Proceedings of the 2014 Industrial and Systems Engineering Research Conference
Y. Guan and H. Liao, eds.

ACCAM Global Optimization Model for the USCG Aviation Air Stations

Abstract ID: I1023

Christie Nelson, Endre Boros, Fred Roberts, Javier Rubio-Herrero, Paul Kantor, Curtis McGinity, Brian Nakamura, Brian Ricks
CCICADA, Rutgers University
Piscataway, New Jersey 08854

Patrick Ball, Chad Conrad, Kevin Hanson, Thomas Rader
United States Coast Guard
Washington DC 20593

Abstract

A model was created for the United States Coast Guard (USCG) to maximize aircraft fleet operational performance subject to budgetary constraints, or, conversely, to minimize aircraft fleet operational costs subject to performance targets. This is a two-stage model: The first stage, prior work, is a simulation model of each USCG Air Station generating performance metrics resulting from various scenarios. Each scenario is determined by a large number of relevant Air Station attributes, including number and types of stationed aircraft, operational level, historical Search and Rescue mission response, historical maintenance processes, deployment requirements, and other mission requirements. The second stage, the subject of this paper, is an optimization model over the structured set of scenarios to determine the optimal deployment assignments, operational levels and aircraft allocation among all USCG Air Stations, under the current infrastructure. Additionally, the optimization model is used to demonstrate the potential efficiencies of proposed infrastructural changes, namely the introduction of a deployment center.

Keywords
Resource allocation, aviation, COAST, optimization, USCG

1. Introduction

United States Coast Guard (USCG) aircraft fleet allocations to USCG aviation Air Stations have long been dictated by assumptions of operational response, aircraft availability, Search and Rescue (SAR response), and planned missions which may or may not still be valid. With the increasing demand that aviation aircraft support forward-deployed surface forces, the USCG would like the means to identify the optimal assignment of aircraft at USCG Air Stations. This requires a modeling capability (previously lacking) to comprehensively analyze response and mission demands at the USCG Air Station level. This model could provide USCG decision makers with alternatives for aircraft allocations, which would be optimized to operational as well as logistical capabilities and predicted mission needs [1].

An optimization model was created for the USCG for allocation of USCG aircraft to USCG Air Stations. The project goal was to minimize aircraft fleet operational costs subject to performance targets of various types. The model can be thought of as two-stage: The first stage, prior work called ACCAM STATION, was first a simulation of each USCG Air Station, which generates resulting performance metrics based on the specific station and what is termed a "scenario." Each scenario was determined by a large number of relevant Air Station attributes, including the USCG Air Station, the allowable aircraft at a given station, the station's operational level, the station's historical SAR mission response, the station's deployment requirements by time period, and the station's other mission requirements. The scenario was also determined by aircraft information such as the aircraft type's historical maintenance processes including both scheduled maintenance and unscheduled maintenance. The second stage of the model, the subject of this paper, was an optimization model over the set of station-specific scenarios to determine the optimal deployment assignments, operational levels, and aircraft allocation among all USCG Air Stations, under the current infrastructure. Each station had different feasible scenarios, of which one was chosen for each model run. Additionally, the optimization model can be used to demonstrate the potential efficiencies of proposed infrastructural changes, such as the introduction of a deployment center, a depot of aircraft that can be deployed to USCG Air Stations for specific missions.

This model falls under the Coastal Operation Analytical Suite of Tools (COAST) Aviation Capability and Capacity Assignment Module (ACCAM). The model, termed the ACCAM Global Optimization Model (GOM), or ACCAM GOM, was a joint effort between researchers at Rutgers University at the Command, Control, and Interoperability Center for Advanced Data Analysis (CCICADA), a US Department of Homeland Security (DHS) university center of excellence (COE), and USCG, with close and consistent communication between the partners and researchers. This project was part of an "Engage to Excel" (E2E) initiative of DHS. Under the E2E program, university COEs work very closely with DHS agencies to accelerate the development of cutting edge solutions to real operational problems by collaborating from the beginning of the problem formulation to transition of complete pieces. Previous related projects under this E2E program for CCICADA with the USCG include two Boat Allocation Models (BAM I and II), which determined assignments of small boats, or "resources," across the USCG boat stations (which are different than USCG Air Stations) so that all station requirements, or "demands," were met. In particular, ACCAM STATION and ACCAM GOM were follow-on projects building on the BAM I and II projects to optimally assign aircraft to USCG stations. COAST is a set of modules, usable and updatable individually, but rationally linked together. BAM I was one of the earliest models developed for COAST. BAM I and II will be described in more detail in Section 2.

The outline for this paper is as follows: first, the background of the project will be described in Section 2. In particular, prior work (BAM I and II) will be discussed in Section 2.1, along with related work in Section 2.2, and a description of the USCG Air Stations and missions in Section 2.3, where we also describe prior work on ACCAM STATION. Next, the model is presented in Section 3, including the project approach, describing in more detail the research question to be answered (Section 3.1), the model overview which describes how the various pieces fit together (Section 3.2), the types of data required (Section 3.3), and then the model itself (Section 3.4). Finally, Section 4 describes future work whereas conclusions are presented in Section 5.

2. Background and Prior Work

As described in Section 2.1, much prior work was completed with the USCG before starting to model the ACCAM GOM. In particular, a boat allocation model was created for the allocating of small boats (BAM I), which was then expanded to allow the "sharing" of boats between stations along with the associated time period in which the boats should be switched between boat stations in BAM II. In addition, a literature search of various resource allocation models is presented in Section 2.2. Last, a background of what USCG Air Stations are, an overview of how they function, and what some of their required missions are, including Search and Rescue as well as other missions is described in Section 2.3, where we also describe the ACCAM STATION model.

2.1 Prior Work: BAM I and BAM II

BAM I [2], [3], [4], [5] sought to create a mathematical model which could produce "good" assignments of small USCG boats, or "resources," across the USCG boat stations so that all station requirements, or "demands," were met. BAM II was a separate model which built on the ideas of BAM I, but addressed the concept of "sharing" resources or boats of a specific type among stations within a USCG Sector. The concept of "sharing" is the allocation a single boat to more than one station, with the goal of reducing program level costs, while still fulfilling mission requirements.

For the boat allocation problems, the USCG boat stations are organized into geographic districts. Each district is made up of sectors, which manage individual boat stations. Every boat station has a number of requirements which must be satisfied, such as having a minimum number of total boat hours set aside for use toward each mission, as well as enough boats to cover the station's designated capabilities. Each capability is simply a requirement asking that a minimum number of boats can be assigned, chosen from those types sharing the station's designation. These requirements may change at the start of a new time period (or fiscal year).

BAM I was modeled as a variant to the well-studied Resource Allocation Problem, an optimization problem found throughout Operations Research literature. BAM I was modeled as a mixed-integer-program (MIP) with binary, integer, and continuous variables. BAM I limited the budget as a constraint and optimized a measure of unmet task "demand," which was formalized as part of the project. BAM I did not focus on individual boats, but rather the resource was the total number of boats of a type at a given station. This helped to reduce the size of the problem, and it also gave the analyst and commanding officers flexibility when considering boats individually. The total number of hours allocated for use across all assigned boats of a type, per station, per mission type, was a second type of resource. BAM I did not consider the assignment of personnel to stations or boats. It was assumed that once boats were assigned to stations, and mission hours to boats, only then will personnel be appropriately stationed by commanding officers. The problem is then restated: to find an assignment of boats to stations and mission hours to such boats that satisfies all USCG "business rules" so that a weighted total of all stations' mission hours gone unmet is minimal. In other words, assignments are penalized if they leave stations with missions that are unable to be completed. The weights of each mission, per station, represent the importance of some over others. Business rules include covering the capability requirements at each station, scheduling hours to assigned boats within a specified range determined by maintenance requirements, allotting enough hours per boat type for adequate training of personnel, and spending no more of the budget than allowed.

Allowing fractional solutions in BAM I led to the idea of sharing boats between stations. The "sharing" of resources in the application of boats and boat stations is the concept that a boat may move between stations over the course of time (e.g. a year), shifting between various stations or else staying stationary at a given station (and not being "shared) during each time period. Even when allowing for a transfer cost to switch boats between stations, this potentially lessens the number of boats required to fulfill USCG business rules and missions and can be cost-saving. BAM I was able to show that "sharing" of boats, which are thought of as resources, could reduce the number of required USCG boats, but did not show exactly how to implement this. In addition to providing information regarding the number of boats required and the hours assigned to them to meet all known constraints of BAM I, BAM II [6] determined when and where a boat is needed by time period for a specific type of boat. This second-stage optimization problem allows the USCG to determine how many time periods they would like to have over the course of a year, and then within each time period, how many boats and boat hours are assigned to each station in that period. BAM II only allows for one boat type, as it is reasonable to assume that the USCG will not start sharing all their boats at one time, but rather, would potentially start with one boat that had potential for savings.

BAM I and BAM II were both successful projects created by our partnership with the USCG and led to the ACCAM GOM which is presented in this paper. All of these projects fall under the COAST suite.

2.2 Related Work

The model to be presented in this paper is classified as a Resource Allocation Problem. In addition to our related work presented in Section 2.1 on BAM I and II, there are many other applications. The well-studied Resource Allocation Problem,is an optimization problem found throughout Operations Research literature. Despite the fact that ours has some notable differences (particularly the idea of choosing "scenarios" for the USCG stations which will be described in more detail in Section 3), there are some similar examples found in the literature. A resource allocation model can be defined as allocating a finite resource to specific activities such that it minimizes the objective value while satisfying all the necessary conditions found in the constraints [7].

One classic example is utilizing resources to effectively fight forest fires, which involves resource planning [8]. When a fire is initially detected, it is attacked quickly in order to minimize damage. A dynamic model is used to continually re-allocate resources as the fire continues. This model differs from the ACCAM GOM as the forest fire problem deals with re-allocating resources, while ACCAM GOM allocates resources just once (until an analyst determines they would like to look at a different scenario). In addition, ACCAM GOM allocates resources to many locations (USCG Air Stations across the United States), not just one location. Another example of a resource allocation problem is planning for a public health emergency such as a bioterrorist attack [9]. This problem determines locations for dispensing aid such as vaccines or drugs, and then allocating resources to the chosen locations. The USCG ACCAM GOM has pre-determined locations to which resources are allocated.

There were a few other modeling projects we found in others' work with the USCG. First, in [10], Purdue University developed a visual system to analyze historical USCG response and assess risks that may occur in the surrounding maritime environment. The system analyzed patterns, trends, and anomalies among USCG SAR operations and associated boats, allowing the USCG to be able to identify areas with elevated risk. [10] obviously differs from ACCAM GOM (as well as from BAM I and II), as it is looking at SAR in particular, and its effect on the surrounding environment through a visual modeling system. ACCAM GOM is instead optimizing the allocation of USCG aircraft. A more qualitative study is found in [11]. [11] is a qualitative thesis describing resource allocation and the USCG, but is more of a political or social science examination at the USCG in the 1970s (when it was written), and does not try to build a model. It differs from ACCAM GOM and even BAM I and II, as it does not attempt to model for any particular resource (i.e. boats or aircraft), and does not perform any quantitative modeling. Instead, the thesis describes the USCG resource system as a whole.

Last, another resource allocation paper [12] seeks to perform resource allocation with large fleets of buses. This research seeks to maintain service standards with a limited budget with an aging fleet. The authors of [12] developed a model for allocating funds among fleet improvement programs by minimizing total cost. Integer programming was used to solve this problem with branch and bound. However, this research is also different than ACCAM GOM, as it does not model the problem in the same way by using scenarios (as it is a very different type of problem).

2.3 Background – USCG Air Stations and Missions

The application of the model to be presented in this paper is allocating USCG aircraft to USCG Air Stations. There are many subtleties to take into account in order to accurately understand the problem. Unlike modeling for the USCG boats, the USCG aircraft have more specifics and complications to consider with the aircraft themselves, as well as their interaction on the station and global levels.

The primary mission of the USCG aviation is the Search and Rescue (SAR) mission. SAR occurs when the USCG station receives a distress call, and must respond quickly with aid. Sometimes these calls are false alarms, but the USCG aims to respond to all SAR events. Each station has an Area of Responsibility (AOR) so that when a SAR distress call is received, the stations know which one is to respond. The frequency of these SAR missions varies by station, as some stations are much busier than others. In addition, the frequency of SAR cases by station has seasonality, which may differ by station as well. SAR case frequency on the station level may depend on the day of the week, the time of the day, the month of the year, etc., which again may vary by USCG Air Station. In addition to the very important SAR mission, the Air Stations have other missions as well, such as Law Enforcement, Ports and Waterways Coastal Security, etc. that while less important than SAR, must also be fulfilled. Air Stations also have required training to maintain pilot proficiency.

To simplify the idea of an Air Station, it can be thought of as similar to a firehouse. The goal is to maintain a full SAR crew of personnel and aircraft and to be ready at all times in case a SAR event happens, as these may come in at all hours of the day or night. Obviously certain times are more common for these events to occur, but the USCG should always be ready for these SAR cases.

In addition, the USCG aircraft have characteristics that need to be considered. Each station has a specific number of aircraft that are required for readiness over a 24 hour period. Each aircraft has a specific number of allowed flight hours per year, which is driven by personnel, parts, cost, etc. Flight hours are needed for SAR response, for station-specific missions, and for training to maintain pilot proficiency. In a sense, aircraft are more complex than boats in terms of modeling a Resource Allocation Problem. Additional factors to be taken into account with aircraft are: their performance including speed, range, and endurance; their reliability and availability; their life cycle costs; and the space / hangar capacity at each station. Each aircraft type has a specified range that it can reach in a certain length of time. Moreover, aircraft availability is a function of maintenance. There are two types of aircraft maintenance: scheduled maintenance and unscheduled maintenance. The former happens at specified intervals, similar in a sense to a car. Most scheduled maintenance events are scheduled either by calendar or by hours flown, and are initiated on or just before these milestones. However, there is some allowable variance to avoid taking an Air Station below its required readiness level during times of heavy unscheduled maintenance. The latter can also happen, and this can occur in a few different ways. Unscheduled maintenance can happen when unexpected issues are discovered during scheduled maintenance that need to be taken care of prior to being able to use the aircraft again. It can also occur when aircraft are being prepared for use, and they must be taken care of prior to being flown.

ACCAM STATION [13] is a discrete event simulation that takes all of these things into account, and in particular includes randomly generated SAR and unscheduled maintenance events. Of particular importance is the probability of not being able to respond to a SAR event within a specific amount of time. ACCAM STATION was a simulation model of each USCG Air Station generating performance metrics resulting from various scenarios. Each scenario was determined by a large number of relevant Air Station attributes, including stationed aircraft, operational level, historical Search and Rescue mission response, historical maintenance processes, deployment requirements, and other mission requirements. In some cases, a simpler binomial model seems to give outputs that work as well as those of ACCAM STATION, in which case that model can be used.

3. Approach

3.1 Approach

The basic goal of the ACCAM GOM aircraft resource allocation model was to keep SAR success rate at a specified high level and also to meet non-SAR mission requirements. The USCG needed a tool to be used by an analyst conducting studies. Some examples of basic questions the USCG would like to answer would be: What class of aircraft to assign to a given station or stations, as some stations may have slack capacity? During the low demand times, can a plane be deployed (i.e. sent away for use)? What would happen if the class of an aircraft were changed at a specific station, or if an aircraft type was being phased out and had to be replaced with something different? What if the USCG wanted to reduce their aircraft fleet size? What if the USCG wanted to move an Air Station to a slightly different location?

When modeling this problem, aircraft performance, reliability, cost, and maximum allowable flight hours had to be taken into account. Recall that reliability is a function of maintenance (either scheduled maintenance or unscheduled maintenance). The model also had to take into account station-specific requirements including mission demands, SAR rate and demands, training hours, capacity of an Air Station, etc. Also, it had to consider global USCG-wide requirements such as budget, deployment requirements, and an upper limit of number of total aircraft of each type. Aircraft might be assigned to a deployment center, a depot of aircraft, but that can be deployed to USCG Air Stations for specific missions. These deployment centers have aircraft solely for non-SAR uses. Also, the aircraft can have deployed days at sea (also considered a deployment center). Therefore, there are various considerations at all levels: SAR, aircraft, station-specific, deployment center-related, and global requirements.

3.2 Model Overview

The ACCAM GOM has multiple pieces. Initially, we considered the Air Stations individually. First, the model must choose a "scenario" for each station, while meeting certain global requirements USCG-wide. Each scenario includes a possible resource assignment for a station. Additional parameters associated with a scenario are obtained from ACCAM STATION or a simpler binomial model, e.g., parameters such as the predicted expected number of unmet SAR events. Each station has many possible scenarios. We must generate many scenarios for each station, and then the model considers all of them and optimally determines the "best" selection, keeping performance above a certain threshold, while minimizing the cost.

Each scenario has a large number of relevant attributes. Some of these are deterministic and provided by USCG analysts; others control the random processes and are determined by ACCAM STATION or a simpler binomial model and historical USCG datasets. As there are a large number of scenarios for each station, their feasibility had to be determined prior to choosing optimal assignments. These scenarios will be described in more detail in Section 3.3.

Once the scenarios were created for each USCG Air Station, global constraints also had to be considered. The model seeks to minimize the total cost, while meeting global requirements. It must not assign more aircraft to the Air Stations than are available, meet training, mission, and deployment requirements, meet acceptably high global SAR response rate, and also meet system-wide aircraft availability requirements. The output of ACCAM GOM is an efficient (cost-minimizing) assignment of aircraft to air stations and to a deployment center. We define the objective function precisely below.

Table 1: Basic Sets

S : The set of USCG Stations; $s \in S$

A : The set of aircraft types; $a \in A$ is an aircraft type

D : The set of deployment locations; $d \in D$ is a deployment location

T : The set of time periods (e.g. 12 months); $t \in T$ is a time period

Δ : The set of deployment demands; $\delta \in \Delta$ is a deployment demand, described as the following 5-tuple:

- **δ(a)** : The required aircraft type
- **δ(n)** : The number of aircraft required
- **δ(d)** : The deployment location
- **δ(t)** : The time period of the deployment
- **δ(h)** : The expected flying hours per aircraft on this deployment

Table 2: Basic Sets and Scenarios Additional Notation

H(a,s,d) : A matrix of roundtrip flying hours from station s to deployment location d for an aircraft of type a

L(a,s) : The minimum number of aircraft of type a to be assigned to station s

U(a,s) : The maximum number of aircraft of type a that can be assigned to station s

N(a) : The total number of aircraft of type a available in the system

F(a,s) : The required flying hours on missions, other than Search and Rescue (SAR), deployment, and training for aircraft of type a assigned to station s

τ(s) : The required training as a fraction of total flying time at station s; e.g. $\tau(s) = 0.4$ means that 40% of the total flying time is required to be training at station s (some of this must be flown on aircraft assigned to the station, some may be performed on a flight simulator)

$\boldsymbol{\tau_{max}(s)}$: The maximum fraction of training time at station s that can be performed on a flight simulator; e.g. $\tau_{max}(s) = 0.5$ means that at most half of the required training hours at station s can be performed on a simulator

$\mathbf{Q_0}$: The base capacity of the flight simulator in hours

$\mathbf{Q_1}$: The additional capacity of the flight simulator that can be created at some hourly cost

$\mathbf{C_Q}$: The hourly cost of the additional flight simulator capacity

$\mathbf{U_{BRV}(s)}$: The maximum number of non-bravo events allowed at station s

$\mathbf{T_{BRV}(s)}$: The minimum fraction of time in a bravo state required at station s

$\mathbf{T_{avg}}$: The lower bound on the system wide average of the fraction of time in a bravo state; typically $T_{avg} > T_{BRV}(s)$

$\mathbf{P_{SAR}(s)}$: The minimum required SAR response rate at station s

$\mathbf{P_{avg}}$: The minimum required SAR response rate system wide; typically $P_{avg} > P_{SAR}(s)$

3.3 Data

There are many inputs for the global model. Tables 1 and 2 describe the basic sets and scenarios for input into this model. A scenario consists of a USCG station ($s \in S$), an aircraft type ($a \in A$), and deployment information in the form of a 5-tuple (required aircraft type, number of aircraft required, deployment location, time period of the deployment, and the expected flying hours per aircraft on this deployment).

Table 3 describes the scenarios for input into the ACCAM GOM. In particular, the station for which the scenario was generated, the number of aircraft of each type assigned to the station, the number of yearly hours for which aircraft of each type are budgeted, the readiness, the yearly cost, and the expected number of SAR events at the station The scenario also includes the predicted expected number of unmet SAR events due to aircraft not being available because of either scheduled or unscheduled maintenance, the expected number of flying hours on SAR missions at that station and the expected number of "no-bravo" events at a station in that particular scenario. (A "no-bravo" is when there are no aircraft available due to scheduled or unscheduled maintenance; this is of concern even if it does or does not "count" as a missed response SAR event – i.e. no matter if a SAR event happens or not, the Air Station is unable to meet one should one occur). The scenarios also include the expected fraction of time in a bravo state (i.e. each state; what fraction of time are only 1 aircraft available, 2 aircraft available, etc.) and the number of aircraft assigned for a deployment (where the deployment is specified in terms of number of aircraft of each type required, time period, location, and flight hours). Input for the scenarios comes from real data provided by USCG, and from ACCAM STATION or a simpler binomial model.

Once the scenarios are created or generated in some manner, they must be checked to make sure that they are feasible. Table 4 shows the feasibility constraints for each of them. These are constraints that do not depend on the decision variables, but rather apply individually to each scenario.

With our notation we have the total flying capacity (in hours) assigned to station $s = \sigma$ (name) in each scenario σ for aircraft of type $a \in A$ after deployments, stated in Equation (1a). Equation (1b) imposes the minimum number of aircraft of type a at station = σ (name) in any time period in scenario σ. The limitations in equations (1c) and (1d) relate to the number of aircraft a station must hold. Scenarios must also satisfy "quality of operations" constraints, including meeting bravo requirements most of the time and having a high SAR response rate, which is stated in equations (1e) and (1f) respectively. All of these can be thought of as "filtering" equations to weed out the infeasible scenarios prior to modeling or optimizing the Resource Allocation model.

Table 3: Scenarios

σ(name) :	The station for which a scenario was generated; $\Sigma(s) = \{\sigma \in \Sigma \mid s = \sigma(\text{name})\}$ is the set of scenarios generated for station $s \in S$.
σ(a) :	The number of aircraft of type $a \in A$ assigned to station $\sigma(s)$ in scenario σ
σ(a,h) :	The number of hours (yearly) that aircrafts of type a are budgeted in scenario σ
σ(r) :	The level of readiness budgeted in scenario σ
σ(b) :	The total cost (yearly) of operations in scenario σ
σ(SAR) :	The expected number of SAR events at station s
σ(UMT) :	The expected number of unmet SAR events
σ(SAR,h) :	The expected number of flying hours on SAR missions
σ(U) :	The expected number of non-bravo events in scenario σ
σ(T) :	The expected fraction of time in a bravo state in scenario σ
σ(δ) :	The number of aircraft assigned for deployment δ; recall that the attributes of $\delta \in D$ specify the time period, location, flight hours, etc.

Table 4: Feasibility Constraints for the Scenarios

$\eta(a,\sigma) = \sigma(a) * \sigma(a,h) - \sum_{\delta \epsilon D:\delta(a)=a} \sigma(\delta) * (H(\delta(a), \sigma(name), \delta(d)) + \delta(h))$	(1a)
$\mu(a,\sigma) = \sigma(a) - \max_{t\epsilon T} \sum_{\delta\epsilon D:\delta(a)=a,\delta(t)=t} \sigma(\delta)$	(1b)
$\sigma(a) \leq U(a, \sigma(name)) for\ all\ a\epsilon A\ and\ \sigma\ \epsilon\ \Sigma$	(1c)
$L(a, \sigma(name)) \leq\ \mu(a,\sigma) for\ all\ b\ \epsilon\ A\ and\ \sigma\ \epsilon\ \Sigma$	(1d)
$\sigma(T)\ \geq\ T_{BRV}(\sigma(name)) for\ all\ \sigma\epsilon\ \Sigma$	(1e)
$1 - \frac{\sigma(UMT)}{\sigma(SAR)}\ \geq\ P_{SAR}(\sigma(name)) for\ all\ \sigma\ \epsilon\ \Sigma$	(1f)

3.4 The Model

Once the scenarios have been created for each USCG Air Station, as described in Section 3.3, then we can introduce the ACCAM GOM resource allocation model. Its formulation consists of the following decision variables: x_σ, binary, which takes on value 1 if scenario σ is chosen and 0 otherwise; $y_s(a)$, which represents the number of training hours assigned to station *s* by aircraft type *a*; $z_s(a)$, which accounts for the number of training hours assigned to station *s* by aircraft type *a* that are able to be performed on a flight simulator; Z, which denotes the global number of training hours that are performed with a flight simulator; and Z^1, which yields the number of training hours performed on a flight simulator in excess of the capacity. With the exception of x_σ, all these variables are continuous. These decision variables are shown in Table 5.

Table 5: Decision Variables

$\mathbf{x_\sigma}$: σ ϵ Σ, binary selector variables; $x_\sigma = 1$ if scenario σ is chosen in the solution, and $x_\sigma = 0$ otherwise
$\mathbf{y_s}$ **(a)** : s ϵ S, the number of training hours at station s flown on aircrafts of type a assigned to station s, ≥ 0
$\mathbf{z_s}$ **(a)** : s ϵ S, the number of training hours for type a aircrafts at station s performed on flight simulator, ≥ 0
Z : total number of training hours performed on flight simulator, ≥ 0
$\mathbf{Z^1}$: The number of training hours performed on flight simulator in excess of the base capacity, ≥ 0

The global optimization model for the USCG, ACCAM GOM (see Table 6), has an objective function to minimize the total cost (shown in Equation 2). Exactly one scenario must be chosen for each station, as shown in Equation (3a). The model cannot assign more aircraft than are available, which is stated in Equation (3b). Equation (3c) ensures that deployment requirements are met. Per equations (3d) and (3e) respectively, each station must meet the training hours and other mission requirements. Furthermore, equations (3f) and (3g) are limitations on the flight simulator, whereas by Equation (3h) the system-wide Search and Rescue (SAR) response rate should be above the specified high response success requirement. In addition, Equation (3i) forces the global bravo requirement to be met.

ACCAM GOM was modeled as a variant to the well-studied Resource Allocation Problem (see Section 2.2). ACCAM GOM was modeled as a mixed-integer-program (MIP) with binary, integer, and continuous variables. A solution to this MIP could be found by the powerful heuristic technique Branch-and-Bound, which attempts to cleverly search through all possible combinations of integer and binary variable choices (continuous ones are "easy" to solve for by other methods) by bounding the best solutions' objective values within an ever-shrinking interval. The upper and lower bounds of such an interval were crucial to a fast algorithm, as they were used to cut down the search-tree formed by the fixing of variables. The problem was encoded in a leading commercial optimization package called Xpress, which has a state-of-the-art Branch-and-Bound method built in. This package comes equipped with its own programming language, Mosel, which allows high-level algebraic modeling as well as lower-level functional programming.

Table 6: The Model

Minimize $\sum_{\sigma \epsilon \Sigma} \sigma(b) * x_\sigma + C_Q * Z^1$ (2)

subject to

$\sum_{\sigma \epsilon \Sigma:\sigma(\mathrm{name})=\mathrm{s}} x_\sigma = 1 \; for\ all\ s \epsilon S$ (3a)

$\sum_{\sigma \epsilon \Sigma} x_\sigma * \sigma(a) \leq N(a) for\ all\ a \epsilon A$ (3b)

$\sum_{\sigma \epsilon \Sigma} \sigma(\delta) * x_\sigma \geq \delta(n) for\ all\ \delta \epsilon \mathrm{D}$ (3c)

$y_s(a) + z_s(a) \geq \tau(s) * \sum_{\sigma \epsilon \Sigma:\sigma(\mathrm{name})=\mathrm{s}} x_\sigma (\sum_{a \epsilon A} \sigma(a) * \sigma(a,h)) \; for\ all\ s \epsilon S\ and\ a \epsilon A$ (3d)

$\sum_{\sigma \epsilon \Sigma:\sigma(\mathrm{name})=\mathrm{s}} x_\sigma * (\eta\,(a,\sigma) - \sigma(SAR,h)) \; \geq F(a,s) + y_s(a) \; for\ all\ s \epsilon S\ and\ a \epsilon A$ (3e)

$\sum_{s \epsilon S} \sum_{a \epsilon A} z_s(a) = Z \; \leq \; Q_0 + \; Q_1$ (3f)

$Z^1 \geq Z - \; Q_0$ (3g)

$\sum_{\sigma \epsilon \Sigma} x_\sigma * \left(1 - \frac{\sigma(UMT)}{\sigma(SAR)}\right) \geq |S| * \; P_{avg}$ (3h)

$\sum_{\sigma \epsilon \Sigma} x_\sigma * \sigma(t) \geq |S| * \; T_{avg}$ (3i)

4. Future Work

ACCAM GOM is expected to be implemented with real data for all USCG Air Stations and produce realistic results. This would aid the team to be able to determine potential cost savings for the USCG. Also, the optimization model can be used to demonstrate the potential efficiencies of proposed infrastructural changes, such as the introduction of a deployment center with certain characteristics.

Moreover, this work may be useful in other resource allocation problems that the USCG is facing as well as with other resource allocation problems outside of the USCG.

Apart from implementing the model, it should be possible to find other functionalities for the model. For instance, implementing ACCAM GOM with a visualization tool would be a nice addition. Another avenue to be explored is whether the “sharing” of aircraft between Air Stations (on top of deployment centers already taken into account in ACCAM GOM) is beneficial.

5. Conclusion

This project was a successful collaboration between us and the USCG. We were able to model the real problem of the USCG Air Stations, and come up with a way to optimize the allocation of USCG aircraft to USCG Air Stations while still meeting USCG business requirements.

This project presented many challenges. Using optimization tools on a real world problem is always difficult, and this project had close coordination between the modelers and the USCG experts. This was required for us to be able to accurately represent the needs of the USCG.

Acknowledgements

We would like to thank the United States Coast Guard for their partnership with this project and prior ones. We would also like to acknowledge the Department of Homeland Security Award 2009-ST-061-CCI002-05 for allowing us to perform this joint research and for funding the partnership between USCG, CCICADA, and FICO, the creator of the Xpress software used in this project.

References

1. USCG Research and Development Center, 2013, “Acquisition Directorate: Modeling Capability Development Plan for the Aviation Capability and Capacity Allocation Model (ACCAM).”
2. Boros, E., Kantor, P., Roberts, F., Nelson, C., and Oster, M., 2013, “Boat Allocation Module (BAM) Final Report.”
3. Boros, E., Nelson, C., Oster, M., Aikins, T., Babcock, K., Ball, P., Conrad, C., Hanson, K., Kantor, P., Pottenger, W. M., Roberts, F., Sharkey, T., Wojtowicz, J., 2012, “A Resource Allocation Model for Assigning Boats to Stations in the United States Coast Guard.” *INFORMS Annual Meeting 2012*, Phoenix, AZ, Oct 2012.
4. Oster, M., 2014, “Two Applications of Combinatorial Optimization,” Dissertation Thesis, Rutgers University.
5. Boros, E., Kantor, P., Roberts, F., Nelson, C., and Oster, M., 2014, “Boat Allocation Model for the United States Coast Guard,” in preparation.
6. Conrad, C., Egan, D., Hanson, K., McGinity, C., Nakamura, B., Nelson, C., and Radar, T., 2014, “BAM II Demo to USCG Office of Boat Forces.”
7. Ibaraki, T. and Katoh, N, 1988, *Resource Allocation Problems: Algorithmic Approaches.* Cambridge, MA: The MIT Press, 1988.
8. Dimopoulo, M. and Giannikos, I., 2004, “Towards an integrated framework for forest fire control,” *European Journal of Operational Research*, 152, 476-486.
9. Lee, E. K., Chen, C. H., Pietz, F., and Benecke, B., 2009, “Modeling and Optimizing the Health Infrastructure for Emergency Response,” *Interface*s 39.5, 476-490.
10. Malik, A., Maciejewski, R., Maule, B., and Ebert, D. S. “A Visual Analytics Process for Maritime Resource Allocation and Risk Assessment,” *IEEE Conference on Visual Analytics Science and Technology (VAST).*
11. Losea, R. J., 1976, “Planning, Programming and Budgeting for Resource Allocation in the United States Coast Guard,” Department of the Navy Naval Postgraduate School Thesis.
12. Mishra, A., Sharma, S., Khasnabis, S., and Mathew, T., 2013, “Preserving an aging transit fleet: An optimal resource allocation perspective based on service life and constrained budget,” *Transportation Research Part A: Policy and Practice*, 47, 111-123.
13. Boros, E., Baron, J., DeMarco, R., Kantor, P., McGinity, C., Nakamura, B., Nelson, C., Oster, M., Pottenger, W. M., Ricks, B., Roberts, F., Vaughn, T., Wojtowicz, J., Ball, P., Conrad, C., Hanson, K., Lehocky, M., Rader, T., and Sharkey, T., 2013, "Aviation Capability and Capacity Allocation Model," *INFORMS Annual Meeting*.