

The ACCAM Model: Simulating Aviation Mission Readiness for U.S. Coast Guard Stations

Curtis McGinty, Endre Boros,
Paul Kantor, Fred S. Roberts,
Brian Nakamura, Christie Nelson, Brian Ricks

CCICADA,
Rutgers University
Piscataway, NJ 08854
Website: <http://www.ccicada.org>

Thomas J. Rader, Kevin J. Hanson,
Patrick J. Ball, Chad M. Conrad,

United States Coast Guard
Headquarters
2100 2nd St. SW STOP 7000
Washington, DC, 20593

Abstract—We present a model and discrete event simulation of USCG Air Stations, accounting for the mission demands and maintenance procedures pertaining to USCG aircraft. The simulation provides aircraft availability distributions and mission performance metrics based on varying input scenarios, including changes in the number of stationed aircraft and maintenance targets. The Air Station model is novel in its relatively simple, easily tunable, renewal process treatment of maintenance procedures, mitigating the need for the modeling of complex maintenance subprocesses and the resulting statistical estimation of numerous parameters. The simulation also models mission requirements such as Search and Rescue that are stochastic in time and space. Simulations are consistent with historical data and offer insights into hypothetical scenarios.

I. INTRODUCTION

United States Coast Guard (USCG) Aviation fleet allocation to USCG Aviation Air Stations has long been dictated by assumptions of operational response, aircraft availability, and mission demands which may or may not still be valid. With the increasing demand for Aviation forces to support forward-deployed surface forces, the USCG would like the means to determine the optimal assignment of aircraft among Air Stations and deployment sites in consideration of the respective mission demands and performance targets. In order to meaningfully identify this assignment, it was necessary to develop a modeling capability of a USCG Air Station and its aircraft that would allow us to comprehensively analyze mission demands, response, and “business rules” at the Air Station level in order to estimate their impact on aircraft operational performance.

The USCG has both fixed-wing and rotary-wing aircraft and, moreover, several types of each. Each Air Station has assigned either a single type of rotary-wing or fixed-wing aircraft, or a single type of each. So far, our model assumes that there is only one class of aircraft at the Air Station, namely rotary-wing, and only one type of rotary-wing aircraft. Each Air Station also has certain required missions. A key one of these is Search and Rescue (SAR), which is complicated to plan for since SAR events occur stochastically in both time and space. We seek to equip an Air Station with aircraft sufficient to guarantee that aircraft will be available to respond to SAR events in a timely way. We also have to take into

account other missions required of the aircraft, such as Law Enforcement (LE), Smuggling Interdiction, etc. Complicating this are the facts that scheduled maintenance is needed to keep the aircraft flyable and unplanned maintenance from breakdowns or mishaps can occur. Ultimately, we have a fleet of aircraft to assign to Air Stations. How we make the assignment in such a way as to attain some global goals is part of another project [5] for which the project described here is a necessary precursor. In this project, we assume we have aircraft assigned to a station and model the availability and performance of those aircraft.

To this end, a discrete event simulation was developed for the USCG to model the capabilities, requirements, and operational procedures of an Air Station and provide operational performance metrics. Relevant performance metrics include the number of events when no mission capable (MC) aircraft are available at the Air Station (a so-called *no-Bravo* event), percentage of SAR missions met within a target response time (the *SAR met rate*), and generally the distribution of the number of available aircraft (an aircraft *state distribution*). Fundamentally, driving all of these performance metrics is *aircraft availability* and reliability; that is, the availability of an aircraft to successfully service a mission demand, should one arise or not, and the likelihood that it will not have an equipment failure that prevents completion of the mission sortie. A *sortie* is defined as an operational flight by a single military aircraft. The meaning of performance metrics will be discussed further in Section II-A.

In order to sufficiently capture the processes driving aircraft availability, the model incorporates several pertinent attributes of Air Station operations, including the type and number of aircraft, aircraft capabilities, primary mission demands, maintenance processes, deployment requirements, aviator training requirements, other mission demands, and established business rules. The nature of the various aircraft demands necessitates a variety of modeling methods. For example, some events occur according to a schedule, which in some cases might be contingent on other demands and the number of available aircraft; some processes display a stochastic component, again possibly conditional on other operations and number of available aircraft; many interdependencies exist among the various processes, and the interactions are ostensibly complex. It is

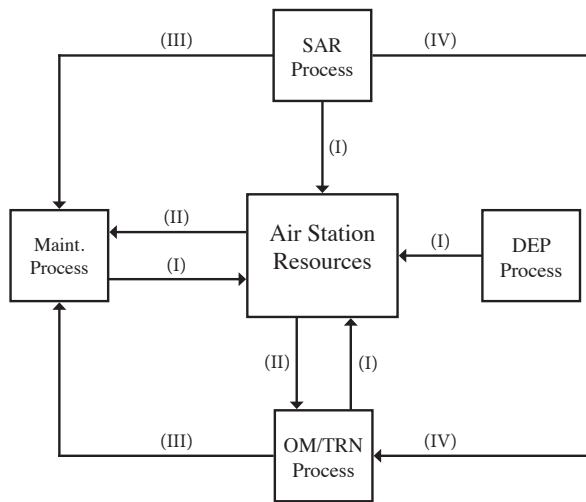


Fig. 1. Interdependence among Maintenance, SAR, Other Missions and Training, and Deployment Processes. Directed edges illustrate the propagation of influence. (I) Aircraft Demands; (II) MC Aircraft State Information; (III) Completed Flight Hours; (IV) Diversion Demand.

therefore appropriate to model each aspect as a discrete event process with rule-based interdependencies among processes.

Of particular importance are maintenance processes. Central to Air Station operations, maintenance processes constitute a primary means of system control. Air Station engineers strategically manage procedures to achieve operational requirements and efficiently utilize resources. The underlying heuristics and ad hoc decisions are sometimes nebulous and difficult to model directly, but play a significant role in Air Station management. Additionally, the element of control over certain aspects of maintenance processes makes this an attractive candidate for potential refinements to improve fleet efficiency. One goal of the simulation is to meaningfully provide performance estimates for proposed changes to maintenance guidelines and the number of stationed aircraft, so-called “what if” scenarios for which no data currently exists.

The model we have developed falls under the Coastal Operation Analytical Suite of Tools (COAST) Aviation Capability and Capacity Assignment Module (ACCAM). The global model, termed the ACCAM Global Optimization Model (or ACCAM GOM), together with the ACCAM Simulation presented here, are a joint effort between researchers at Rutgers University at CCICADA, and USCG, with close and consistent communication between the partners and researchers. This project is part of an “Engage to Excel” (E2E) initiative of DHS. Under the E2E program, university centers of excellence 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). ACCAM Simulation and ACCAM GOM are follow-on projects to optimally assign aircraft to USCG stations. COAST is a set of modules, usable and updatable individually, but rationally linked together.

II. AIR STATION SIMULATION

A. System Components

We design our simulation to incorporate the following aspects:

- 1) The simulation should have the capability to meaningfully investigate unprecedented hypothetical scenarios (so-called “what if” scenarios), including changes to the number of stationed aircraft, the aircraft capabilities, and the target availability of aircraft at the Air Station. The *target availability* of stationed aircraft refers to the desired percentage of time a single aircraft is in a MC or PMC state; engineers use the target availability of individual aircraft as a proxy to guarantee the availability of at least one aircraft at the Air Station.
- 2) USCG datasets contain a substantial quantity of detailed data on missions and maintenance. In some cases, however, the data is manually entered, contains reporting bias that may vary by Air Station, and occasionally contains data requiring interpretation by USCG domain experts. Some USCG datasets, such as MISLE, are sufficiently large and extensive to admit accurate statistical estimation of parameters after correcting for these factors, whereas others, such as ALMIS, are less so.
- 3) Aircraft maintenance processes have both quasi-periodic and random components, and intrinsically depend on aircraft flight hours, the number of stationed aircraft, and complex engineer decision-making regarding load balancing and PMC aircraft. The structure of datasets¹ render statistical analysis and subprocess modeling challenging and potentially fallacious.
- 4) SAR missions hold a particular priority among USCG Aviation, often require multiple sorties to complete the mission, and display seasonality at the monthly, weekly, and daily timescales.
- 5) Other Missions and Training are predominantly performed for a fixed period of time during routine working hours subject to aircraft availability. Training missions will be diverted to SAR or LE response missions, if required.
- 6) Forward Deployment is determined by fleet-wide considerations. Deployment demands typically consist of a 30-day duty period and a week-long period of extensive maintenance prior to deployment, during which the deployed aircraft is effectively removed from Air Station operations.

To take account of these aspects, the ACCAM Simulation implements Processes which place demands on the stationed aircraft, analogous to a multi-server queue (however without queuing). The Processes include:

- Maintenance Process
- SAR Process
- Other Missions and Training Process
- Deployment Process,

with interdependencies illustrated conceptually in Figure 1.

¹In this case, the ALMIS dataset

Each Process is updatable individually, and propagating interdependencies change implicitly. Four types of interdependence exist among the Processes, as shown in Figure 1:

I Aircraft Demands

Each Process places demands on stationed aircraft for a fixed period of time, which varies by the Process. Missions (SAR and Other Missions and Training (OM/TRN)) require aircraft to complete one or more sorties, each for a random amount of time sampled from the appropriate distribution (see Section II-C); Deployments occupy an aircraft for the duration of the deployment (see Section II-D); Maintenance occupies an aircraft for the duration of the maintenance event (see Section II-B).

II MC Aircraft State and Flight Hour Information

Information about the current number of mission capable (MC) (the state) is relevant for all Processes which are to some degree controllable by the Air Station, namely OM/TRN and Maintenance. OM/TRN incorporates this information by adjusting the number of Training and non-essential missions flown when the state is low. Analogously, Maintenance processes are adjusted by conducting more maintenance when the state is high and delaying non-essential maintenance when the state is low, when possible. OM/TRN also use consumed flight hour information to adjust the number of missions flown heuristically to achieve approximately a constant monthly consumption.

III Completed Flight Hours

The number of flight hours completed on missions (SAR and OM/TRN) contributes to required maintenance as detailed in Section II-B.

IV Diversion Demand

SAR missions hold a precedence over OM/TRN, such that if a SAR distress call arrives during another mission or Training, the aircraft will divert to complete the SAR mission.

B. Maintenance Model

The Maintenance process is central to the accurate modeling of the Air Station, simultaneously incorporating several aspects in an elegant, tunable model. The maintenance procedures described above involve a number of aspects which are difficult to model precisely, for various reasons. Systematically identifying unscheduled maintenance in USCG datasets is challenging, not always possible, and when possible not scalable to larger datasets or for other Air Stations. This is problematic in terms of identifying potential approximations to a random distribution. Even as the MRL provides guidelines for flight-hour-based maintenance, tracking flight hours of individual aircraft requires an assignment scheme for load balancing. However, the significant role of complex engineer decision making in load balancing is particularly difficult to capture via a deterministic, rule-based approach and does not strictly generalize to all Air Stations.

We present a model based on the notion of “maintenance-unaccounted” flight hours of all aircraft at the Air Station and a stochastic component. We observe in the data that maintenance primarily occurs during specific times of the workday: early morning, morning, afternoon, and evening. We identify each as a scheduled maintenance time t , and consider the possibility




Fig. 2. Sample series of maintenance-unaccounted flight hours g at scheduled maintenance times t .

of generating one or more maintenance events at each time sequentially. We introduce the following parameters: the number of flight hours triggering a maintenance event of type k (obtained from the MRL), denoted T_k ; the cumulative number of flight hours consumed by all aircraft at the Air Station up to time t , denoted $F(t)$; and the cumulative number of maintenance events on all aircraft at the Air Station up to time t , denoted $m(t)$. We define the “maintenance-unaccounted” flight hours at the Air Station at time t :

$$g(t) := F(t) - m(t)T_k. \quad (1)$$

Moreover, since $F(t)$ denotes the flight hours of the Air Station, these hours could be distributed among the a aircraft in any way. Considering the case of perfect load balancing, where all a aircraft fly T_k hours without triggering a maintenance event, we see that aT_k is the largest $g(t)$ which may not trigger a maintenance event. Similarly, considering the other extreme of load balancing where $F(t)$ is flown only by a single aircraft, the least value of $g(t)$ to possibly trigger a maintenance event should be T_k . Therefore the probability of a maintenance event occurring at time t is given by

$$\text{Prob. of MTN} = \begin{cases} 0, & \text{if } g(t) < T_k, \\ P(g(t)), & \text{if } T_k \leq g(t) \leq aT_k, \\ 1, & \text{if } g(t) > aT_k, \end{cases} \quad (2)$$

where $P(g(t))$ is a given cumulative distribution function. We illustrate this in Figure 3.

Using the parameter $g(t)$ and a probability distribution $P(t)$, we have a simple method for generating maintenance events: At every scheduled maintenance time t , generate a maintenance event with probability $P(g(t))$. If we let $\tau(a)$ be the maximum number of maintenance events to occur in parallel and β be the current number of Bravo aircraft at the Air Station, we can augment the procedure by repeating the method $\min\{\tau(a), \beta\}$ times, each time recomputing $P(g(t))$.

The model is easily calibrated to data through the parameter T_k and the parameters of the distribution P . The ACCAM Simulation uses a Gamma CDF and $T_k \equiv T$, for all k , calibrated to closely approximate historical maintenance data,




Fig. 3. Depicts the probability of a maintenance event $m(t)$ occurring at scheduled time t for three possible Gamma distributions. Each distribution represents a varying degree of load balancing.

in terms of the average total number of maintenance events and the variance in frequency (time between successive events). We note that letting $T_k = T$, for all k reflects the “pacing item” or “masking” phenomenon in [1] (and observed in ALMIS).

For a fixed distribution, the choice of T in some sense captures target availability. Larger values of T allow for longer times between maintenance events on average. As T increases, the Gamma scale parameter can be left constant and shape parameter adjusted to achieve closer agreement with (2). This approach allows an increase in both tails, near T and aT , respecting the possibility of random failures between less frequent maintenance events and reflecting the expected increase in PMC aircraft as target availability increases. Note that due to the renewal property of the parameter g , the calibration depends on the particular mission demands and aircraft of an Air Station.

C. Search and Rescue Model

SAR arrives independently of all other processes and state variables with some seasonal dependence, and therefore we model SAR arrivals as a non-homogeneous Poisson process. Chi-squared analysis of USCG datasets substantiated this model and supported the hypothesis that the seasonal dependence was captured by treating the month of the year (MoY) and the day of the week (DoW) as independent effects. We therefore have

$$\text{SAR} \sim \text{Poisson}(\lambda(\text{DoW}, \text{MoY})). \quad (3)$$

The rate parameters for each month and day of the week are obtained from data by maximum likelihood estimation, and the non-homogenous rate parameter λ is simulated by thinning of the Poisson process via the method of Lewis and Shedler [10].

The durations of SAR events necessarily depend on several random components, both spatial and temporal, as well as logistic constraints. For example, the travel time to the event location, search time, weather conditions, operational requirements, and fuel limitations potentially play a role. The distributions of several of these factors depend on the Air

Algorithm 1 Maintenance Check Procedure

```

1: for every scheduled time  $t$  do
2:   for  $\tau \leq \min\{\tau(a), \beta\}$  do
      $m \leftarrow \text{Unif}[0, 1]$ 
      $\mu \leftarrow P(g(t))$ 
3:   if  $m < \mu$  then
      $t_m \leftarrow \text{Exp}(1/\lambda)$ 
     generate maintenance at time  $t + t_m$ 
      $g(t) \leftarrow g(t) - T_k$ 
4:   end if
      $\tau \leftarrow \tau + 1$ 
5:   end for
6: end for

```

Station. Estimating search and travel times empirically from data precludes the need for spatial modeling of SAR events. A Gamma distribution is often used to model cumulative event durations [12], although data recording bias rendered a good fit elusive. Thus the ACCAM Simulation samples SAR durations from empirical distributions.

Additionally, SAR events often require multiple sorties to complete the mission. To incorporate this aspect, MISLE datasets were analyzed by the Case ID number of individual SAR events to establish the conditional probabilities of having an additional sortie, given having already made $1, 2, \dots, 4$ sorties prior. Flight hours are logged for the duration of each completed sortie servicing a SAR event.

D. Additional Missions Models

Other Missions and Training account for the majority of flight hours at Air Stations [8], and are logged for each sortie. Through elicitation sessions with USCG domain experts, it was determined that the vast majority of such missions and Training occur for a fixed duration according to a schedule during routine working hours. However, the schedule is state dependent, in the sense that the number of FMC and PMC aircraft determine the number of missions and Training to be flown on a given day. Additionally, the Training schedule depends on the current consumption of flight hours on SAR and other missions. Training sorties are adjusted (reduced if missions are above average, increased if below average) to achieve a roughly constant monthly consumption of flight hours in an ad hoc fashion. Moreover, as mentioned above, aircraft currently on Training flights divert, if possible, to SAR missions should one occur concurrently.

Forward Deployment to surface vessels is an important component in aircraft availability at USCG Air Stations due to the relative duration of deployment periods, however its implementation is rather straightforward. Through elicitation sessions with USCG domain experts, it was determined that deployment demands occur for a period of approximately 30 days, with a period of heavy maintenance the week prior. Limitations on the frequency of deployment are also enforced, and the ACCAM Simulation allows for particular deployment periods to be selected randomly or by user input.

III. SIMULATION OUTPUT AND CONCLUSIONS

The ACCAM Simulation effectively models aircraft availability at USCG Air Stations, providing a distribution of the

One Year of SAR Missions and State Transitions




Fig. 4. A visualization of simulation outputs based on hypothetical data, illustrating mission capable aircraft state transitions (bottom) and mission completions (top) over time. Unmet mission demands are identified by dotted lines.

available aircraft at an Air Station and SAR mission success rates for variations in the number of aircraft and target aircraft availability. The simulation behaves as expected throughout its parameter space and provides nontrivial results for “what-if” scenarios.

An example visualization of the ACCAM Simulation output is shown in Figure 4 for an Air Station over the course of one year. The visualization depicts state transitions over time and arriving SAR missions along the top; an unmet SAR demand is identified by a vertical dotted line. ACCAM Simulation results for a particular instance of a single Air Station are presented in Figure 5. The results show that the ACCAM Simulation closely matches the historical data over the same time period. Namely, experiments indicate that available aircraft state distributions are well-approximated by a binomial distribution with parameter depending in a complicated way on the target availability and number of stationed aircraft. The scenario presented in Figure 6 shows the state distribution due to a hypothetical change in maintenance processes reducing the target availability of aircraft. In this scenario, the ACCAM Simulation output shows significant deviation from the binomial distribution with parameter equal to the reduced target availability, but is well approximated by a best-fit binomial whose parameter is larger than the reduced target availability.

We note one limitation of the model is that it does not allow for the simultaneous consideration of different types of rotary-wing aircraft at an Air Station. Parameters may be

adjusted to simulate a different type of rotary-wing aircraft, but when two or more types are present simultaneously additional effects need to be considered and implemented. As mentioned above, due to maintenance overhead considerations the USCG currently operates, with a small number of exceptions, Air Stations with either fixed-wing or rotary-wing aircraft, and additionally with a single type of fixed- or rotary-wing aircraft. Future investment in infrastructure may wish to investigate multiple types at the same Air Station, and the simulation would require extension in this case.

Future work will proceed in a number of directions. First, a simplified version of the ACCAM Simulation will be used in the scenario generation of the ACCAM Global Optimization Model [5].

Second, the model will be extended to include additional classes of aircraft, such as fixed-wing aircraft. We will add other types of rotary-wing aircraft next, and then allow for more than one type and/or class of aircraft at an Air Station.

Moving beyond specific aircraft, we plan to consider abstract aircraft defined by attributes such as speed, range, fuel capacity, etc., thus allowing the model to do a variety of “what if” experiments with different types of investments in future performance of aircraft.

Finally, a novel renewal process maintenance model is presented and utilized effectively as a lightweight alternative to previous aviation maintenance models in the literature. The renewal process model provides an effective and convenient

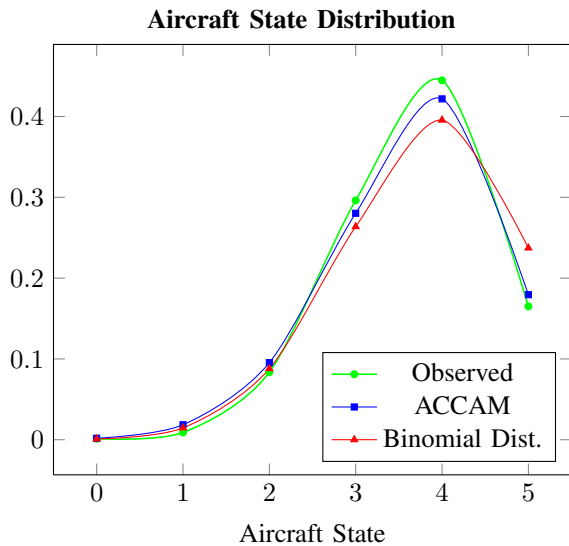


Fig. 5. Aircraft state distribution for a single Air Station: observed, simulated, and approximated by a binomial distribution.

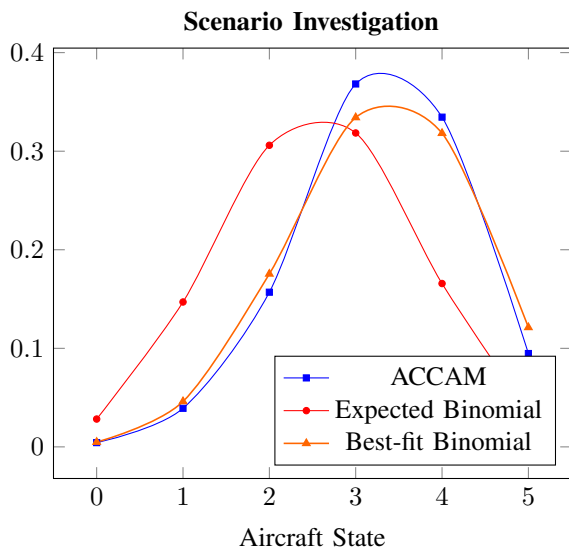


Fig. 6. Aircraft state distribution for a single Air Station in the scenario of a hypothetical maintenance refinement.

proxy for complex maintenance operations without granular modeling of individual aircraft or maintenance subsystems and is particularly suited to discrete event simulation. The model meaningfully incorporates aircraft utilization and historical maintenance data and allows for effects of engineer decision-making, PMC aircraft, and the adjustment of target availability through a modicum of parameters. The maintenance model will be studied further through application to a variety of scenarios as directed by USCG analysts. In due course, the model will be analyzed as a renewal process in order to better understand the parameters and the model's use as a simulation tool.

ACKNOWLEDGMENT

We would like to thank the United States Coast Guard for their partnership with this project and prior collaborations.

We would also like to thank the Department of Homeland Security under Award 2009-ST-061-CCI002-05 to Rutgers University for affording us the opportunity to conduct this joint research and for funding the partnership between USCG and CCICADA.

REFERENCES

- [1] H. S. Balaban, R. T. Brigantic, S. A. Wright, and A. F. Papatyi. A Simulation Approach to Estimating Aircraft Mission Capable Rates for the United States Air Force. In *Proceedings of the 2000 Winter Simulation Conference*, pages 1035–1042. Winter Simulation Conference, December 2000.
- [2] M. Bazargan and R. N. McGrath. Discrete Event Simulation to Improve Aircraft Availability and Maintainability. In *Proceedings of the Annual Reliability and Maintainability Symposium*, pages 63–67. IEEE, May 2003.
- [3] N. Brown and S. Powers. Simulation in a Box (A Generic Reusable Maintenance Model). In *Proceedings of the 2000 Winter Simulation Conference*, pages 1050–1056. Winter Simulation Conference, December 2000.
- [4] F. W. Ciarallo, R. R. Hill, S. Mahadevan, V. Chopra, P. J. Vincent, and C. S. Allen. Building the Mobility Aircraft Availability Forecasting (MAAF) Simulation Model and Decision Support System. *Journal of Defense Modeling and Simulation*, 2(2):57–69, April 2005.
- [5] E. Boros, P. Kantor, C. McGinity, B. Nakamura, C. Nelson, M. Oster, W. Pottenger, B. Ricks, F. S. Roberts, J. Wojtowicz, P. Ball, C. Conrad, K. Hanson, M. Lehocky, T. Rader, and T. Sharkey. ACCAM Global Optimization Model for the USCG Aviation Air Stations. Submitted for publication, 2014.
- [6] E. Boros, P. Kantor, F. S. Roberts, C. Nelson, and M. Oster. Boat Allocation Model for United States Coast Guard. Submitted for publication, 2014.
- [7] M. Oster. *Two Applications of Combinatorial Optimization*. PhD thesis, Rutgers University, 2014.
- [8] A. Iyer and V. Deshpande. Executive Summary: MiDAP - A Final Report. Private Communication, October 2013.
- [9] M. Lehocky. Modeling Capability Development Plan for the Aviation Capability and Capacity Allocation Model (ACCAM). Technical report, USCG Research Development Center, 2013.
- [10] P. A. W. Lewis and G. S. Shedler. Simulation of Nonhomogeneous Poisson Processes by Thinning. Technical report, Naval Postgraduate School, 1978.
- [11] A. Malik, R. Maciejewski, Y. Jang, S. Oliveros, Y. Yang, B. Maule, M. White, and D. S. Ebert. A visual analytics process for maritime response, resource allocation and risk assessment. *Information Visualization*, October 2012.
- [12] V. Mattila, K. Virtanen, and T. Raivio. A Simulation Model for Aircraft Maintenance in an Uncertain Operational Environment. In *Proceedings of 17th European Simulation Multiconference*, pages 296–304. SCS Europe BVBA, June 2003.
- [13] K. Mooney and G. Sanchez. Improved Aviation Readiness and Inventory Reductions through Repair Cycle Time Reductions Using Modeling and Simulation. Master's thesis, Naval Postgraduate School, 1997.
- [14] B. U. Nguyen and K. Y. K. Ng. Modeling of Canadian Search and Rescue Operations. *Military Operations Research*, 5(1):5–16, 2000.
- [15] K. S. Upadhy and N. K. Srinivasan. Availability of Weapon Systems with Air-attack Missions. *Journal of Defense Modeling and Simulation*, 1(2):111–121, April 2004.
- [16] U.S.C.G. U.S. Department of Homeland Security. Aeronautical Engineering Maintenance Management Manual, April 2011.
- [17] U.S.C.G. U.S. Department of Homeland Security. U.S. Coast Guard Addendum to the United States National Search and Rescue Supplement (NSS) to the International Aeronautical and Maritime Search and Rescue Manual (IAMSAR), January 2013.
- [18] S. E. Vigus. A Simulation-based Analysis of the Impact of In-sourcing a Major Process Element on the Coast Guard HH-60J Depot Maintenance Process. Master's thesis, Air Force Institute of Technology, 2003.

- [19] A. Zavala and E. Ramos. Optimal cyclic replacement policy in MSS maintenance via Binomial-PSO. In *Proceedings of the 14th International Conference on Genetic and Evolutionary Computation*, July 2012.