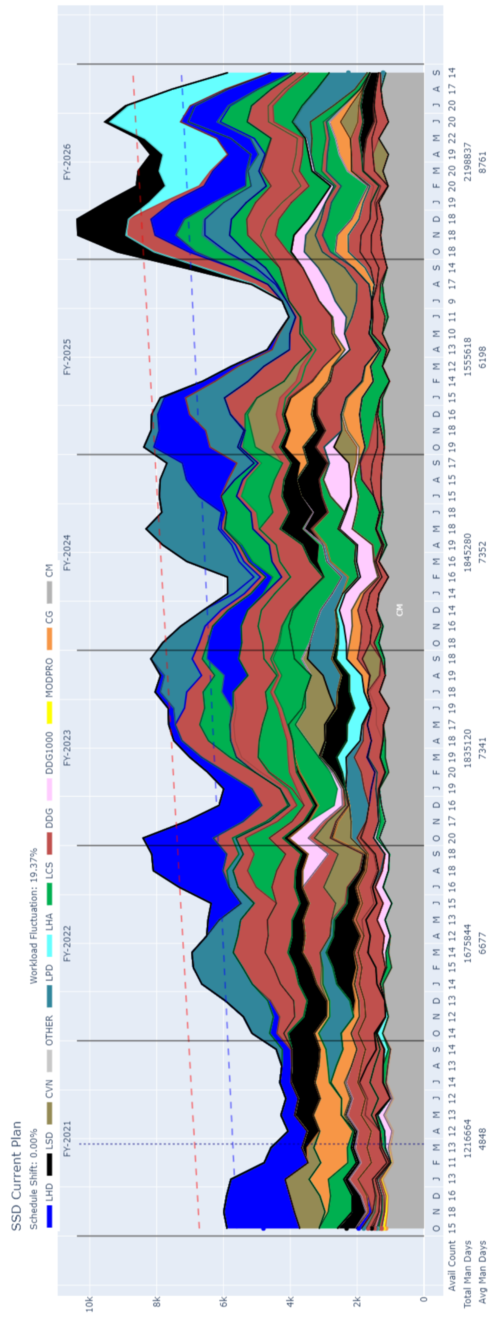

Figure 3.1. SSD workload with original schedule.

3.1.2 Labor Distribution

Other relevant data include the number of production days in each month, the distribution of labor over time based on the duration of the availability, and the distribution of labor for continuous maintenance within each calendar year. The number of production days in a month varies by month and by year. Our data includes the number of production days by month for every year of the case study. Table 3.2 shows a sample of this data for the year 2021. Table 3.3 includes the distribution of labor for availabilities up to six months in duration. Our data includes distributions for availabilities up to 48 months. Table 3.4 shows the distribution of continuous maintenance in a single year. This distribution remains the same in every year of the case study. We combine these data with the data in Section 3.1.1 to determine the RPD required by each availability, where RPD is measured by man-days. Recall that we are seeking to get a steady RPD in order for the shipyards to maintain a steady workforce. The RPD will help us calculate the data parameters $\text{LBR}_{a,s,t}$ required for the model.

Table 3.2. The number of production days in 2021.

Jan	Feb	Mar	Apr	May	Jun
19	19	23	22	20	22
Jul	Aug	Sep	Oct	Nov	Dec
21	22	21	20	20	21

Table 3.3. The monthly distribution of workload for an availability based on its duration, for availability durations up to six months.

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

Table 3.4. 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

3.2 Building Model Parameters

In order to run the multi-port loading model, we need to convert the raw data from NAVSEA to the parameters needed to run the MILP in Chapter 2.

3.2.1 Computing Set Parameters

Of the sets in the MILP, the set A of availabilities, the set T of monthly time periods in the planning horizon, and the set P of ports in the AOR are relatively straightforward. The set A consists of all the availabilities listed in the schedule described in Section 3.1.1. In this case study, which starts on 10/1/2020 and covers 6 fiscal years, the set T includes 72 months from October of 2021 to September of 2026. The set P of all ports in the AOR consists of the ports in the West Coast AOR, namely, SSD, SEA, and POR.

The set S_a consists of all feasible start dates for an availability a. It is defined by (1) the scheduled start date in the proposed schedule given by NAVSEA, (2) the number of days the schedule is allowed to shift, as required by NAVSEA, (3) the planning date, 3/22/2021, and (4) the resolution of our desired schedule. In these scenarios, the maximum number of days an availability can shift is 30 days if the start date is between 1 and 2 years from the planning date and 45 days if the availability is more than 2 years from the planning date. The number of days needed to award a contract is 420 days, so an availability scheduled more than 420 days from the planning date cannot shift left into that window. The planning date in this case study is 3/22/2021, so any availabilities that are originally scheduled on or after 5/16/2022 (420 days after 3/22/2021) cannot be shifted left to 5/15/2022 or earlier. The resolution of our schedule refers to the precision with which we look at potential start dates. It limits the size of S_a for a single availability to improve computational tractability. For example, S_a for a given availability will have fewer potential starting dates with a resolution of one week than with a resolution of one day. The larger the set S_a for each availability a, the longer it takes for the solver to compute an optimal solution.

To demonstrate the determination of S_a, consider the availability for CVN-B in Table 3.1. The original start date of this availability is 6/10/2022. This start date is between one and two years from the planning date, so the window of possible start dates would normally be 5/11/2022–7/10/2022. However, since the original start date is after 5/16/2022 and thus outside the 420 day lead time necessary to award a contract, $S_{\text{CVN-B}}$ cannot include any dates before 5/16/2022. The new window becomes 5/16/2022–7/10/2022. If the resolution is one day, $S_{\text{CVN-B}}$ would contain every day between 5/16/2022 and 7/10/2022: $\{5/16/2022, 5/17/2022, 5/18/2022, \ldots, 7/10/2022\}$ for a total of 55 possible start dates. If the resolution is one week, $S_{\text{CVN-B}}$ would contain dates starting with 05/16/2022 and skipping every seven days through 7/10/2022: $\{5/16/2022, 5/23/2022, 5/30/2022, 6/06/2022, \ldots, 7/4/2022, 7/10/2022\}$ for a total of nine possible start dates. When running the model, we can see how iterating through an eight-date set is much faster than iterating through a 55-date set. Instead of running the model once with a resolution of one day, one way to reduce the run time is to run the model several times with decreasing solutions each time.

For this case study, we use a resolution of seven days for the first run, which we use to assign each availability to one of its eligible ports. The first run returns a new proposed schedule for each port. The new schedules for each port become the input for our second run, in which we run the port loading model for each port separately. The second run refines the schedules for SSD, SEA, and POR separately to assign optimal start dates with a resolution of one day.

The set P_a consists of the ports that availability a can be assigned to, and it is determined by the CWB setting of each availability, as shown in Table 3.1. If CWB is 1, then P_a for availability a is the set of all ports in the AOR, namely $\{\text{SSD}, \text{SEA}, \text{POR}\}$ for the West Coast. If CWB is 0, then P_a only includes the ship's home port.

3.2.2 Computing Data Parameters

The four data parameters in this model, $\text{LVL}_{p,t}$, $\text{TGT}_{p,t}$, $\text{CAP}_{p,t}$, and $\text{LBR}_{a,s,t}$, relate to the labor required or available at the ports.

The parameter $\text{LVL}_{p,t}$ is the ideal leveled workload, which we obtain from NAVSEA. Each port (SSD, POR, SEA, etc.) has an initial amount of resources per day (RPD) as of October 2020—the first month of FY21—which corresponds to the number of workers available at that port. The RPD grows at a linear rate to account for anticipated growth in the maintenance work. These values for initial RPD and yearly growth are summarized in Table 3.5. Ideally, we want to adjust the start date of the availabilities so that their aggregated workload matches $\text{LVL}_{p,t}$ for each month t and each port p.

Table 3.5. Resources available by port, according to SWRMC and NWRMC.

Port	Initial RPD	Yearly Growth
SSD	5600	5.0%
POR	1200	3.0%
SEA	900	2.0%

The next parameter, $\text{TGT}_{p,t}$, is the target workload for future availabilities that will start at port p during month t. The target workload $\text{TGT}_{p,t}$ is the level workload $\text{LVL}_{p,t}$ minus continuous maintenance and the maintenance work during that month t for availabilities that have already started. We calculate the continuous maintenance and the maintenance work from availabilities that have already started using the labor and duration data from Table 3.1, the number of production days per month from Table 3.2, and the labor distribution from Table 3.3. To obtain $\text{TGT}_{p,t}$, we subtract the continuous maintenance and maintenance that has already begun from the ideal leveled load $\text{LVL}_{p,t}$ for that month.

The parameter $\text{CAP}_{p,t}$ is the maximum labor capacity at port p for month t. The Navy wants to avoid having the aggregated workload exceed the ideal workload of a port by more than 20%, so we set $\text{CAP}_{p,t} = 1.2 \times \text{LVL}_{p,t}$

Recall from Section 2.3 that the parameter α is the penalty incurred due to scheduling availabilities away from a ship's home port. In our case study, we set α equal to 0.1. We use a small number here because we want to prioritize the objective of minimizing workload fluctuation. The value of α can increase to prioritize scheduling ships for maintenance in their home ports.

3.2.3 Running the Model

We implement the model in Python and Pyomo, and use the CBC 2.10.3 solver (Forrest 2000) on a personal computer with a 2.60 Intel Core i7 CPU and 8 GB of RAM. The case study consists of 128 availabilities in the planning, among which 45 are eligible for CWB. The first optimization stage, which assigns a port and a tentative start date to each availability, has 405 constraints and 1324 variables, among which 1602 are integer variables. The solver returns a solution with a 0.3% optimality gap with a time limit of 60 seconds.

3.3 Results

In our results, we investigate the differences in workload fluctuation among our four scenarios. The first scenario is a proposed schedule from NAVSEA that meets fleet requirements. The second is NAVSEA's proposed schedule adjusted to minimize workload fluctuation by shifting each availability's starting date. The third is a schedule that minimizes workload fluctuation by adjusting each availability's start date and each eligible availability's port assignment simultaneously. The fourth scenario locks the port assignments and start dates of availabilities that require drydocks and then adjusts the rest of the schedule by shifting availabilities in time and across ports. Table 3.6 summarizes the different features of each scenario.

Table 3.6. The four scenarios and their features.

	Scenario	Shifts Schedule	Allows CWB	Locks Drydocking Availabilities
1	Original schedule			
2	Level loading over time	X		
3	Level loading over time and ports	X	X	
4	Level loading over time and ports with drydock constraint	X	X	X

When we incorporate the SSDSP into our model in Scenario 4, we assume that the contracts for availabilities that require drydocks have already been awarded. Once a contract is awarded, the availability cannot shift in time or across ports. In practice, when we incorporate the drydock schedule into our multi-port loading model, we *lock* the ports and start dates of those drydocking availabilities. Essentially, the set P_a of ports an availability could be assigned to contain just the port where the drydock is located, and the set S_a of possible start dates for that availability contains just the one start date per the contract. These drydock availabilities are a subset of the availabilities in the original schedule. The other availabilities that are coast-wide bid eligible in the original schedule remain coast-wide bid eligible.

We again represent the results of these scenarios in layer-cake charts. In these graphs, we see lines to represent the ideal leveled load, $\text{LVL}_{p,t}$ and the labor capacity at the port, $\text{CAP}_{p,t}$. The blue dotted line represents $\text{LVL}_{p,t}$, or the ideal workload over the entire planning period. It starts at the initial RPD values and increases at the rate given in Table 3.5 to represent the Navy's projected increase in Naval maintenance work. The further the aggregate workload strays from this line, the more fluctuation in the workload. The red dotted line represents $\text{CAP}_{p,t}$, which is 1.2 times $\text{LVL}_{p,t}$ to represent the Navy's guidelines to not exceed the leveled load by more than 20%. We do not want aggregate workload to spike above this line, which indicates that the workload exceeds the Navy's specified limit. Our optimization model will try to shift each colored layer, or availability, to smooth out the graph around the blue $\text{LVL}_{p,t}$ line and prevent spikes above the red $\text{CAP}_{p,t}$ line.

Metrics of interest in this case study include workload fluctuation, schedule shift, and proportion of time ships spend in maintenance away from their home ports. Workload fluctuation is defined in (2.12) and is calculated for every port and the AOR in each scenario. Schedule shift measures the difference between a new maintenance schedule and the original schedule proposed by NAVSEA. Recall that D_a stands for the duration of availability a, and write δ_a for the number of days availability a is shifted away from its original starting date in the new schedule. Define schedule shift as

$$\frac{\sum_a \delta_a}{\sum_a D_a},$$

which is the ratio between the total number of days availabilities are shifted and the total duration of all availabilities. We calculate schedule shift for every port and the AOR in each scenario. Finally, the proportion of time ships spend in maintenance away from their home ports is defined in (2.13) and is calculated for the entire AOR in Scenarios 3 and 4, which involve moving ships away from their home ports.

We present the schedules for each of the four scenarios by port, and then we present the numerical metrics for the entire AOR.

3.3.1 San Diego Schedules

San Diego is the largest Navy base, and its results are displayed in Figure 3.2. Figure 3.2a represents Scenario 1, the workload in San Diego over FY21–FY26 before any optimization, and Figure 3.2b represents Scenario 2, the workload in San Diego after schedule optimization without considering coast-wide bid. The workload fluctuation of Scenario 1, is 19.4%. When we minimize the workload by shifting availability start dates in Scenario 2, the fluctuation decreases to 11.1% with an 8.20% shift in schedule. In the layer-cake chart in Figure 3.2b, less work exceeds the maximum capacity of San Diego shipyards.

When we allow coast-wide bid in Scenario 3, as seen in Figure 3.2c, the workload fluctuation further decreases to 5.65%, and schedule shift decreases to 7.98%. While the schedule optimization in Scenario 2 results in less work that exceeds the red line for $\text{CAP}_{p,t}$ capacity of San Diego shipyard than the original schedule in Scenario 1, the aggregate workload in Scenario 3 only exceeds the red $\text{CAP}_{p,t}$ line once and it hugs the blue $\text{LVL}_{p,t}$ line much more closely.

Figure 3.2d shows Scenario 4, the San Diego schedule optimized over time and ports when we use SSDSP to lock the drydock schedule. The workload fluctuation increases from 5.65% in Scenario 3 to 8.95% in Scenario 4, and the schedule shift decreases from 7.98% to 5.58%. We expect this increase in fluctuation and this decrease in schedule shift because the docking schedule prevents some availabilities from moving in time at all. Locking the dates and ports is a restriction on the workload fluctuation minimization problem, but keeps the schedule closer to the original. The fluctuation in Scenario 4 is still an improvement from Scenario 2, which has a fluctuation of 11.1%.

3.3.2 Seattle Schedules

Seattle is home to five Navy surface ships, and its results are displayed in Figure 3.3. Figure 3.3a represents Scenario 1, the workload in Seattle before any optimization, and Figure 3.3b represents Scenario 2, the workload in Seattle after optimizing by shifting availabilities in time only. From Scenario 1 to Scenario 2, there is a minor shift in schedule, 6.54%. There is a small improvement in workload fluctuation, which decreases from 69.8% in the original

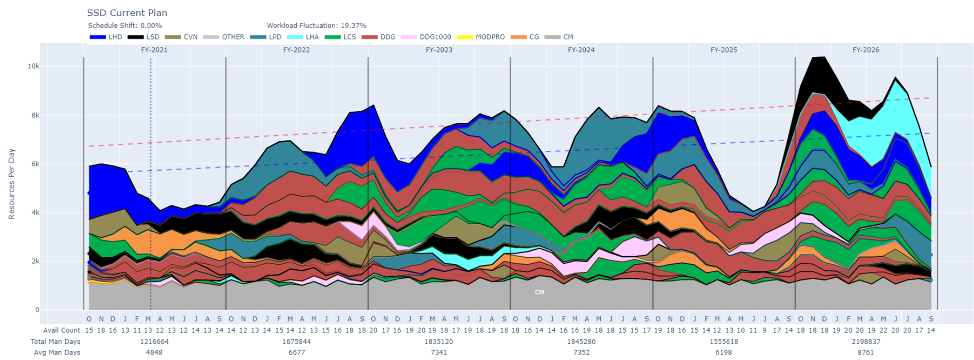

(a) Scenario 1: Original schedule.

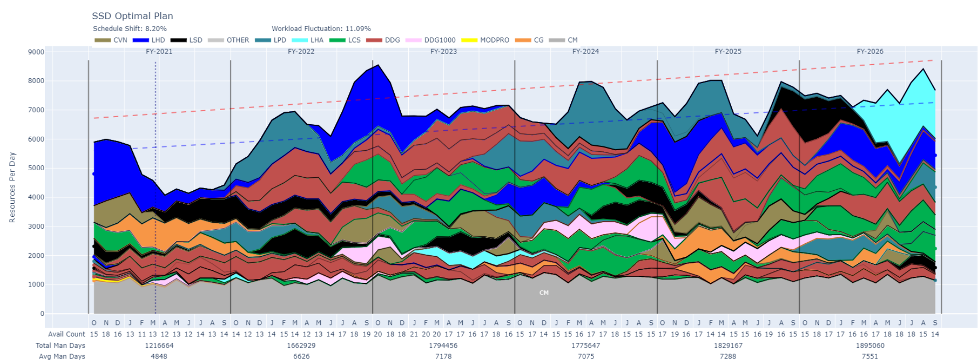

(b) Scenario 2: Level loading over time.

Figure 3.2. SSD workloads.

schedule of Scenario 1 to 61.5% in the optimal schedule of Scenario 2. In Scenario 2, there is also less labor exceeding the capacity of the port, especially in early FY26.

Figure 3.3c shows the workload for Scenario 3, when we allow CWB-eligible availabilities to leave their home ports. Scenario 3's workload fluctuation decreases to 26.7% with a schedule shift of 8.65%. The aggregate workload exceeds the capacity of the port by far less in Scenario 3 than in Scenarios 1 or 2 during the second half of FY21 and the first half of FY22. For this FY21–FY22 time period in Scenarios 1 and 2, there is a huge workload in the original schedule. When we allow availabilities to move out of their home ports in Scenario 3, three LCS availabilities move from Seattle to San Diego, which eases the workload on Seattle and levels out the peak during FY21–FY22. Similarly, in FY25–FY26, the Seattle workload decreases because a CVN availability moves from Seattle to Portland.

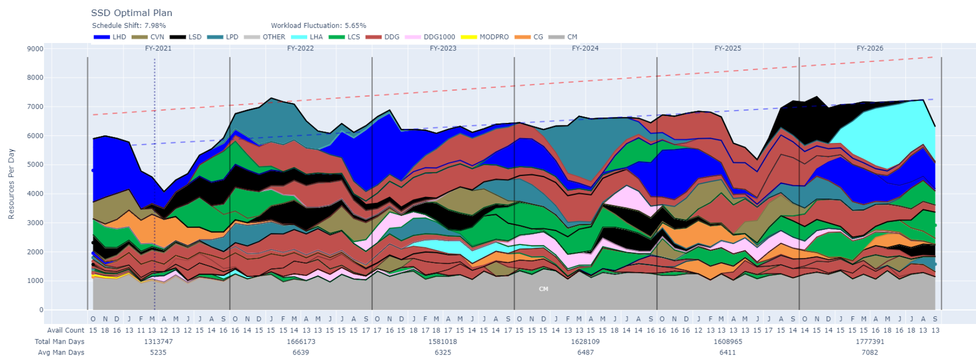

(c) Scenario 3: Level loading over time and ports.

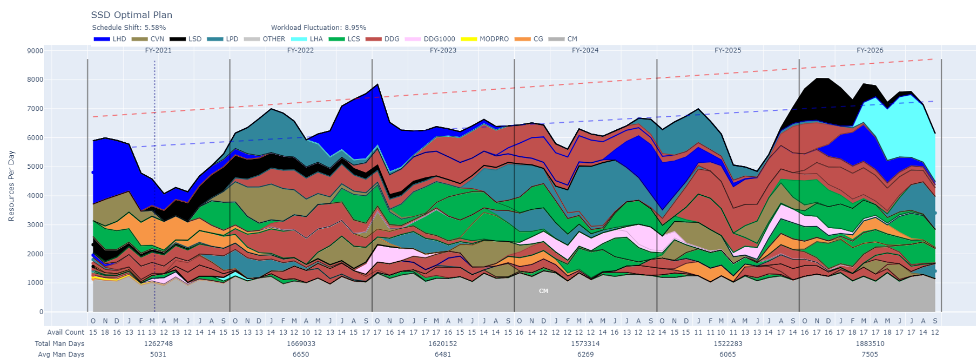

(d) Scenario 4: Level loading over time and ports with drydock constraint.

Figure 3.2. SSD workloads (cont.).

In Figure 3.3d, we see the Seattle schedule optimized over time and ports when we use SSDSP to lock the drydock schedule. The workload fluctuation increases to 40.8% from 26.7% in Scenario 3. The schedule shift decreases to 4.59% from 8.65% in Scenario 3 from Scenario 3's 8.65% to 4.59%. We also see in the layer-cake chart that, from FY21-FY22, the workload in Scenario 4 begins to exceed the port's capacity by more than it does in Scenario 3. However, the fluctuation is still an improvement from the 61.5% fluctuation in Scenario 2.

3.3.3 Portland Schedules

While Portland is not home to any surface ships, it is home to shipyards that can work on surface ship maintenance availabilities. Figures 3.4a and 3.4b represent Scenarios 1 and

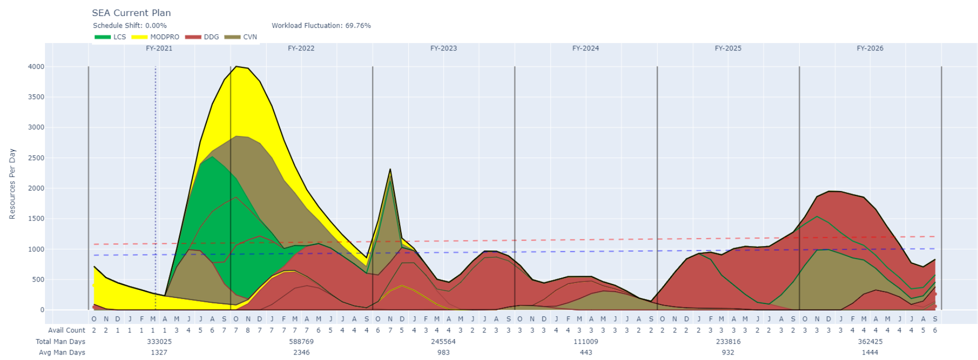

(a) Scenario 1: Original schedule.

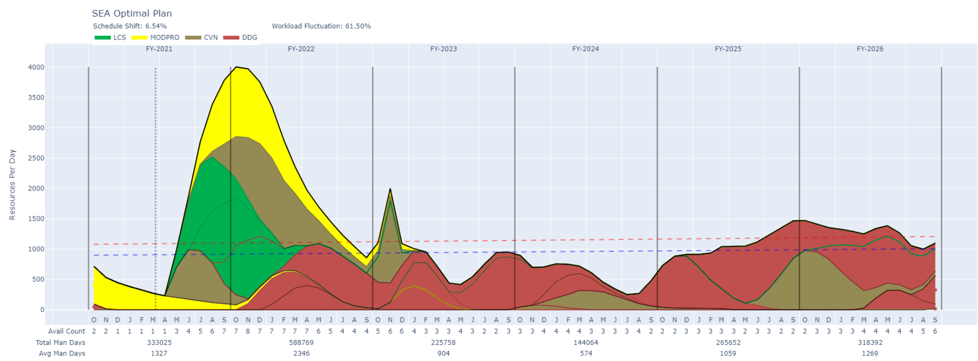

(b) Scenario 2: Level loading over time.

Figure 3.3. SEA workloads.

2, the workload over time in Portland before and after optimization, respectively, without coast-wide bid as an option. The layer-cake charts are identical in Figures 3.4a and 3.4b because the Navy does not have ships with a home port in Portland. The one availability shown had already been scheduled in Portland before the planning date 3/22/2021. Scenario 2 does not allow coast-wide bid, so no other availabilities are relocated to Portland when we optimize the schedule without coast-wide bid. Although this schedule does not ever exceed the ideal leveled load, we see a high workload fluctuation at 89.9%, which does not improve when we optimize over time because the single-port optimal schedule is the same as the original schedule.

Figure 3.4c shows Scenario 3, the workload over time when we allow coast-wide bid. Allowing coast-wide bid lowers our workload fluctuation to 24.1%, with a schedule shift of

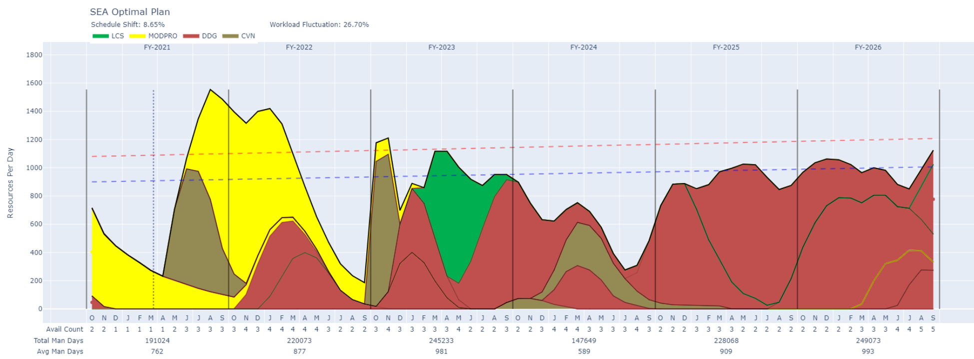

(c) Scenario 3: Level loading over time and ports.

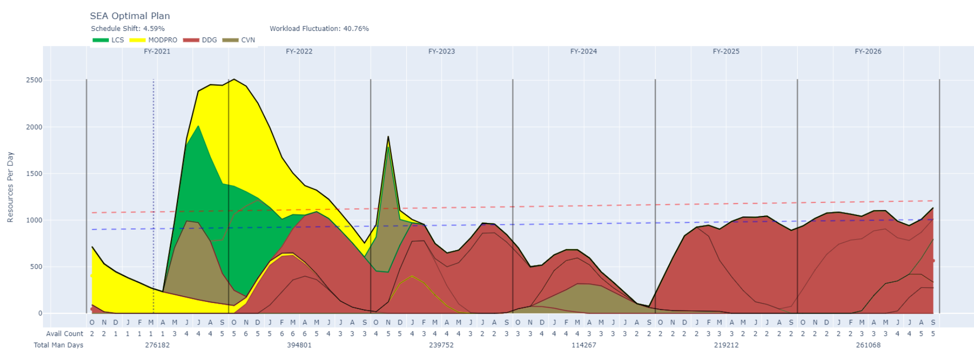

(d) Scenario 4: Level loading over time and ports with drydock constraint.

Figure 3.3. SEA workloads (cont.).

7.34%. Allowing coast-wide bid also helps better utilize the resources available in Portland. In Scenario 3, 10 availabilities that are originally from San Diego and three availabilities from Seattle are reassigned to Portland.

Figure 3.4d represents Scenario 4, the Portland schedule with minimized workload fluctuation by adjusting availability start dates and ports with a locked drydock schedule. The workload fluctuation increases from 21.4% in Scenario 3 to 40.8%. The schedule shift decreases from 7.34% in Scenario 3 to 6.45%. In this scenario, 13 availabilities originally from San Diego and one availability from Seattle are reassigned to Portland. Although the workload is not as leveled in Scenario 4 as it is in Scenario 3, it is still a significant improvement from Scenario 2.

Overall, while shifting availabilities in time helps lower the workload fluctuation in a single

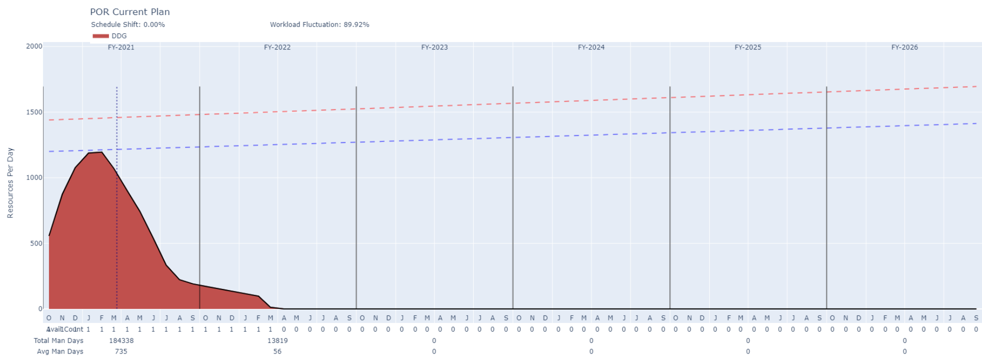

(a) Scenario 1: Original schedule.

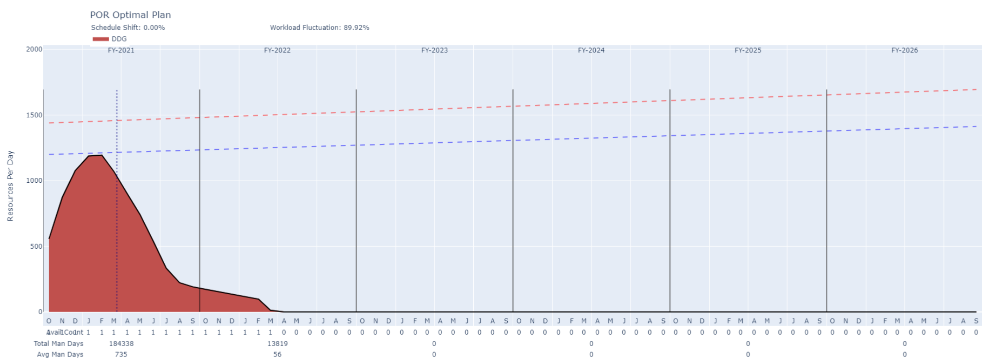

(b) Scenario 2: Level loading over time.

Figure 3.4. POR workloads.

port, allowing coast-wide bid in addition to time shifts decreases workload fluctuation even more. Meanwhile, although the schedules for each port in Scenario 4 have higher workload fluctuations, they tend to have less schedule shift, which shows how our model can be adapted to more specific requirements beyond coast-wide bid eligibility. Tables 3.7 and 3.8 summarize the metrics for workload fluctuation and schedule shift, respectively.

3.3.4 Area of Responsibility

Recall that our formulation in Section 2.3 involves competing objectives between minimizing workload fluctuation and minimizing the proportion of time ships in the AOR spend outside of their home ports. Table 3.7 displays the workload fluctuations for the AOR and individual ports in each scenario. When we calculate the proportion of time ships spend

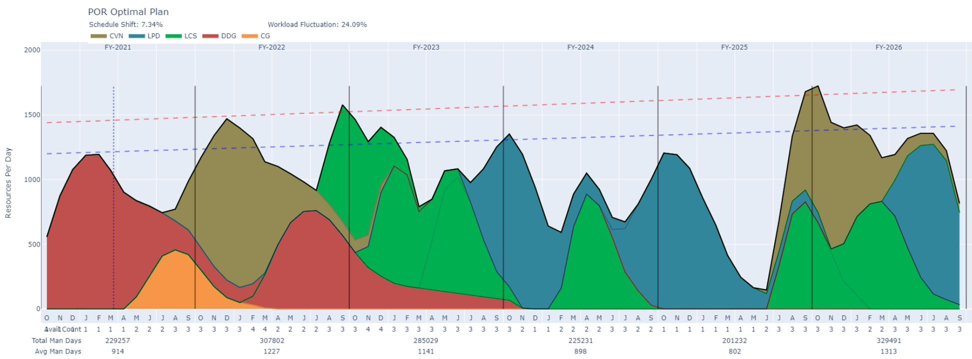

(c) Scenario 3: Level loading over time and ports.

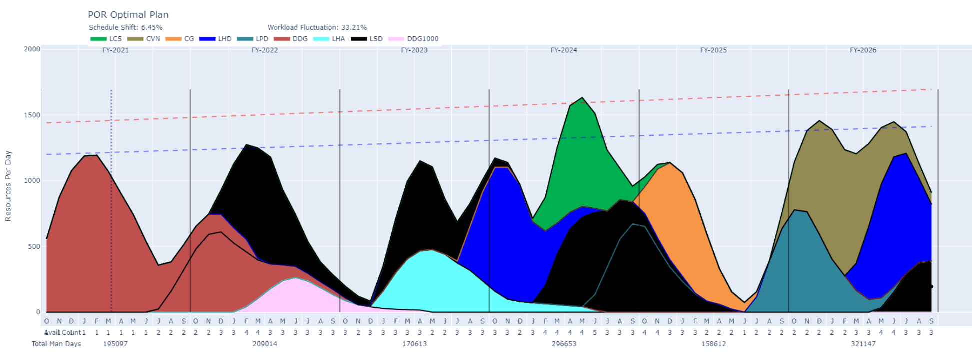

(d) Scenario 4: Level loading over time and ports with drydock constraint.

Figure 3.4. POR workloads (cont.).

outside their home ports, we only include time in maintenance for availabilities that start on or after the planning date. In this case study, although Portland has one availability with a home port of Seattle already in progress, we do not include it when we calculate the proportion of time ships spend out of their home ports. With that said, for Scenarios 1 and 2, none of the availabilities that start on or after 3/22/2021 are assigned to a port other than their home port. When we allow coast-wide bid in Scenario 3, this proportion of time increases to 11.2% to allow the workload to level across ports. In Scenario 4, the proportion increases to 12.7%. The value of α can be adjusted to compare optimal solutions that trade off workload fluctuation and out-of-home-port maintenance.

To measure the difference between each new schedule and the original schedule, we also calculate schedule shift. Table 3.8 displays the schedule shifts for the AOR and for each port

Table 3.7. The workload fluctuation for each scenario (%).

Scenario	Location			
	West Coast AOR	SSD	SEA	POR
1	14.0%	19.4%	69.8%	89.9%
2	10.9%	11.1%	61.5%	89.9%
3	8.00%	5.65%	26.7%	24.1%
4	10.7%	8.95%	40.8%	33.2%

Table 3.8. The schedule shift for each scenario (%).

Scenario	Location			
	West Coast AOR	SSD	SEA	POR
1	0.00%	0.00%	0.00%	0.00%
2	7.86%	8.20%	6.54%	0.00%
3	9.06%	7.98%	8.65%	7.34%
4	5.54%	5.58%	4.59%	6.45%

in the four scenario.

The proposed schedule from NAVSEA contains 153 availabilities, of which 25 are already in progress by 3/22/2021. Of the remaining 128 availabilities, 17 availabilities in Scenario 3 receive port assignments that are not their home ports. In Scenario 4, 16 availabilities receive port assignments that are not their home ports. Table 3.9 summarizes these assignments. The numbers of availabilities that move from one home port to any of the three ports in the AOR are comparable between Scenarios 3 and 4.

Table 3.9. Number of availabilities that receive maintenance away from the ships' home ports for scenarios 3 and 4.

Scenario	Home Port	Assigned Port		
		SSD	SEA	POR
3	SSD	94	1	10
	SEA	3	19	3
4	SSD	92	0	13
	SEA	2	20	1

CHAPTER 4: Conclusion and Future Research

This thesis presents a model that schedules maintenance availabilities to minimize workload fluctuation over time and across multiple ports simultaneously. The model's input is a proposed maintenance schedule with availabilities and their start dates, durations, home ports, coast-wide bid eligibility statuses, and labor requirements. The model uses mixed integer linear programming to shift availability start dates and assign ports to return an optimal schedule. The competing objective to minimize the proportion of time ships spend in maintenance outside of their home ports can be used to balance the costs and benefits of low workload fluctuation and moving ships away from their home ports.

In Chapter 3, we demonstrate the effectiveness of the multi-port loading model. By itself, the model can produce an optimized schedule that considers both the Navy's operational requirements and private shipyards' capacities. We demonstrate this capability in two scenarios, when we produce optimal schedules from a proposed NAVSEA schedule. We compare the workload of the proposed NAVSEA schedule to the workload when we shift availabilities in time only, shift availabilities in time and allow coast-wide bid, and shift availabilities in time, allow coast-wide bid, and lock drydock availabilities according to the SSDSP in Hilliard (2019) .

Fleet planners can use this model to gain more insight about the maintenance scheduling process. Below are a few questions fleet planners may ask: How might the Navy need to adjust its maintenance schedule if the Navy's private contractors do not increase their labor capacity as the Fleet grows? How can the Navy adjust its maintenance schedule if a ship suddenly requires unplanned maintenance (due to a fire, collision, etc.)? Where does labor capacity fall short that might require the Navy to investigate other options to make maintenance more efficient? To address these questions, fleet planners might change the parameters of the model, such as the initial resources of a port and its yearly growth. By adjusting the weight coefficient α from (2.14), it is also possible to compute different Pareto optimal solutions to study the tradeoff between workload fluctuation and the proportion of time ships spend in maintenance outside of their home ports.

The MILP techniques used to schedule ship maintenance in this thesis can be extended to other Navy processes, such as scheduling submarine or aircraft maintenance, personnel assignment, and fleet management. The multi-port loading model developed in this thesis demonstrates the potential how data analytics and mathematical programming can help the Navy improve other tactical operations, such as search, combat, logistics, and surveillance.

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.

Forrest J (2000) CBC, version 2.10.3. Accessed March 10, 2020, https://projects.coin-or.org/Cbc.

Hart W, Watson JP, Woodruff D (2008) Pyomo, version 5.6.1. Accessed March 10, 2020, 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.

Hilliard A, Lin KY, Sparks RD (2020) Surface ship drydock schedule planner. *Naval Engineers Journal* 132(1):63–72.

Mackin M (2017) Navy Ship Maintenance: Action Needed to Maximize New Contracting Strategy's Potential Benefits Action Needed to Maximize New Contracting. Technical Report GAO-17-54, GAO.

Martin B, McMahon M, Riposo J, Kallimani J, Bohman A, Ramos A, Schendt A (2017) *A Strategic Assessment of the Future of U.S. Navy Ship Maintenance: Challenges and Opportunities*. ISBN 9780833099235, URL http://dx.doi.org/10.7249/rr1951.

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.

Oakley SS (2020) Navy Ship Maintenance: Evaluating Pilot Program Outcomes Could Inform Decisions to Address Persistent Schedule Challenges. Technical Report GAO-20-370, GAO.

Office of the Chief of Naval Operations (2019) Maintenance Policy for Navy Ships. Technical Report OPNAVINST 4700.7M, Washington, DC, https://www.secnav.navy.mil/doni/Directives/04000 Logistical Support and Services/04-700 General Maintenance and Construction Support/4700.7M.pdf.

O'Malley L, Lin KY (2020) A port loading model to minimize workload fluctuation for surface ship maintenance jobs. *Military Operations Research* (in press).

Parmer C (2016) Dash, version 0.39.0. Accessed March 10, 2020, https://plot.ly/products/dash/.

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 10, 2020, https://www.python.org.

Initial Distribution List

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California