Air Force Crew Scheduling: An Integer Optimization Approach

by

Matthew J Koch

B.S. Operations Research, U.S. Air Force Academy

Submitted to the Sloan School of Management
in partial fulfillment of the requirements for the degree of

Master of Science in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2021

Author . . . .
Sloan School of Management
June 2021

Approved by . . . .
Allison Chang
MIT Lincoln Laboratory
Technical Supervisor

Certified by . . . .
Hamsa Balakrishnan
William E. Leonhard (1940) Professor of Aeronautics & Astronautics
Thesis Supervisor

Accepted by . . . .
Georgia Perakis
William F. Pounds Professor of Management Science
Co-director, Operations Research Center

Research was sponsored by the United States Air Force Research Laboratory and the Department of Air Force-MIT Artificial Intelligence Accelerator and was accomplished under Cooperative Agreement Number FA8750-19-2-1000. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the United States Air Force or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

Air Force Crew Scheduling: An Integer Optimization Approach

by

Matthew J Koch

Submitted to the Sloan School of Management
on June 2021 in partial fulfillment of the
requirements for the degree of
Master of Science in Operations Research

Abstract

Air Force flight, training, and crew scheduling is a labor-intensive and largely manual process across all flying squadrons. Complex training requirements and dependencies, operational constraints, numerous qualifications, and unforeseen missions confound the schedule development process. We develop multiple optimization formulations for the Air Force crew scheduling problem. Furthermore, we present multiple objective functions aiming at mimicking reality to account for pilot qualification upgrades and their ability to stay current and mission ready. To compare candidate schedules, we identify numerous metrics that show the impact of the different objective functions. Finally, we briefly discuss how to incorporate scheduler preferences and focus on creating human-interpretable schedules so that the scheduler can select the most desired schedule for the squadrons' current needs.

Technical Supervisor: Allison Chang
MIT Lincoln Laboratory

Thesis Supervisor: Hamsa Balakrishnan
Professor and Associate Department Head of Aeronautics and Astronautics

Acknowledgements

There are many people in my life that I would like to thank that not only made this part of my life possible, but made this experience a truly incredible one.

First and foremost, I would like to thank my MIT advisor, Hamsa Balakrishnan, for her endless support. I am fortunate to have developed, both intellectually and personally, from such a passionate and knowledgeable expert.

I also want to thank my MIT Lincoln Laboratory advisor, Allison Chang, for her feedback and support throughout my time at MIT. Additionally, I owe a ton of credit to the entire MIT AIA team for the friendships and support they provided. I especially want to thank Chris Chin for all of his hard work, assistance, and feedback for the algorithms and formulations in this thesis.

Thank you to all the faculty, students, and members of the ORC for helping me along the way. I want to specifically thank the "Military Crew," Dan, Jon, and Thomas, for all the help in courses and research, and for their friendships and providing a fun environment to work.

It is without question I would not have been awarded this opportunity without the support of the United States Air Force and MIT Lincoln Laboratory. I thank them both for giving me this opportunity and I looking forward to serving and continuing to focus on making a positive impact.

Last but not least, I am forever grateful for the support of my family, particularly my parents, for their endless love and dedication in all my aspirations. I especially thank Haley for her continuous love and patience over the years.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The United States Air Force has a complex resource scheduling problem. They operate globally, even in remote regions of the world, with a diverse set of missions. There are fighter jets for combat operations, cargo planes to transport people and equipment, tankers to provide in-flight fuel to the other aircraft, and many others. Some Air Force missions have the benefit of a regular demand pattern, but most arise unexpectedly. Thus scheduling pilots to fly missions, while also ensuring that they keep up with their training requirements, and are allowed sufficient leave time is a complicated optimization problem. Like many legacy practices, flight scheduling is currently an extremely non-standardized, manual, and tedious process that relies on several different data sources to ensure a flight is legally able to fly. In every flying squadron, schedulers are often pilots themselves, and when they are not flying, typically spend eight hours per day assigning crews to flights.

With the recent push within the Department of Defense and the Air Force for automated methods and artificial intelligence, there is an opportunity to significantly enhance the scheduling process [7, 10]. Puckboard is a web-based software application aimed at automating scheduling for C-17 crews, while having all necessary data in one centralized location. Puckboard is one of many recent initiatives across the Air Force to move towards a more digital and agile Air Force. It was created by TRON, which is a part of the Aloha Spark Office at Hickam Air Force Base (AFB) in Honolulu, Hawaii. TRON's focus is to build software that will reduce administrative tasks and allow more time for Airmen to execute their missions. The objective of our research is not to replace schedulers, but to design optimization-based algorithms, to be transitioned to the Air Force through Puckboard, that will greatly decrease the time spent on the labor-intensive and largely manual process of scheduling so that the Airmen can focus more on their primary missions. Furthermore, the longer-term goal of Puckboard is to create an application used in every flying squadron across the Air Force and eventually be used to coordinate with outside squadrons and agencies.

1.1 Contributions

We make the following contributions in this thesis:

- We present our baseline integer optimization formulation for the Air Force crew scheduling problem, with a few different objective functions for the scheduler to choose from.

- Next, we provide an extension to our baseline formulation that incorporates a pilot's flight training requirements such that they are able to work towards qualification upgrades and staying constantly qualified. This formulation thus generates schedules that more closely mimic an actual schedule in a typical flying squadron. Similar to our baseline model, we also present a few different objective functions to include the pilot's flight training requirements.

- Lastly, we present metrics and graphics to allow the schedulers to analyze the objective functions to see how a specific schedule affects a squadron as time progresses.

1.2 Thesis Structure

The structure of this thesis is as follows. In Chapter 2, we describe the basic considerations that go into assigning crews to a flight schedule. We also compare how the process currently looks in an Air Force squadron, how Puckboard works right now, and then our team's goals and approach. Additionally, we provide a brief literature review. We discuss all the data and information necessary for creating a feasible schedule in Chapter 3, while further detailing our team's approach and methodologies. In Chapter 4, we detail the mathematical formulations to solve the Air Force crew scheduling problem. Next, in Chapter 5, we look at a specific squadron's personnel and historical flights from a six month time period to analyze the effects of our methodologies based on proposed relevant metrics and visualizations. Finally, in Chapter 6, we present our conclusions and key takeaways, and provide some thoughts and avenues for future work.

Chapter 2

Background and Literature Review

Puckboard is currently focused on only one type of air-frame, the C-17, which is the closest plane the Air Force has to a commercial airliner with respect to the size and capabilities of the plane. It can be used to transport people and/or cargo. A C-17 is staffed with mainly two different types of positions: pilots and loadmasters. Pilots fly the aircraft, and loadmasters are in charge of the part of the aircraft where the people and cargo are located. Both pilots and loadmasters have different levels of qualifications that designate if they are able to fly on various flights. The three general flight events that will be discussed throughout this thesis include simulators, training (commonly referred to as locals), and missions, which all require specific qualifications of personnel depending on the task(s) to be executed on the flight. The objective for the scheduler is then to take all of this into consideration to create a feasible and legally qualified flying schedule.

Typically, simulators and locals are scheduled months ahead of time. In particular, simulators are managed through contractors, and the Air Force covers a non-refundable cost for simulator flights well in advance, so the schedule is also set well in advance. Missions, on the other hand, are usually given to a squadron scheduler only two to three weeks in advance, and sometimes there are missions that pop up within just a few days of notice. When these urgent missions appear, the scheduler is short on time and the scheduling process becomes very hectic. Due to these unforeseen flights, which occur fairly often, schedulers usually only schedule crew to flights at most two weeks in advance to prevent the need to have to adjust a large number of flights. For these reasons, an automated tool like Puckboard can be very helpful and make a great impact in the reduction of time a scheduler spends scheduling crews.

The remainder of the chapter is organized as follows. In Section 2.1, we highlight the original planning process and where the name Puckboard came from. We detail how most squadrons' current scheduling processes work in Section 2.2. We explain the current use of Puckboard in Section 2.3, along with challenges and shortfalls. We initially propose our planning process in Section 2.4. Lastly, we conclude in Section 2.5 with a literature review of the pertinent research describing solution approaches to similar problems.

2.1 Original Planning Process

Every squadron operates and schedules flights based on what works best for their personnel. However, for many years, flying squadrons and their leadership would congregate around a white board and lay out all the flights for the upcoming week(s). Each member in the squadron was then assigned a designated "puck," and then the leaders and schedulers spent an inordinate amount of time assigning personnel to all the possible flights, deconflicting everyone's schedules by hand. The notion of moving "pucks" around a white board is where the name Puckboard came from.

2.2 Current Planning Process

Our team was fortunate to be able to visit a C-17 squadron at Hickam AFB and get a glimpse into one specific squadron's operating procedures. We had the chance to interview schedulers and witness the way their current scheduling process is handled. Like most flying squadrons across the Air Force, they currently use a Microsoft Excel based scheduling tool, which displays the flights, the dates, and who is assigned to the flight. There is a separate data source that displays the details for simulator flights, and different databases that are home to necessary personnel information. Typically, an Airman can request to be on a flight if there are open positions, but if a scheduler needs to fill empty slots, they have to call, text, or personally talk to a specific pilot or loadmaster to see if they are available. This can be a very tedious and time consuming task, especially when a squadron receives an unexpected mission that is to be flown in the next few days. A scheduler must then cross reference all of these different data sources and contact available personnel to successfully schedule these urgent missions.

2.3 Puckboard Planning Process

Puckboard currently can assign crews to a batch of flights within a desired time window accounting for three constraints: 1) a crew member cannot fly when they are unavailable due to leave, temporary duties (TDYs), appointments, etc., 2) each person can only be in one place at a time, and 3) each flight must have the correct number of qualified personnel to fill the different roles required to be legally safe to fly. When it comes to the third constraint, Puckboard treats all loadmasters as if they have the same qualification and only considers two different qualifications for pilots: whether or not they are qualified to be an Aircraft Commander (AC). In reality there could be upwards of 30 different combinations of qualifications at a specific squadron for pilots depending on the flight and task(s) they are assigned to accomplish on that particular flight (more detail on this to come in subsequent sections). Additionally, the application allows for the schedulers to lock in crew members when a specific member requests to be on a flight, or when a scheduler confirms with a member that they will accept a specific flight assignment. Puckboard then uses a simulated annealing based heuristic through OptaPlanner to automatically assign available crew to the desired flight.

However, the basic constraints and modeling put in place right now do not provide results that mimic the reality of a schedule actually used in practice. Thus, the squadrons that have already adopted Puckboard, mostly just use the basic functionality of schedule visualization and the benefit of all necessary data located in one central location. Therefore, there is still plenty of room to improve the automation capabilities and the optimization algorithms to provide more meaningful solutions.

2.4 Proposed Planning Process

Our goals for the overall Air Force scheduling problem are two-fold: 1) enhance Puckboard's automation capabilities for the crew scheduling problem, and 2) expand Puckboard to automatically schedule the flight events and the specific event types to enhance schedule efficiency and maximize training accomplished. In this thesis, we analyze and propose optimization methods for the first task, commonly referred to as the crew rostering problem. Throughout this paper, we focus only on the pilot problem space as it is more complicated from the vast number of qualifications and there are more pilots required for each flight compared with loadmasters. However, our methods can easily expand to include loadmasters as well. See the subsequent chapters for a more detailed explanation of the team's approach.

2.5 Literature Review

Workforce scheduling is a classic optimization problem which dates back to the 1950s. The development of optimization methods that can identify or approximate optimal solutions for complex scheduling problems in practical computation times remains an active area of research. The challenge when it comes to these types of problems in the real world is that each problem typically provides its own unique characteristics not captured in standard formulations. With Puckboard, some unique challenges include the complicated hierarchical qualification structure, the crew requirements for the combination of different flights and tasks, urgent missions that disrupt previously scheduled missions, pilot training requirements, and the fact each pilot possesses additional duties outside of flying.

Very little work has been documented on the military scheduling problem and most prior efforts have studied the Naval flight training scheduling problem. One of the first papers that studied the military scheduling problem [15], encountered many of the same constraints relevant to the scheduling problem today; Honour used a network formulation, which is not used in this thesis, but could potentially be a valid alternative explored to represent our assignment of pilots to flights. Hall *et al.* [12] analyzes current scheduling practices in the Navy, many which are analogous to the Air Force scheduling problems, however, their analysis is primarily focused on a pilot training scenario where their main objective is to train pilots in the fastest way possible. Both Jacobs [16] and Slye [21] presented mixed integer linear optimization formulations to different Navy scheduling problems, where we were able to leverage some of the basic constraints in our formulations. Jacobs focused on the Navy pilot training problem with really only two different kinds of qualifications (student pilots and instructor pilots) and an objective of minimizing the time each student spends at pilot training. Additionally, Slye more heavily focused on the scheduling of flight events due to constrained resources rather than the crew scheduling aspect of the problem. These efforts have looked at slightly different circumstances and there has not been an attempt to elicit and accommodate user preferences in the process or to generate human-interpretable schedules, both of which we believe are essential for practical adoption. These articles provided a good starting point for certain constraints, but we quickly diverged from these efforts due to the unique challenges mentioned above.

There has been significantly more research discussed on scheduling in the commercial airline industry. Typically their scheduling process occurs in roughly three stages: 1) demand modeling, 2) flight scheduling (number of flights, flight legs, flight times, etc.), and 3) crew scheduling. Here we focus on the last stage of their scheduling problem, which is usually broken up into crew pairing and crew rostering [17]. Crew pairing involves assigning a pairing of flights or legs that start and end at the same base, and crew rostering refers to assigning crew to these pairings under specific constraints. Most of the crew scheduling research has been done on the crew pairing phase and not necessarily the crew rostering stage, although more and more research on rostering is starting to appear. Caprara *et al.* [5] propose generic models and formulations for a variety of crew rostering problems (CRP), that provide useful insights into different constraints necessary in such problems, and offer ideas in modeling them as linear functions. Zhang *et al.* [26] provide more detailed research on specifically the airline crew rostering problem (ACRP) and offer several models and formulations in solving the ACRP. Additionally, they provide algorithms that help balance the trade-offs between hard and soft constraints present in all crew scheduling problems. The overall goal of the ACRP is usually to minimize the overall cost for the specific company, but Zhou *et al.* [27] describe ways to account for crew fairness and maximizing the satisfaction of their employees. Furthermore, they provide useful insights into multi-objective optimization formulations, which is briefly discussed in our research, where we ultimately try to have a reasonable work-life balance for the military members. While the military scheduling problem has some direct parallels to the commercial airline's scheduling problem, there are some stark differences and challenges that we are presented with. Ultimately, the commercial airline industry is able to leverage pattern based flight schedules and the crew pairing phase, both of which simplify the crew rostering phase, that the military scheduling problem does not have the luxury of leveraging. Also, the complex qualification hierarchy, abundant training requirements, and the pop-up missions present in the military scheduling problem provide further complications not necessarily present in the commercial airline scheduling problem.

Scheduling problems are not unique to only airline problems. Other common industries and applications where scheduling literature arises is healthcare [13, 18, 20], aircraft maintenance [14, 23], distribution and call centers [3, 19], and transit industries [8, 9]. Additionally, there are popular literature surveys on all types of scheduling problems as a good starting point for more research [17, 22].

Chapter 3

Scheduling Problem Inputs

As mentioned in previous chapters, information and relevant data are needed from various data sources to solve the scheduling problem. Data is currently being pulled and used from multiple databases within the Air Force, which include Aviation Resource Management System (ARMS), LeaveWeb, Graduate Training Integration Management System (GTIMS) and Global Reach. Global Reach is a C-17 specific database, so we heavily rely on a manually provided subset of the data for our modeling approaches. The relevant data within all these different data sources include historical flights (dates/type/crew type/etc.), legally flyable crew requirements for each type of flight, pilot training requirements, pilot evaluation dates, pilot unavailability, and pilots' qualifications. The following provides a more detailed discussion on the intricacies within the given and necessary data.

We begin this chapter by introducing the necessary inputs and data to consider in order to solve the scheduling problem in Sections 3.1 and 3.2. In Section 3.3, we introduce some assumptions to simplify the problem. We conclude in Section 3.4 with a discussion on how we solve the Air Force crew scheduling problem qualitatively, explaining the relevance of the objective functions and necessary constraints.

3.1 C-17 Flight Data

The three flight events are simulators, locals, and missions. These flights can be broken down further into different flight types: Air Land (AL), Airdrop (AD), and Special Operations Low Level II (SOLL II). AL flights are the most common type of flight (usually the assumed type if not specified) and do not require special qualified pilots, whereas the AD and SOLL II flights require the pilots to have special qualifications. Additionally, all types of flights may require Air Refueling (AR), if the flight happens to have a long flying time. This leads to the discussion of required crew to fly on each one of these flights. Normally there are two types of crews: either augmented or basic. Augmented crew flights require a minimum of three pilots to fly a flight and basic crew flights have a minimum requirement of two pilots. Within the augmented crew types, there are special augmented designations if the flight requires AD, SOLL II, or AR. Thus, the pilots with special qualifications are required for these more specialized flights.

The schedulers are given information from parent organizations like Air Mobility Command (AMC) for the different flights a particular squadron needs to accomplish. For simulator and local flights, the scheduler is usually given the date, start time, and type of flight training. For each mission, the scheduler is given the entire plan including required crew, any other special requirements, itinerary, full schedule of stops and locations, etc. Then once a scheduler has all of this information, they must take into account all personnel information to assign necessary crew to meet the minimum flight requirements.

3.2 Personnel Data

The most important personnel data discussed above are unavailability dates and pilot qualification designations. There are various reasons a pilot might be unavailable: leave, temporary duties (TDYs), additional duties, appointments, etc. Unavailability is usually described with a start and end date, and sometimes could include specific time periods if it is simply a short medical appointment. A pilot's qualification is typically a five-character designation (e.g., *FPCC5*), where the first three letters refer to their AL qualification, the fourth letter is their training level, and the last digit is their special duty qualification. The training level qualification (usually designated as level A, B, or C) is typically set by the squadron commander to establish the more experienced pilots and those who require fewer flight hours to stay current or become recurrent. The special duty qualifications are for flights that are tasked to execute AD, SOLL II, or other atypical tasks embedded within a squadron's specific mission set (e.g., landing on ice in Antarctica). For C-17 crews there are generally four distinct groups of qualifications. From highest to lowest qualification group they are: evaluator pilots (EPs), instructor pilots (IPs), mission pilots (MPs), and flight pilots (FPs). Based on discussions with C-17 pilots, a typical C-17 squadron has about 70 pilots, where there is roughly 7 EPs, 10 IPs, 15 MPs, and 38 FPs.

Also discussed above in regards to personnel data are the pilot's semi-annual or periodic training requirements and their different evaluation dates. Pilots have required types of flights they must fly semi-annually or periodically to stay current with their qualifications and ability to fly. So usually, pilots prefer to fly on flights that help them to constantly work towards meeting their training requirements. For instance, a specific pilot would be against flying the same type of simulator over and over again because that repetition would not help them progress in their flying career. Furthermore, pilots typically like to be in control of their own schedule for various reasons, so they will voluntarily request to fly flights they are interested in or are necessary to continue to move up the ranks.

3.3 Assumptions

In this thesis we make a few assumptions to simplify the proposed methods. First, we assume each pilot can only fly one event per day. In reality, a pilot can fly two simulators in the same day, but it does not happen very often and really only happens when a specific squadron is struggling to find available pilots. In determining pilot availability, we also have to account for crew rest, which is the rest time a pilot is legally required to take after completing a flight. There is typically no crew rest for simulator flights, but for training and mission flights, the necessary crew rest can be from hours to days [1]. We assume crew rest is taken care of for the locals based on our first assumption that only one flight can be flown per day. However, for the missions, we are given historical data on how long a pilot is on crew rest for based on the day the postmission debrief occurs, so we assume that accounts for crew rest for those flights. Scheduling a future event has a projected duration, but of course this time is truly uncertain. Thus, adding in a buffer time, like in the commercial airline industry, and using crew rest regulations allow us to obtain a projected duration that a pilot is unavailable both for the flight and the accompanied crew rest.

3.4 Model Description

Integer optimization is a natural way to model the CRP, where our goal is to assign pilots to flights. The three most important constraints given a set of flights are that we need sufficient pilots, both the quantity and the correctly qualified pilots, and that no pilot can be assigned to overlapping flights.

Talking to subject matter experts, many whom are or were schedulers in their respective squadrons, describe their conflicting and evolving priorities when creating schedules. They explain their main goal is to make sure each flight is fully staffed, but after that they consider two other main priorities. First, they aim to keep their most qualified pilots, EPs and IPs, close to their home base as much as possible, as these pilots are the only ones able to sign off on training flights. Second, schedulers try to schedule pilots to a variety of different types of flights so that they are able to stay current with their different flight training requirements. Taking these priorities into account, we offer different objective functions into our optimization formulations so that schedulers can choose the objective function that best tailors to their immediate needs. Moving forward with some of these ideas, we present our mathematical formulations in the following chapter.

Chapter 4

Formulation

The Air Force crew scheduling problem can be solved as an assignment based integer optimization problem. The goal of this problem, given a set of flights, is to properly assign personnel to these flights based on Air Force flight regulations [2], the crew that best meets the squadron's current objectives, and the crew that best executes the specific flight's tasks. To further complicate the problem, we must include the military specific constraints that we have previously mentioned.

We introduce the baseline formulation in Section 4.1, which is the minimum set of constraints needed to have a legally flyable crew. We present an extension of the baseline formulation in Section 4.2, where it aims to consider additional constraints a scheduler accounts for to create the best possible schedule and to ensure personnel are meeting necessary flight training requirements.

4.1 Baseline Model Formulation

This section presents the mathematical formulation of the baseline problem, formulated as a integer optimization problem. The baseline problem is also presented in [6].

4.1.1 Indices and Sets

$i \in I$	Pilots
$f \in F$	Flights
$p \in P$	Periods
$j \in J$	Qualifications

4.1.2 Subsets

$I^{EP} \subset I$	The set of pilots qualified to be EPs
$I^{IP} \subset I$	The set of pilots qualified to be IPs
$U_f \subset F$	The set of flights that overlap with flight f for $f \in F$
$J_f \subset J$	The set of required qualifications for flight $f \in F$
$P_f \subset P$	The set of periods in flight $f \in F$

4.1.3 Parameters

I_f^{min}	Minimum number of pilots required for flight $f \in F$
I_f^{max}	Maximum number of pilots allowed for flight $f \in F$
a_{ip}	1 if pilot $i \in I$ is available during period $p \in P$, zero otherwise
z_{ij}	1 if pilot $i \in I$ satisfies qualification $j \in J$, zero otherwise
q_{fj}	Number of pilots with qualification $j \in J_f$ or higher needed on flight $f \in F$
W_i	Negative weight $\in [-1, 0)$ attached to deploying pilot $i \in I$ on a flight

4.1.4 Decision Variables

X_{if} 1 if pilot $i \in I$ is assigned to flight $f \in F$

4.1.5 Mathematical Formulation

$$\min \sum_{i \in I} \sum_{f \in F} X_{if} \tag{4.1}$$

subject to:

$$\sum_{i \in I} X_{if} \geq I_f^{min} \qquad \forall f \in F \tag{4.2}$$

$$\sum_{i \in I} X_{if} \leq I_f^{max} \qquad \forall f \in F \tag{4.3}$$

$$\sum_{i \in I} z_{ij} X_{if} \geq q_{fj} \qquad \forall f \in F, j \in J_f \tag{4.4}$$

$$X_{i\overline{f}} + X_{if} \leq 1 \qquad \forall i \in I, f \in F, \overline{f} \in U_f \tag{4.5}$$

$$X_{if} \leq a_{ip} \qquad \forall i \in I, f \in F, p \in P_f \tag{4.6}$$

$$X_{if} \in \{0, 1\} \qquad \forall i \in I, f \in F \tag{4.7}$$

The first equation, the objective function, minimizes the number of pilots scheduled to fly in the desired time window. Additionally, (4.1) can be used for feasibility as sometimes in last minute situations, a scheduler might solely be in search for a legally flyable crew. We name each objective function so we can easily refer to them in future sections and we name (4.1) "Min # Assignments."

Equations (4.2) and (4.3) ensure that the number of required pilots on flight f are filled; all flights have a minimum required number of pilots, but some also have a maximum number of pilots allowed to be on a flight.

Equation (4.4) ensures that each flight has enough pilots with the appropriate qualifications. Consider a flight where you need two pilots; one pilot with a minimum qualification level of A and the other pilot with a minimum qualification level of B, where A is a higher qualification than B. Then, $q_{fj} = 1$ for qualification $j = A$ and $q_{fj} = 2$ for $j = B$. The way this constraint is formulated, if $q_{fj} = 1$ instead of $q_{fj} = 2$ for $j = B$, it could be the case that the second pilot the optimization model assigned to that flight would not be qualified enough to fill the seat for $j = B$.

Equation (4.5) ensures each pilot can only be in one place at a time; the set U_f also takes into account unavailability due to crew rest requirements (12 hours) after each flight and the assumption pilots can only fly one event per day. Equation (4.6) ensures that a pilot cannot fly if they are unavailable and Equation (4.7) establishes the domain of the decision variables to be binary.

4.1.6 More Objective Functions

$$\min \sum_{i \in I^{IP}} \sum_{f \in F} X_{if} + \sum_{i \in I^{EP}} \sum_{f \in F} X_{if} \tag{4.8}$$

$$\max \sum_{i \in I} \sum_{f \in F} X_{if} W_i \tag{4.9}$$

Equation (4.8) or "Max EP/IP", minimizes the number of IPs and EPs assigned to flights, in order to account for the scheduler's priorities discussed in Section 3.4 of maximizing the availability of the highly qualified pilots.

Equation (4.9) or "Penalize Overqual", is similar to the previous equation, but attempts to assign the lowest qualified pilot feasible for each pilot seat. Lower ranked pilots will have less of a negative impact on the objective value, and will thus be used before higher qualified pilots who have a higher negative penalty. This allows the more qualified pilots to be available for important missions and administer training flights. W_i is used in the objective function to assign the lowest qualified, feasible pilot to each slot. More qualified pilots have a more negative W_i, so in a maximization problem, they will be used sparingly. We normalize these weights to be in the interval $[-1, 0)$, so to get W_i values in the appropriate range, we divide the integer qualification levels by the number of unique qualification levels in the squadron.

4.2 Training Requirements Model Formulation

Now we expand upon the baseline formulation to incorporate training requirements, so that pilots stay current and legally qualified. Additionally, this will prevent pilots from flying the same type of flight repeatedly.

All constraints and information from Section 4.1 are also included in this model, but for the sake of simplification we choose to not rewrite them. Instead we present additional constraints that are to be added to the baseline model, and offer substitute objective functions for the scheduler to consider. Like in the baseline model, this model's mathematical formulation can also be formulated as a integer optimization problem and is described in the following subsections.

4.2.1 Additional Indices and Sets

$s \in S$ All types of flight training requirements

4.2.2 Additional Sub Sets

$F_s \subset F$ Flights that are of type requirement $s \in S$

$S_i \subset S$ Set of requirements that pilot $i \in I$ needs to satisfy

4.2.3 Additional Parameters

R_{is} Number of requirements of type $s \in S$ needed for pilot $i \in I$

t_{is} Time until requirements of type $s \in S$ are due for pilot $i \in I$

4.2.4 Additional Variables

r_{kis} 1 if $k \in 1 : R_{is}$ of requirement $s \in S_i$ is satisfied by pilot $i \in I$

4.2.5 Mathematical Formulation

$$\max \sum_{i \in I} \sum_{s \in S_i} \sum_{k \in 1:R_{is}} r_{kis} \tag{4.10}$$

subject to:

$$\sum_{k=1:R_{is}} r_{kis} \leq \sum_{f \in F_s} X_{if} \qquad \forall i \in I, s \in S_i \tag{4.11}$$

$$r_{kis} \leq r_{k-1,i,s} \qquad \forall i \in I, s \in S_i, k \in 2 : R_{is} \tag{4.12}$$

$$r_{kis} \in \{0, 1\} \qquad \forall i \in I, s \in S_i, k \in R_{is} \tag{4.13}$$

Equation (4.10) or "Training Equal", the objective function, maximizes the number of training requirements satisfied, weighting all requirements equally.

Equation (4.11) says that the number of requirements *completed* of type s by pilot i is less than or equal to the number of flights flown of type s by pilot i. Equation (4.12) ensures that the binary variable representing one requirement satisfied of type s by pilot i is 1 before the binary variable representing two requirements satisfied is 1, and so on. Finally, Equation (4.13) establishes the domain of these additional auxiliary variables to be binary.

4.2.6 More Training Requirements Based Objective Functions

$$\max \sum_{i \in I} \sum_{s \in S_i} \sum_{k \in 1:R_{is}} r_{kis}(R_{is} - k + 1) \tag{4.14}$$

$$\max \sum_{i \in I} \sum_{s \in S_i} \sum_{k \in 1:R_{is}} r_{kis}(R_{is} - k + 1)(\frac{1}{t_{is}}) \tag{4.15}$$

$$\max \sum_{i \in I} \sum_{s \in S_i} \sum_{k \in 1:R_{is}} r_{kis}(R_{is} - k + 1) + \sum_{i \in I} \sum_{f \in F} X_{if} W_i \tag{4.16}$$

In contrast to Equation (4.10), Equation (4.14) or "Training Linear", weights the training requirements linearly by rewarding more if there are more outstanding requirements of that type of flight for a pilot. For instance, if for a pilot there are 3 of Type 1 requirements and 2 were flown in a period, the 1st requirement flown would have a reward of 3 and the 2nd would have a reward of 2.

Equation (4.15) or "Training Time", has the same structure as (4.14), but now we account for time such that the pilots who have more urgent requirements to finish before they become non-current are rewarded more heavily.

Equation (4.16), referred to as "Training Penalize Overqual", is a multi-objective function that accounts for both training requirements and overqualification. The first part of the objective function rewards the training requirements linearly like in Equation (4.14), whereas the second part of the objective function, similar to (4.9), assigns the lowest qualified pilot feasible for each slot. Lower ranked pilots will have less of a negative impact on the objective value, and will thus be used before higher qualified pilots who have a higher negative penalty. We note that the training requirements part of the objective function is more heavily prioritized than the overqualification part of the objective function.

Furthermore, any combination of any of the objective functions proposed in Sections 4.1 and 4.2 are valid substitutions to the already described equations, but more research and exploration needs to be done on multi-objective optimization and we briefly explain this further on in the thesis.

Chapter 5

Computational Experiments and Findings

To test our methodologies and formulations, we applied them to a provided dataset from a single C-17 squadron. The dataset consisted of 87 pilots (9 EPs, 37 IPs, 8 MPs, and 33 FPs) and about six months worth of flights (486 simulator flights and 314 local/mission flights). The dataset also had 245 records of unavailability, where certain pilots were linked to multiple unavailability periods. Throughout this chapter we present the effectiveness of our formulations along with evaluations of the different objective functions on this data.

The rest of the chapter is organized as follows. In Section 5.1, we further describe the dataset and the relevant features important to obtain our results. We describe the computational tools used and the efficiency of our optimization algorithms in Section 5.2. We propose important metrics discussed by users and subject matter experts to evaluate schedules in Section 5.3. In Section 5.4, we display the results of the objective functions for the proposed metrics, as well as visualizations to help schedulers analyze candidate schedules. Lastly, in Section 5.5, we conclude this chapter summarizing the results and offering a few recommendations.

5.1 Dataset Description

We further detail all the data within the provided dataset in this section. First, the pilot data provides their name or pilot ID and their qualification for which types of flights they are able to fly. Next, is the pilot unavailability data file, which gives the start date, end date, and reason for being unavailable. For the simulator flights data, each entry consists of the specific type of simulator and the date the simulator flight was accomplished. We note that in this dataset, there are nine unique types of simulator flights. The locals and missions were combined in one file, where one entry in the file consists of the type of flight, the crew type, the premission brief date, the departure date, the return date, the postmission debrief date, whether the flight was at night, and whether the flight had AR or not. The main thing to note from this file is that if a pilot is assigned to a flight, they are completely assigned to that flight from the premission brief date to the postmission debrief date, and not only from the departure date to the return date. We were also provided the training requirements necessary for each qualification level that we extrapolated out to each pilot, where certain qualifications require more training requirements of a specific type of flight than others. In our modeling we consider 9 simulator training requirements and 5 locals/missions training requirements, for a total of 14 different flight training requirements. Finally, we note that the data provided did not actually contain which pilots were assigned to which flights.

Now to preprocess all the data to be used in the model, we first set the pilots to be unavailable for all their record dates in the unavailability dataset. Next, based on the event type (simulators and locals/missions), we are able to determine the quantity of pilots needed for each flight based on Air Force documentation and details provided by schedulers. Once we know the event type, we further determine the necessary crew qualifications makeup from the type of flight and crew type information in the data, which is again based on Air Force documentation and provided details. We are also able to determine which flights overlap with each other based on the dates provided in the data, so that our algorithms do not schedule pilots on overlapping flights. Lastly, we determine the necessary training requirements for each pilot based on the data provided such that our algorithms can track and reward pilots for completing their necessary training.

5.2 Computational Efficiency

To carry out our experiments, we used the programming language Julia, the modeling language JuMP for optimization, and the commercial solver Gurobi to solve our integer optimization problems. Typically with combinatorial and integer optimization problems, heuristics are needed to ensure the optimization problems are solved to optimality. However, because we were only solving the problem every two weeks as recommended by the schedulers and leveraging the functionality of Gurobi, we were able to solve each two week problem instance to optimality in only a matter of seconds. If a commercial solver like Gurobi or CPLEX is unavailable, then further heuristic methods might need to be explored to solve the problem instances efficiently and to ensure the problems are solved to optimality.

5.3 Metrics Description

Recall the two main priorities in a schedule for the schedulers after feasibility are keeping their highly qualified pilots available, or preventing overqualification, and maximizing the number of training requirements completed. We propose methods to quantify each of these priorities into metrics to display to the scheduler in order for them to decide what is best for the squadron at that time. Furthermore, talking to pilots in different flying squadrons, they explain an additional metric they care about in their own schedule is the total number of locals and missions flown. We directly optimize the metrics for overqualification and training requirements completed in our formulations, but the other metric is used for schedule evaluation after the fact and to confirm the objective functions work as intended. We analyze the total number of locals and missions flown per pilot for the different qualification groups, where we would expect the more experienced pilots to fly fewer flights when we optimize for overqualification.

5.3.1 Overqualification

To maximize the availability of highly trained pilots, we introduce the notion of overqualification. We define overqualification to be the act of assigning a more qualified pilot then necessary to fill an individual pilot seat on a particular flight. The idea of preventing overqualification is exactly why Equations (4.8) and (4.9) were introduced in Section 4.2.

We further define how to quantify the notion of overqualification by introducing an overqualification score (OQS). The OQS for one single pilot assignment is the number of qualification levels greater than the minimum qualification necessary for a specific pilot seat. For example, if the minimum qualification level needed for a particular seat was 1, but a pilot with qualification 5 was assigned to that seat, then the OQS for that seat would be: $(5-1) = 4$. We then take the cumulative OQS for every single pilot assignment to get the total OQS. Additionally, we also present the average OQS, by dividing the total OQS by the total number of pilot assignments. We also note that the W_i's in the objective functions are set in such a way to minimize OQS. We set the lowest qualified pilots to be the smallest integer value and the highest qualified pilots to be the number of unique pilot qualifications so the higher qualified pilots are penalized more than the inexperienced pilots in a maximization problem.

We refine this previous metric by introducing an advanced OQS (AOQS). The calculation is the same for each pilot assignment as OQS, but instead we use a subset of qualifications for the AD and SOLL II events. Only the special qualified pilots are able to fly AD and SOLL II events, so it would be unnecessary to consider the entire array of qualifications for these types of flight events in our scoring. Again, aggregating across all pilot assignments gives us the total AOQS, and dividing the total AOQS by the total number of pilot assignments gives us the average AOQS.

5.3.2 Training Requirements

To compare the number of training requirements accomplished by the different objective functions in Section 4.2, we count the total number of training requirements completed for each pilot. For each training requirement a pilot accomplishes, they receive a single point, however it is possible to accomplish more than one training requirement on a particular flight. We then aggregate the total number of training requirements accomplished for each pilot across all flights and time periods to get the total number of training requirements accomplished.

5.3.3 Total Locals and Missions Flown

Work-home balance is a crucial aspect of a pilot's life. More qualified and older pilots, who have families and training to evaluate, tend to like to stay close to home. By contrast, younger and more inexperienced pilots usually prefer to jump on missions and locals to gain more real flying experience while being able to fly to new places around the world. Additionally, locals and missions tend to be more rigorous and time consuming due to all the preparation needed before and after each flight. These are all reasons pilots or users are interested in the total number of locals and missions each one of them are flying.

5.4 Metrics Results

To gather results for each objective function, we solve the problem for the first two weeks, save the optimal flight-pilot assignments, take that into account for the next two weeks, solve the problem for that next two weeks, and continue that process until we assigned pilots to all of the flights in the dataset. We then aggregate all of this information across all thirteen two week time periods to gather the six months worth of assignments for each objective function to compare metrics.

5.4.1 Overqualification

We present the different scores for the three baseline objective functions in Table 5.1. We see that Equation (4.1) has the highest scores, which is expected because this objective function does not capture overqualification. Now both Equations (4.8) and (4.9) were used to directly minimize the OQS. However, we see that Equation (4.9) had the lower score because it attempted to assign the lowest qualified pilot for each pilot seat, whereas Equation (4.8) only focused on preventing overqualification for the highest two qualification groups.

Equation	Total OQS	Ave. OQS	Total AOQS	Ave. AOQS
Min # Assignments (4.1)	18,652	9.869	17,186	9.093
Max EP/IP (4.8)	7,840	4.148	6,691	3.540
Penalize Overqual (4.9)	4,188	2.216	3,288	1.740

Table 5.1: Total and average OQS and AOQS for each objective function.

5.4.2 Training Requirements

Before we discuss training requirements results, we illustrate the difference between Equations (4.10) and (4.14) by referencing Figures 5.1 and 5.2. Figure 5.1 corresponds to the objective function where each training requirement is weighted equally and Figure 5.2 corresponds to the objective that weights the training requirements linearly based on the explanation in Section 4.2.5. Each color represents the number of outstanding training requirements for a particular flight event and the height of each line represents the frequency of outstanding training requirements across all pilots and all flight events. For example, if one pilot has four training requirements left for all fourteen training requirement categories, then the frequency for that particular pilot would be fourteen for four outstanding requirements and zero for the others. The x-axis is the time period, the y-axis is the frequency of outstanding training requirements, and the stacked bars represent the number of outstanding training requirements at the beginning of each time period.

The key takeaway is that the linear objective function heavily favors completing the two, three, and four outstanding requirements before completing the one outstanding requirement. This is most clearly seen in how the blue bars are larger in Figure 5.2 compared with Figure 5.1, especially between periods 3 through 10; as pilots complete events with at least two outstanding requirements, there are more pilot-event combinations left with just one outstanding requirement. Furthermore, we see that the linear based objective function in Figure 5.2 has almost eliminated all three or four outstanding requirements after time period 8, whereas the same is not accomplished until about time period 11 for the equally weighted objective function in Figure 5.1. Thus, the linear based function prioritizes the pilot-event combinations that have the most outstanding requirements initially, which ultimately will make it easier for pilots to finish all their training requirements in time.

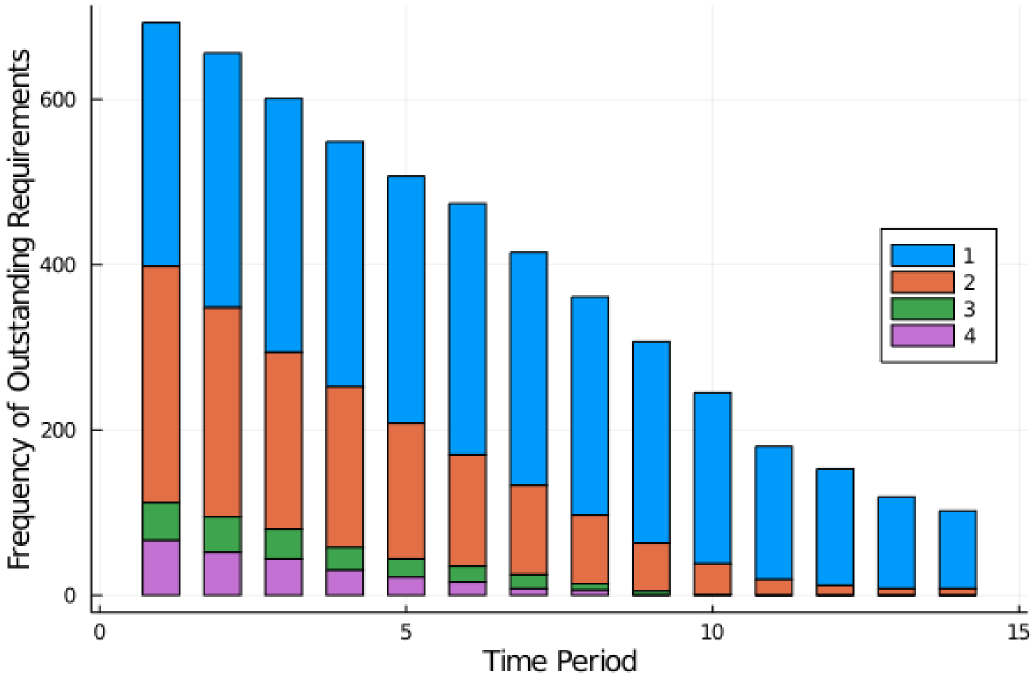

Figure 5.1: Outstanding Training Requirements for Equation (4.10).

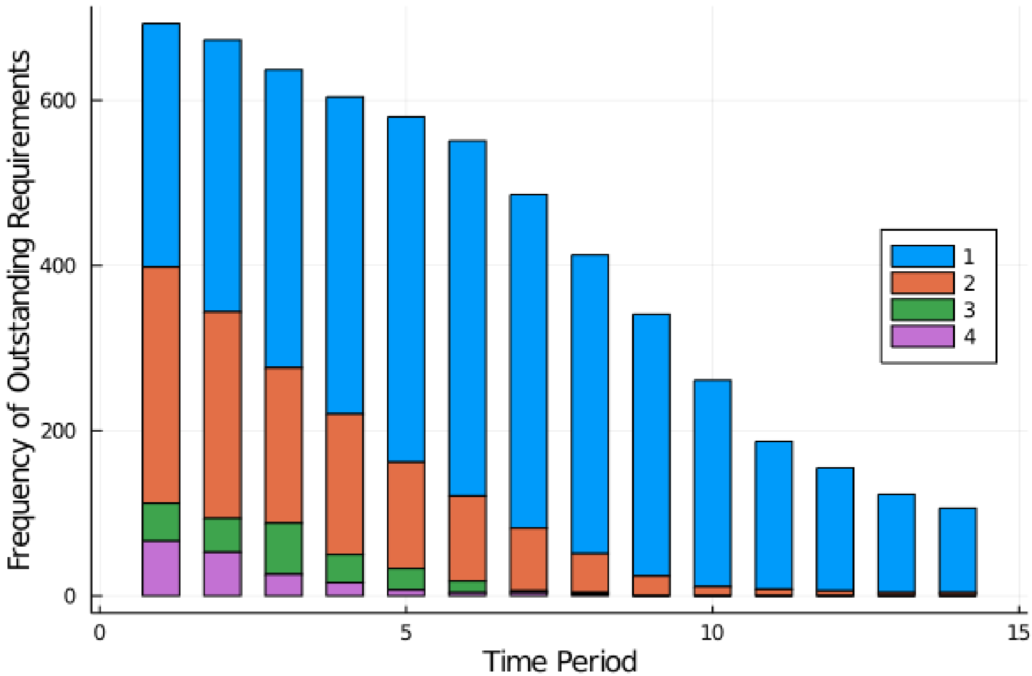

Figure 5.2: Outstanding Training Requirements for Equation (4.14).

Now we display the total number of training requirements accomplished for all the objective functions in Section 4.2, which can be found in Table 5.2. We see from the results that the total number of requirements completed is comparable across the training requirements objective functions, but the objective function that incorporates time (Equation (4.15)) completes a few less requirements. This is due to the fact that training requirements are not prioritized until closer to their due date. An objective function like this could be valuable for the few types of training requirements that renew on a periodic or rolling horizon basis, which is best to fulfill the training towards the end of the time horizon to avoid flying that type of flight so often.

Equation	**Training Requirements**
Training Equal (4.10)	1,288
Training Linear (4.14)	1,288
Training Time (4.15)	1,282
Training Penalize Overqual (4.16)	1,290

Table 5.2: Total Training Requirements completed by objective function.

We further illustrate the difference between the training requirements objective functions by looking at a time series plot in Figure 5.3. The x-axis is each one of the thirteen time periods and the y-axis is the number of training requirements completed for each time period. Each line closely follows one another, but we see the most obvious differences between the objective functions are in time periods 6, 7, and 10-12, which ultimately lead to the differences seen in Table 5.2.

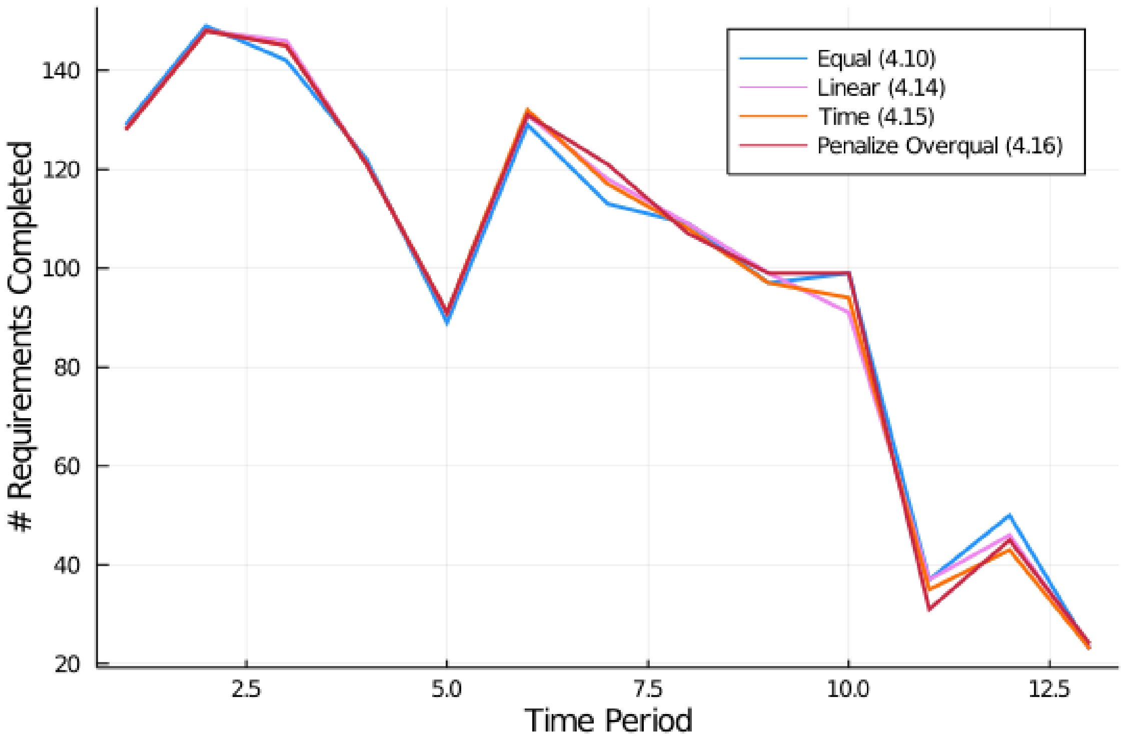

Figure 5.3: Training Requirements completed over time by objective function.

5.4.3 Overqualification vs. Training Requirements

It is imperative to compare the trade-offs between the OQS and the total number of training requirements completed to see how favoring one of these priorities impacts the other. Table 5.3 displays both of the metrics for every proposed objective function proposed in Chapter 4.

Equation	Average OQS	Training Requirements
Min # Assignments (4.1)	9.869	808
Max EP/IP (4.8)	4.148	657
Penalize Overqual(4.9)	2.216	504
Training Equal (4.10)	9.277	1,288
Training Linear (4.14)	9.279	1,288
Training Time (4.15)	9.279	1,282
Training Penalize Overqual (4.16)	5.873	1,290

Table 5.3: OQS vs. Total Training Requirements completed for each objective function.

To further investigate this trade-off, we introduce the parameter $B \in [0, \infty)$ to Equation (4.16), such that we have the following objective function:

$$\max \sum_{i \in I} \sum_{s \in S_i} \sum_{k \in 1:R_{is}} r_{kis}(R_{is} - k + 1) + B * \sum_{i \in I} \sum_{f \in F} X_{if} W_i. \tag{5.1}$$

We iterate over several values of B and obtain results for the OQS and the total number of training requirements completed for each B. We visualize the findings of the trade-offs in Figure 5.4 for both the proposed objective functions in Chapter 4 and the results for varying B. The x-axis is the average OQS for each iteration, the y-axis is the total number of training requirement accomplished for that iteration, and the solid black line represents the maximum number of training requirements possible for the entire six month time window.

We observe in Figure 5.4 that a Pareto frontier appears between Equations (4.9) and (4.16). Additionally, we note that as B tends to infinity, we approach Equation (4.9), when $B = 1$ we get exactly Equation (4.16), and when $B = 0$ we get exactly Equation (4.14). Lastly, we note that Equations (4.1), (4.8), (4.10), (4.14), and (4.15) are all Pareto dominated.

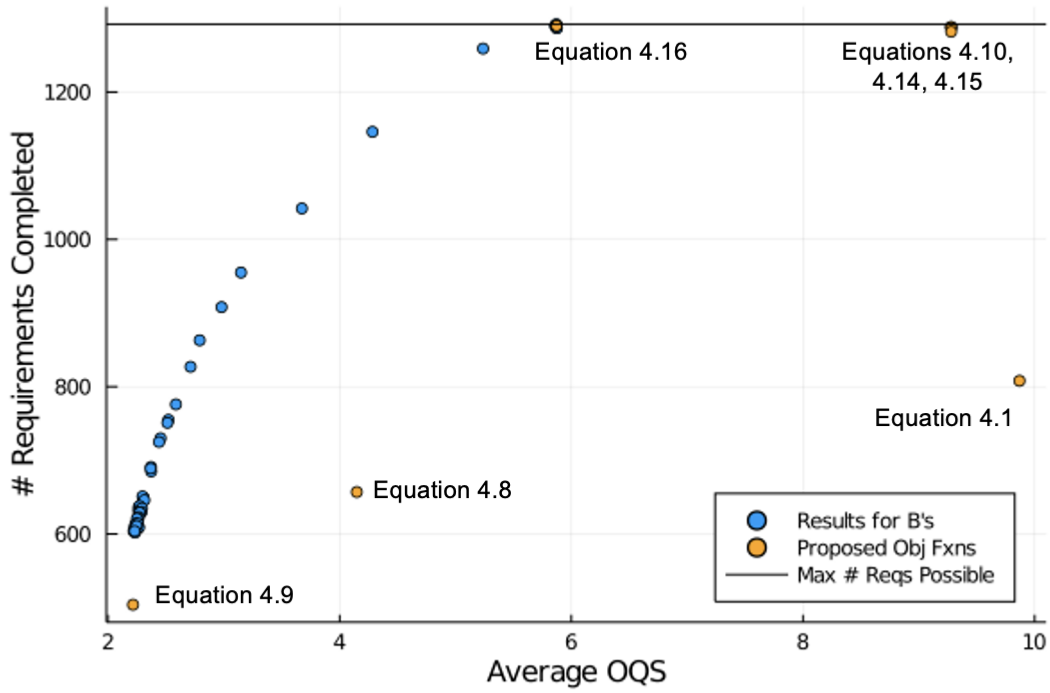

Figure 5.4: OQS vs Training Requirements trade-off.

5.4.4 Total Locals and Missions Flown

To display the total number of locals and missions each pilot is to fly under a specific objective function we refer to the histograms in Figures 5.5 and 5.6. We choose to display histograms for only two of the objective functions proposed: the first one is for Equation (4.1) and the second graph is for Equation (4.9). The y-axis is the total number of locals and missions flown and the x-axis represents each pilot ordered in increasing order of total flights flown. The color of each bar represents the qualification group, where the lighter colors are the less qualified pilots and the the darker pilots are the more qualified pilots.

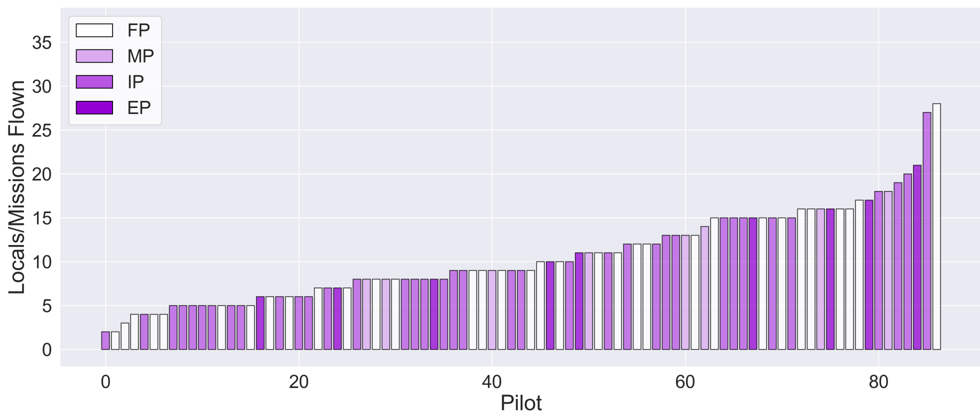

Figure 5.5: Histogram of total locals and missions flown by pilot for Equation (4.1).

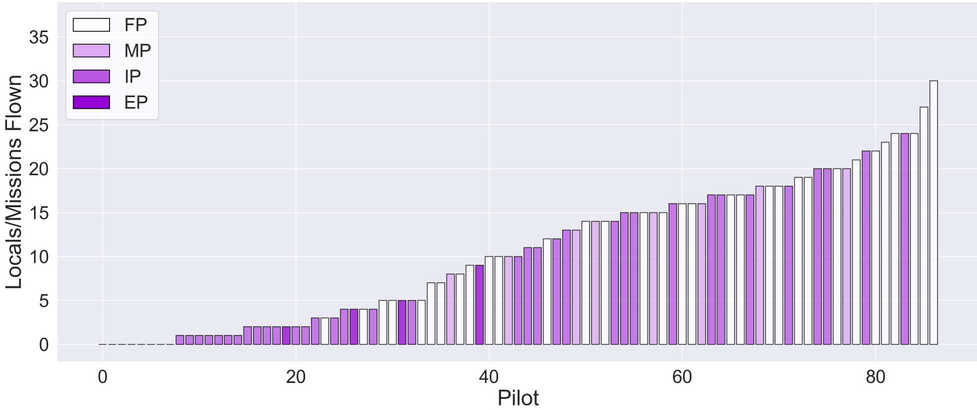

Figure 5.6: Histogram of total locals and missions flown by pilot for Equation (4.9).

To quantify these histograms, we present the summary statistics for each objective function we have proposed thus far in Table 5.4. We display the minimum and maximum number of locals and missions flown by any one pilot, the average number of locals and missions flown by each pilot, and then the standard deviation (σ) of the number of locals and missions flown for each pilot. We note the average number of locals and missions flown for each pilot are all identical because the total number of flight assignments are all the same.

Equation	**Min**	**Max**	**Average**	$\boldsymbol{\sigma}$
Min # Assignments (4.1)	2	28	10.552	5.203
Max EP/IP (4.8)	1	32	10.552	6.405
Penalize Overqual (4.9)	0	30	10.552	7.959
Training Equal (4.10)	1	26	10.552	4.266
Training Linear (4.14)	2	27	10.552	4.341
Training Time (4.15)	3	24	10.552	4.157
Training Penalize Overqual (4.16)	1	37	10.552	7.135

Table 5.4: Total locals and missions flown per pilot for each objective function.

In addition to the histograms and summary statistics, we are also able to display the distribution of total locals and missions flown for each pilot qualification group. Again, we only display the distributions for the objective functions in Equations (4.1) and (4.9).

In the objective function used to create the results for Figure 5.8 we are minimizing overqualification, whereas the objective function for Figure 5.7 is more randomly assigning pilots to flights. This is evident by noting the difference between the two figures by focusing on the EP qualification group because the average number of locals and missions flown for the second figure is much lower than that of the first figure. This helps us confirm that objective function (4.9) is working as intended by placing the more qualified pilots on less locals and missions to keep them more available. These visualizations, as well as the others above, can help schedulers and users to evaluate schedules.

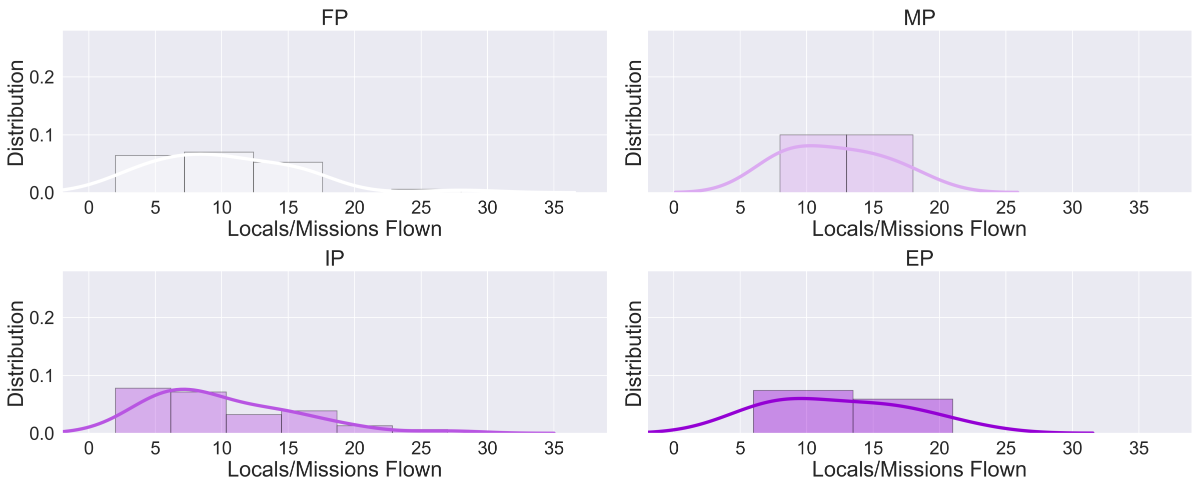

Figure 5.7: Distribution of total locals and missions flown by qualification group for Equation (4.1).

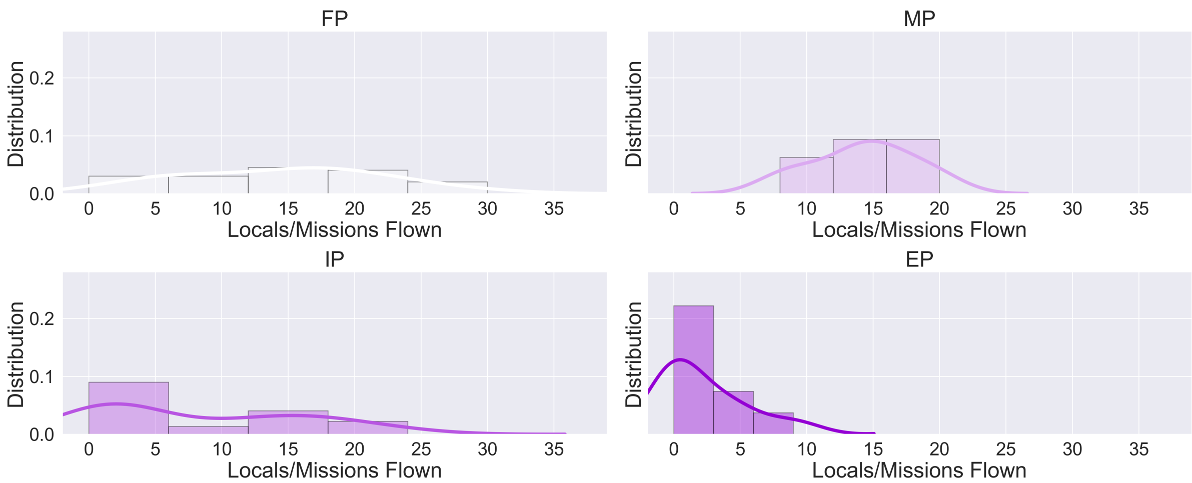

Figure 5.8: Distribution of total locals and missions flown by qualification group for Equation (4.9).

5.4.5 Example Schedules

In addition to the visualizations above, we are also able to display the schedule in a personnel-centric manner. Due to the complex nature of this visualization, we choose to only present a small example of a personnel-centric schedule in Figure 5.9. The y-axis represents each pilot ordered by increasing qualification and the x-axis is the progression of time. Each row then contains all flight assignments or unavailability for a given pilot, where F represents a local or mission and S represents a simulator flight.

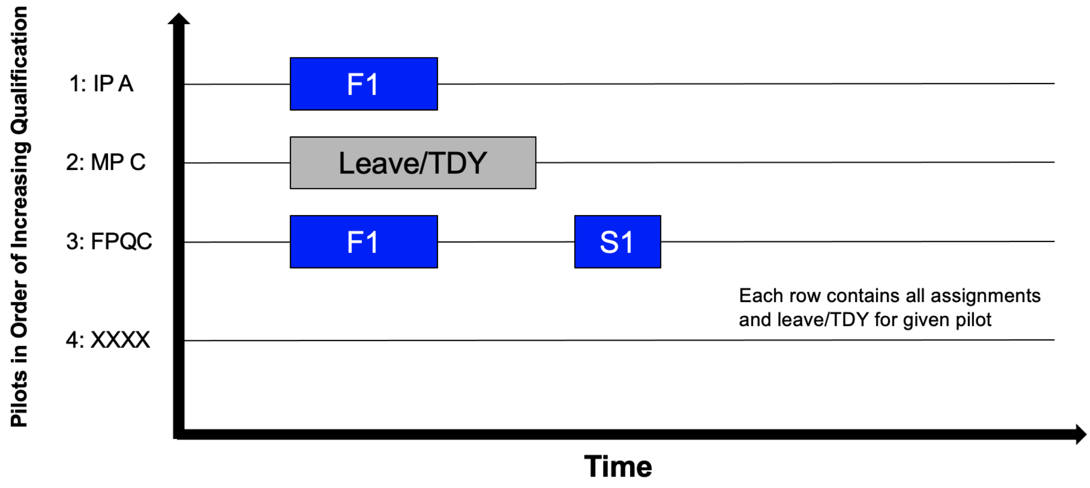

Figure 5.9: Small example of a personnel-centric gantt chart.

5.5 Summary and Recommendations

We presented many results in this chapter and now provide a brief summary. Equation (4.9) provides the lowest OQS, with Equation (4.8) being a valid alternative when optimizing for overqualifications. All the proposed training requirements objective functions produce a similar number of training requirements completed, although Equation (4.16) has the lowest OQS of the group. In the end, each one of the proposed objective functions have their pros and cons, but nevertheless they provide suitable options for a scheduler to choose from depending on their squadron's current situation.

We offer the following recommendations and situations for when the scheduler should possibly choose one objective function over another:

1. If the sole priority is to minimize overqualification, then undoubtedly Equation (4.9) is the best fit. However, if it is only important to minimize overqualification for the EP and IP qualification groups, then Equation (4.8) might be a better choice.

2. If the priority is to maximize training requirements completed, then any of the requirements based objective functions would be a valid choice. In spite of that, Equation (4.16) not only maximizes the number of training requirements completed, but offers the lowest OQS of the requirements based objective functions. Therefore, we would recommend using Equation (4.16) in most circumstances, but as previously mentioned, more research is necessary into the prioritizations for the multi-objective optimization domain.

Chapter 6

Conclusions

Using optimization-based algorithms for the military scheduling problem continues to be an active area of research, and the recent initiative for AI and automation within the Department of Defense keeps us optimistic about the path forward. In this thesis, we have presented mathematical formulations that encapsulate multiple objective functions, and introduced metrics and visualizations that could be used by schedulers. Our methods ultimately provide feasible solutions significantly faster than the manually created schedules.

When we first started developing our research, Puckboard was only being used in one or two C-17 squadrons across the Air Force, but now it is being used in over 120 (and growing) flying squadrons in the Air Force. Needless to say, as we on-board our algorithms into Puckboard's application, we have the ability to make a huge impact in the lives of Airmen, who will then be able to focus primarily on their mission on a daily basis.

6.1 Future Work

Our thesis presents optimization-based algorithms to solve the Air Force crew scheduling problem, but nevertheless, there is still scope for extensions and further research. We detail some ideas for future work.

- As we gather user feedback and begin the process of on-boarding our algorithms to Puckboard, continued refinement and development of the core algorithms will be crucial in the utility of our formulations moving forward.

- In optimization problems, fairness and equity can sometimes be difficult to achieve. An interesting idea would be to propose further objective functions to help optimize for these ideas. Possible examples would be to prevent a pilot from flying only simulator events, or avoid assigning the same pilot to all of the less desirable missions.

- As scheduler and pilot priorities continue to emerge, there could be an investigation of having user inputs for objective function weightings in multi-objective settings. Correctly implementing this will involve the normalization of each part of the objective function [11]. However, allowing a very large degree of customization by users can sometimes lead to unforeseen challenges and consequences.

- Finally, disruptions and uncertainties in a schedule are more prevalent than ever in a military setting. For example, the weather could delay flights, aircraft maintenance issues could reduce resources, mission times could be prolonged in enemy territories, unexpected missions could arise, pilots could become unavailable, and so on. Therefore, it is necessary to develop algorithms and formulations to account for these uncertainties. One way to do this is for multi-stage optimization, where our formulations are the first stage, and the optimization formulations developed in Chin [6] are used as the second stage. Another option is to develop Robust Optimization (RO) formulations to account for uncertainty and guard against the worst case scenarios. In particular, recent research on RO applied to airline flight and crew scheduling [4, 24, 25] are relevant to our problem.

Bibliography

[1] Air Force Instruction (AFI) 11-202V3. *General Flight Rules*, 2016.

[2] Air Force Manual (AFMAN) 11-2C-17V3. *C-17 Operations Procedures*, 2019.

[3] J. F. Bard, D. P. Morton, and Y. M. Wang. Workforce planning at usps mail processing and distribution centers using stochastic optimization. *Annals of Operations Research*, 155(1):51–78, 2007.

[4] D. Bertsimas, D. B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM review*, 53(3):464–501, 2011.

[5] A. Caprara, P. Toth, D. Vigo, and M. Fischetti. Modeling and solving the crew rostering problem. *Operations research*, 46(6):820–830, 1998.

[6] C. Chin. Disruptions and robustness in air force crew scheduling. Master's thesis, Massachusetts Institute of Technology, 2021.

[7] Department of Defense. *Summary of the 2018 Department of Defense Artificial Intelligence Strategy*, 2018.

[8] M. Desrochers and F. Soumis. A column generation approach to the urban transit crew scheduling problem. *Transportation science*, 23(1):1–13, 1989.

[9] A. T. Ernst, H. Jiang, M. Krishnamoorthy, H. Nott, and D. Sier. An integrated optimization model for train crew management. *Annals of Operations Research*, 108(1):211–224, 2001.

[10] D. L. Goldfein and M. P. Donovan. *The United States Air Force Artificial Intelligence Annex to The Department of Defense Artificial Intelligence Strategy*, 2019.

[11] O. Grodzevich and O. Romanko. Normalization and other topics in multi-objective optimization. 2006.

[12] B. T. Hall, H. T. Hargrove, and J. M. Willis. *Analysis of Current Flight Scheduling Practices and Recommendations to Efficiently Reduce Deviations from Syllabus Time-To-Train*. PhD thesis, Naval Postgraduate School, 2011.

[13] P. R. Harper, N. Powell, and J. E. Williams. Modelling the size and skill-mix of hospital nursing teams. *Journal of the Operational Research Society*, 61(5):768–779, 2010.

[14] J. Herbers and J. Hromkovic. Models and algorithms for ground staff scheduling on airports. Technical report, Fakultät für Mathematik, Informatik und Naturwissenschaften, 2005.

[15] C. Honour. A computer solution to the daily flight schedule problem. Master's thesis, Naval Postgraduate School, 1975.

[16] R. S. Jacobs. Optimization of daily flight training schedules. Master's thesis, Naval Postgraduate School, 2014.

[17] N. Kohl and S. E. Karisch. Airline crew rostering: Problem types, modeling, and optimization. *Annals of Operations Research*, 127(1):223–257, 2004.

[18] P. Punnakitikashem, J. M. Rosenberber, and D. F. Buckley-Behan. A stochastic programming approach for integrated nurse staffing and assignment. *Iie Transactions*, 45(10):1059–1076, 2013.

[19] T. R. Robbins and T. P. Harrison. A stochastic programming model for scheduling call centers with global service level agreements. *European Journal of Operational Research*, 207(3):1608–1619, 2010.

[20] D. Sitompul and S. Randhawa. Nurse scheduling models: a state-of-the-art review. *Journal of the Society for Health Systems*, 2(1):62–72, 1990.

[21] R. J. Slye. Optimizing training event schedules at naval air station fallon. Master's thesis, Naval Postgraduate School, 2018.

[22] J. Van den Bergh, J. Beliën, P. De Bruecker, E. Demeulemeester, and L. De Boeck. Personnel scheduling: A literature review. *European journal of operational research*, 226(3):367–385, 2013.

[23] J. Van den Bergh, P. De Bruecker, J. Beliën, L. De Boeck, and E. Demeulemeester. A three-stage approach for aircraft line maintenance personnel rostering using mip, discrete event simulation and dea. *Expert Systems with Applications*, 40(7):2659–2668, 2013.

[24] R. Vujanic, P. Goulart, and M. Morari. Robust optimization of schedules affected by uncertain events. *Journal of Optimization Theory and Applications*, 171(3):1033–1054, 2016.

[25] C. Yan. *Airline scheduling and air traffic control: incorporating uncertainty and passenger and airline preferences.* PhD thesis, Massachusetts Institute of Technology, 2017.

[26] Z. Zhang, M. Zhou, and J. Wang. Construction-based optimization approaches to airline crew rostering problem. *IEEE Transactions on Automation Science and Engineering*, 17(3):1399–1409, 2019.

[27] S.-Z. Zhou, Z.-H. Zhan, Z.-G. Chen, S. Kwong, and J. Zhang. A multi-objective ant colony system algorithm for airline crew rostering problem with fairness and satisfaction. *IEEE Transactions on Intelligent Transportation Systems*, 2020.