Improving the Deployment and Employment of Naval Air Connectors

Daniel Reich[1], Moshe Kress[1], Javier Salmeron[1], and Giovanni Macaluso[1]

[1]Department of Operations Research, Naval Postgraduate School, Monterey, California

Abstract

The Navy Air Logistics Office (NALO) is tasked with planning air cargo and passenger transport missions that fulfill demands around the world. These missions are executed by the Fleet Logistics Support Wing (FLSW), which consists of 11 transportation squadrons spread over 10 bases in different locations in the continental U.S. and Hawaii. Using historical air cargo demand data provided by NALO, we analyze FLSW's operational practices and squadron allocation, with the objective of identifying potential improvements in aircraft deployment. We develop two variants of a mixed-integer program: (1) our *employment* model optimizes assignments of aircraft to missions, and (2) our *deployment* model reallocates aircraft to bases. Both models explicitly minimize the cost of assigning aircraft to missions. We compare the performance of employing the aircraft as currently allocated against the proposed aircraft redeployment allocations. Our computational experiments identify inefficiencies in current deployment and show potential directions for improvement.

1 INTRODUCTION

Navy Air Logistics Office (NALO) is tasked with planning and coordinating Navy Unique Fleet Essential Airlift (NUFEA) missions, a category of air cargo and passenger transport missions that fulfills demands from around the world, primarily servicing the U.S. Navy but also the greater Department of Defense (DOD). The execution of NUFEA missions is the responsibility of Fleet Logistics Support Wing (FLSW), the Navy's sole intra-theater airlift capability. FLSW operates a fleet of 41 aircraft of two different types, C-40A and C-130T, based at 11 fixed-wing aircraft transportation squadrons (VRs) that FLSW oversees, as shown in Table 1.

Squadron	Base
C-40	
VR-51	Marine Corps Air Station Kaneohe Bay, Hawaii (PHNG)
VR-56	Naval Air Station Oceana, Virginia (KNTU)
VR-57	Naval Air Station North Island Coronado, California (KNZY)
VR-58	Naval Air Station Jacksonville, Florida (KNIP)
VR-59	Naval Air Station Joint Reserve Base Fort Worth, Texas (KNFW)
VR-61	Naval Air Station Whidbey Island, Washington (KNUW)
C-130	
VR-53	Joint Base Andrews, Maryland (KADW)
VR-54	Naval Air Station Joint Reserve Base New Orleans, Louisiana (KNBG)
VR-55	Naval Air Station Point Magu, California (KNTD)
VR-62	Naval Air Station Jacksonville, Florida (KNIP)
VR-64	Joint Base McGuire-Dix-Lahkehurst, New Jersey (KWRI)

Table 1: FLSW squadrons and corresponding bases.

DOD entities, mainly the U.S. Navy and Marine Corps, submit to NALO a variety of air transportation requests including delivery of maintenance parts, supplies, and personnel transport. The operators of NALO usually handle 40 to 50 requests per week, and group them into NUFEA missions for fulfillment. One mission, which may comprise several legs, could be fulfilling anywhere from 1 to 10 requests, and be completed in 1 to 9 days. When planning a mission, the operators at NALO need to consider the available mission capable aircraft from all of the squadrons; then, decide from which base and with which aircraft type to execute it. This manual process can be burdensome for the planners, further complicated by aircraft availability (scheduled and unscheduled maintenance), and crew scheduling. NALO and FLSW are seeking more efficient ways to manage and employ the current aircraft, including the allocation of aircraft across bases.

Utilizing mixed-integer program (MIP) models, we analyze FLSW's current practices and aircraft deployment to bases with the goal of finding potential avenues for improvement. Our effort comprises three stages: First, given the existing deployment of aircraft in bases, we compare past actual assignments of aircraft to missions to cost-minimized assignments obtained from our *employment* model. Second, we propose an improved deployment of aircraft via our *deployment* model, which represents the strategic decision of interest and is the primary focus of our work. Third, we run the employment model again, based on the improved deployment obtained at the second stage, and compare our results with historical data. In this final step, the operational assignment model serves only as an evaluation tool, allowing us to assess the effectiveness of a redeployment strategy using consistent, cost-minimizing aircraft-to-mission assignments. The results of the first stage therefore primarily serve to provide a baseline for measuring gains from optimized deployment, but have the secondary benefit of indicating potential for FLSW to reduce cost in their current deployment by optimizing mission assignments.

Several related studies have been published in the open literature. Cummins and Wilborn (2009) perform a cost-benefit analysis to determine the composition of the future FLSW fleet and they recommend to continue the C-40A program. Their work helps quantify the impact of acquiring more aircraft. Our study on optimal deployment and employment of FLSW's aircraft addresses budgetary realities that may delay such acquisitions, by focusing on avenues for improving the utilization of existing resources.

Homsi et al. (2021) solve a knapsack assignment model for transporting items in a military context, focusing on operational use. Our work, on the other hand, is motivated by strategic planning: it focuses on the statistical analysis of fleet utilization, transportation cost, and potential for aircraft reallocation. While we do not explicitly enforce daily-dependent aspects of aircraft availability, we use an abundance of historical operational data to incorporate those aspects implicitly in our simulated aircraft availability approach.

Another related military problem is studied in Brown et al. (2013), which develops a decision aid for air cargo planning implemented in operations that took place in Iraq and Afghanistan. Their MIP model maximizes the flow of priority cargo: passengers and air freight pallets. Like Homsi et al. (2021), they too focus on an operational model, rather than one designed specifically for strategic planning.

A relatively large body of research exists addressing related problems in commercial airlines. Barnhart and Cohn (2004) describe a common 4-stage framework that begins with (stage 1) selecting routes, follows with (stage 2) assigning aircraft to those routes, then addresses (stage 3) aircraft maintenance routing and (stage 4) crew scheduling. UPS employs the first 2 stages of this framework in Armacost et al. (2004). They adopt a composite variable formulation that uses column generation to identify the set of decision variables over which to optimize, and these variables account for selecting both routes and aircraft. Garg et al. (2024) proposes an integrated model that addresses all 4 stages, using Benders Decomposition and Lagrangian Relaxation to improve computational tractability. In addition to their technical contributions, the authors provide an up-to-date, comprehensive review on work relating to aircraft and crew schedule planning.

Literature rooted in commercial airlines has a common motivation of emphasizing profitability. Accordingly, methods are not constructed to consider, and potentially comply with both (1) an existing baseline solution for aircraft assignments to a home location (base); and (2) limited ability to modify fleet composition combined with a requirement to satisfy all demand signals. The strategic planning problem we address in this paper is designed to examine possible repositioning of aircraft based on a large data set comprising historical demand figures, in a manner consistent with operational and administrative requirements.

Our statistical approach is unique in capturing the historical mission data in probability distributions for aircraft availability that once sampled can be used to obtain meaningful results through independent daily optimizations. Typically, when a time-dependent operational infeasibility may arise on a particular day, problems cannot not be decoupled and are solved with a multi-day planning model. However, given our models are not intended for daily operations use and are constructed solely for strategic planning, our approach that bypasses time dependence is suited for its military application.

The remainder of the paper is organized as follows. Section 2 presents the data provided to us by NALO, discusses its use in our modeling, and introduces our sampling approaches for generating operational scenarios. Section 3 describes our optimization models. Section 4 presents our computational experiments and discusses the results. Section 5 summarizes our conclusions and proposes future research directions.

2 DATA AND SCENARIO CONSTRUCTION

NALO provided data on all completed transport mission requests from 2015 to 2023. Table 2 presents the mission parameters, extracted from the data, that are relevant for our study. NALO historical operations data capture many, but not all parameters of interest for the objective of minimizing the cost of assigning available aircraft to missions. In particular, NALO data neither provides information about the day when a mission became available to execute nor the available aircraft at each base and on each day. It is not possible to infer whether the grounded status of any aircraft is due to being available and unassigned or out of service for maintenance. To bridge this gap in the data we introduce two estimation approaches, which we refer to as *upper bound* and *lower bound* sampling. The two approaches randomly generate daily aircraft availabilities, which complete the missing data. Both approaches sample probability distributions constructed using historical NALO data to capture empirical mission durations and aircraft utilization patterns. This design allows the simulated operational environment to be statistically consistent with past FLSW activity. Note that we use the term *bound* loosely to indicate a statistical characterization rather than an absolute limit.

Parameter Name	**Description**
Flying Unit	Squadron executing the mission
A/C	Aircraft type assigned to the mission (C-40 or C-130)
Start Date (Z)	Start date of the mission
End Date (Z)	End date of the mission
Itinerary	Flight path of the mission (sequence of airport codes)
Flt Hrs	Total number of flight hours the mission required
Msn Hrs	Total number of mission hours the mission required (flight hours, cargo loading/unloading times, refueling time, crew rest time)

Table 2: Mission data descriptions.

We sample aircraft availability independently across days (in both our sampling approaches) even though the context clearly has intrinsic time dependencies. For example, an aircraft on mission in any given day may remain on mission the following day(s). In our sampling, we do not account for such effects directly. However, the underlying distributions for each approach are estimated from years of historical operations data; therefore, they already incorporate the aggregate effects of multi-day missions, maintenance cycles, and temporal dependence of aircraft utilization. We deem this sufficient for our long-term, strategic deployment decisions, whereas it could be problematic for daily operational planning.

Aside from detail on aircraft availability, the historical mission data does not contain other administrative constraint factors, such as crew qualifications, squadron-specific administrative responsibilities, and runway length. Accordingly, we assume that available aircraft are assignable and able to execute missions. Because NALO does not adjust aircraft deployment on a seasonal or short-term basis, our analysis focuses on long-horizon strategic deployment rather than time-varying repositioning. Accordingly, we draw mission samples from the entire 2015–2023 period. While this approach may smooth seasonal fluctuations, it is consistent with the longer-term strategic time scale of the deployment decisions under consideration.

2.1 Aircraft Availability: Upper-Bound Sampling Approach

Our upper-bound sampling approach (UBSA) approach estimates the number of aircraft available (not grounded or unavailable aircraft) on any given day using historical mission duration data. The term "upper bound" refers to a statistically derived optimistic estimate of available aircraft, based on the distribution of mission lengths obtained from the NALO data. The over 29,000 recorded missions took from 1 to 9 days to complete. The proportion of missions lasting s days can be treated as a probabilistic estimate of mission duration, p_s, and is shown in Table 3.

Mission duration (days), s	1	2	3	4	5	6	7	8	9
Estimated probability, p_s	0.114	0.271	0.120	0.113	0.154	0.115	0.0	0.075	0.038

Table 3: Mission duration estimated probabilities, p_s.

By scaling the p_s values by the corresponding reciprocals of the missions durations (essentially using the rule of total probability), we obtain an estimate p for the probability that an aircraft is not flying a mission:

$$p = \sum_{s=1}^{9} \frac{1}{s} p_s = 0.382. \tag{1}$$

Equation (1) shows that on any given day from 2015 to 2023, on average, approximately 15.6 of the 41 aircraft were not flying missions. NALO have assessed this estimate as reasonable. We assume in UBSA the best-case scenario in which no aircraft are undergoing maintenance, so the above value represents a statistical upper bound on the average available aircraft for new mission assignment on any given day. In theory, we could reduce this number by accounting for grounded aircraft, but such relevant data was not available. Hence, it is treated as an upper bound on the availability of aircraft, which has the potential to translate into optimistic optimization results.

On each simulated day, we determine the number of available aircraft for each base, by sampling from a binomial distribution for each of the 11 squadrons with the following parameters: number of trials, n, equal to the nominal number of aircraft deployed at the base; and, probability of success in each try, $p = 0.382$.

2.2 Aircraft Availability: Lower-Bound Sampling Approach

Our lower-bound sampling approach (LBSA) generates scenarios for aircraft availability on each day of a selected time horizon using triangular distributions. Like the UBSA approach, these distributions are constructed directly from historical NALO data, so the resulting availability samples are statistically informed. For each base and aircraft type, for the current allocation of aircraft to bases, we create a triangular distribution with a minimum of 0, a mode of the average number assigned, and a maximum equal to the maximum number assigned on any day from 2015 to 2023. For example, base KNIP has an average of 0.86 C-40A aircraft assigned to new missions per day and a maximum of 3 C-40A aircraft, so we create a triangular distribution with parameters 0, 0.86, and 3.

We refer to this method of estimation as a lower bound because it uses departing aircraft as a proxy for those actually available, but there may have been other available aircraft that were not used. The historical data do not track maintenance status or unused but mission-capable aircraft, so LBS may underestimate the true availability on a given day. LBSA, however, is not a lower bound in the strict sense because our triangular distribution allows for a sampled number on any day to be equal to the maximum missions recorded. With this method, out of the 41 aircraft, only $\sim$9.5, on average (the sum of the distribution means from all bases) are available for assignment to new missions on any given day – a conservative estimate.

When considering redeployment of aircraft, a method for adjusting the triangular distributions for affected bases and aircraft experiencing redeployment is required. We apply a fractional scaling rule tied to the change in the number of aircraft of each type at each base. If the allocation for type t at base i changes from $n_{i,t}^{\text{current}}$ to $n_{i,t}^{\text{new}}$, we scale the triangular mode and maximum by $\alpha_{i,t} = n_{i,t}^{\text{new}}/n_{i,t}^{\text{current}}$, when $n_{i,t}^{\text{current}} > 0$. When no current allocation exists, we apply a per–aircraft utilization benchmark to define the new triangular distribution mode as $n_{i,t}^{\text{new}} \sum_i \text{mode}_{i,t}^{\text{current}} / \sum_i n_{i,t}^{\text{current}}$ and its maximum parameter as $n_{i,t}^{\text{new}} \sum_i \max_{i,t}^{\text{current}} / \sum_i n_{i,t}^{\text{current}}$. This preserves our assumption for LBSA that aircraft availability is inferred from historical utilization rather than nominal fleet size, even when redeploying aircraft to adjust allocations assigned to some bases.

2.3 Estimated Mission Cost

Optimizing mission-aircraft assignments requires mission cost parameters for alternate routes not included in the NALO data. We estimate these costs by multiplying hourly flight cost of \$9,557 for the C-130T aircraft and \$4,406 for the C-40A aircraft by flight hours required to complete each mission, which are computed as the ratio of distance traveled to average aircraft airspeeds of 328 mph and 570 mph for C-130T and C-40A, respectively. For calculating origin-dependent mission costs, we remove the actual (historical) origin base of a mission route and swap in all 11 bases to produce 11 potential mission costs that differ according to the base of origin. For example, the sequence of airport codes KNTU - KBED - KMKE - KCHS - KNTU is a typical description of a mission, referred to as a *flight path*. We remove KNTU to get the intermediate flight path, i.e., KBED - KMKE - KCHS. Then we start and end the intermediate path mission at all 11 bases and, for each, compute the cost using great circle distance.

We use flight hours rather than mission hours to compute cost because flight time is the primary factor relating to mission cost that varies across bases. Operational overhead, including cargo loading and unloading time, refueling, and crew rest would be incurred regardless of where an aircraft is stationed and therefore adds a base-agnostic offset to each mission time. While such overhead could in principle exhibit minor regional variation, our data do not capture these effects. We expect using mission hours instead of flight hours would not materially affect the relative ranking of bases or the resulting redeployment recommendations, because regional variation of overhead is a small relative to overall mission cost.

3 MATHEMATICAL MODELING

Our overarching goal is to minimize the costs of assigning aircraft to missions. We construct and implement two variants of an optimization model: *employment* model and *deployment* model. The employment model takes a given deployment of aircraft to bases – initially, as it currently is, and later under reallocation scenarios – and optimizes the employment of aircraft to missions. The deployment model proposes a possible improved reallocation of aircraft to bases. Each of these optimization models solves a mission-to-aircraft assignment problem for a single day of NALO operations. The employment model is used solely as an operational evaluation mechanism for comparing alternative strategic aircraft deployments.

We randomly sample N days of mission demands from the NALO data between 2015 and 2023. For each of the N days, we simulate a number of aircraft available a_{it} of type t at base i, using both UBSA and LBSA described above, which are input into separate optimization instances, and compare aggregated results over the entire time horizon. The samples of both new mission demands and available aircraft are all independently generated. However, we preprocess each day's samples to ensure missions required do not exceed aircraft availability, and we carryover missions to the following day if capacity limits are reached. Aside from such capacity issues, we run the models for each day of our planning horizon independently.

Generating aircraft availability as independent samples may produce sequences of days that are not actually possible because time dependence affects aircraft availability. For example, if a base has only a total of 3 aircraft and independent samples on consecutive days specify 2 available aircraft on each day, then multi-day missions could not be feasibly assigned on the first day to both aircraft. However, we can ignore time dependence when running daily optimizations for strategic planning and obtain statistically meaningful results, because the aircraft availability generated by our UBSA and LBSA statistical approaches captures the effects of multi-day missions in their underlying probability distributions.

We assume all available aircraft can be assigned to any mission regardless of the aircraft being a C-40A or a C-130T[a]. We do not consider crew scheduling, and assume prepared crews are available as needed.

3.1 Employment Model

We formulate our employment model for day $n \in \{1, \cdots, N\}$ as in integer program with binary decisions for assigning aircraft to missions.

Index Sets

I	set of bases;
$J(n)$	set of missions on day n (including preprocessed out of day $n-1$ due to capacity limits);
T	set of aircraft types.

Parameters

$a_{it}(n)$	number of aircraft of type t available in base i on day n (sampled by UBSA or LBSA);
c_{itj}	cost of assigning aircraft of type t from base i to mission j.

Decision Variables

$$x_{itj}(n) = \begin{cases} 1, & \text{if an aircraft from base } i \text{ of type } t \text{ is assigned to mission } j \text{ on day } n; \\ 0, & \text{otherwise} \end{cases}$$

Model

$$C(n) = \min \sum_{i \in I} \sum_{t \in T} \sum_{j \in J(n)} c_{itj} x_{itj}(n) \tag{2}$$

$$\text{s.t.} \sum_{j \in J(n)} x_{itj}(n) \leq a_{it}(n) \qquad \forall i \in I, t \in T \tag{3}$$

$$\sum_{t \in T} \sum_{i \in I} x_{itj}(n) = 1 \qquad \forall j \in J(n) \tag{4}$$

$$x_{itj}(n) \in \{0, 1\} \qquad \forall i \in I, t \in T, j \in J(n) \tag{5}$$

The objective function (2) minimizes the total cost of aircraft assigned to missions. Constraint (3) assigns aircraft according to the sampled aircraft availability at each base, under a given force layout allocation that does not vary by day. Constraint (4) ensures that every mission that needs to be executed has exactly one aircraft assigned to it, and requires preprocessing mission carryover to the next day when infeasibility would arise from insufficient aircraft capacity. Specifically, we transfer $\max\{0,\ |J(n)| - \sum_{i \in I} \sum_{j \in J(n)} a_{it}(n)\}$ missions from day n to day $n+1$ to ensure feasibility on the n-th day.

3.2 Deployment Model

The deployment model for day n differs from the employment model only in a single constraint. Instead of (3) (which requires that all missions are fulfilled by aircraft as they are currently deployed), the deployment model allows aircraft to be repositioned as they are needed, by pooling them together and assigning them to missions:

$$\sum_{i \in I} \sum_{j \in J(n)} x_{itj} \leq \sum_{i \in I} a_{it}(n) \qquad \forall t \in T \tag{6}$$

The deployment model improves the assignment of aircraft to missions by allowing the assignment of available aircraft to the bases that are the most cost effective for a given day. While this daily reallocation is unacceptable in practice (because aircraft cannot be shuffled around among bases daily), it is still useful for determining alternate aircraft allocations. We post-process solutions from the deployment model to compute a (base, aircraft type) index called individual ratio (IR). This index reflects the relative average utilization rate of aircraft of type t at base i across all days.

$$\mathrm{IR}_{it} = \frac{\sum_{n=1}^{N} \sum_{j \in J(n)} x_{itj}(n)}{\sum_{n=1}^{N} \sum_{\tilde{i} \in I} \sum_{j \in J(n)} x_{\tilde{i}tj}(n)}.$$

The IR_{it} values are used to redistribute the available aircraft of type t across the bases i (approximately) proportionally to those values (see Macaluso 2024). Specifically, we generate an *unconstrained* allocation that may not be implementable in practice because of existing administrative constraints, but provides a what-if analysis that can be used to assess the cost of compliance with such constraints. To ensure viability, we also generate a *compliant* allocation that adheres to the following NALO constraints: (i) no closure of existing bases; (ii) a squadron must have at least 2 aircraft of the same type, or not exist at all; and (iii) each base is assigned at least one squadron. Additional what-if analyses can be constructed that satisfy some but not all of these requirements, although we limit our computational experiments below to consider only the fully compliant allocation and the fully unconstrained analysis.

4 COMPUTATIONAL EXPERIMENTS, RESULTS, AND ANALYSIS

Our computational experiments are summarized as follows:

1. Generate 30 scenarios of missions, by randomly selecting 100 days from historical

NALO operations.

2. For each scenario and each day, generate aircraft availability using UBSA and LBSA.

3. Solve the deployment model, which ignores the current aircraft allocation, and produce IR_{it} values that provide a statistical basis for an improved aircraft allocation.

4. Apply a rounding heuristic to IR_{it} values to obtain a non-fractional, unconstrained aircraft allocation.

5. Construct a second allocation, focusing on compliance, based on the IR_{it} values, but enforcing NALO's administrative constraints.

6. Compute the cost of the historical employment decisions to obtain a basis C_1 for comparison[b].

7. For each scenario, solve the employment model using the current, compliant, and unconstrained aircraft allocations.

8. For both UBSA and LBSA, compare historical performance against improved employment, using averages of associated operational costs over 30 scenarios. We shall denote by C_2, C_3, and C_4 the corresponding cost under current, compliant and unconstrained allocations, respectively.

In the last step, we note C_2 represents an efficient assignment of available aircraft sampled according to UBSA or LBSA under current allocation, so we expect it to be an improvement over C_1. C_3 benefits from a compliant redeployment strategy, which should be a cost improvement compared to C_2. C_4 allows an unconstrained redeployment strategy, which should be a further cost decrease compared to C_3. These expectations hold true for both the UBSA and LBSA results.

4.1 Comparison of Sampling Approaches

We present the unconstrained and compliant allocations based on the deployment model solutions using UBSA and LBSA. Table 4 compares the costs C_1 to C_4 for both sampling approaches. Across the 30 simulated scenarios, the relative variability (standard deviation divided by the mean) is consistently around 5% for all cases, indicating that the resulting cost estimates are statistically stable despite the inherent randomness of 100-day samples drawn from a multiyear demand history. As expected, optimized employment performs significantly better than historical employment, reducing cost by around 35-40%. Compliant and unconstrained deployment reallocation further reduce cost by 5% and 14%, and another 3% and 7%, in UBSA and LBSA respectively.

Cost Metric	**UBSA**	**LBSA**
C_1: Current Allocation and Employment	\$105M	\$105M
C_2: Current Allocation with Optimized Employment	\$64M (5.7%)	\$69M (5.6%)
C_3: Deployment Reallocation (Compliant)	\$61M (4.9%)	\$59M (5.8%)
C_4: Deployment Reallocation (Unconstrained)	\$59M (5.0%)	\$55M (4.9%)

Table 4: Cost comparison under UBSA and LBSA averaged on 30 scenarios of 100 days each with relative variability in parentheses.

An unexpected result, given the higher average availability of aircraft in UBSA compared to LBSA, is the lower cost achieved by LBSA compared to UBSA under both compliant and unconstrained reallocation. This can be explained by analyzing aircraft availability. First, aircraft availability for UBSA is roughly proportional to the current allocation, yielding (on average) approximately 6 out of 17 C-40A and 9 out of 24 C-130T available daily. However, historical assignments (35% for C-40A, and 13% for C-130T, on average) drive the triangular distribution parameters in LBSA, which uses effective historical utilization as a proxy for actual capacity. In fact, while the triangular distribution in LBSA generates significantly fewer total aircraft than the binomial in UBSA, it may actually yield higher C-40A availability because of the historical usage imbalance between C-40A and C-130T combined with the proportionality embedded in UBSA, thereby decreasing average assignment cost.

Scenario	UBSA Carryover			LBSA Carryover		
	Avg	Max	End	Avg	Max	End
Current	0.03	2.17	0.00	1.36	9.67	1.60
Unconstrained	0.03	2.07	0.00	1.10	9.07	0.90
Compliant	0.03	2.03	0.00	0.97	8.23	1.03

Table 5: Carryover comparison for UBSA and LBSA averaged on 30 scenarios of 100 days each.

The impact of our conservative assumption on aircraft availability in LBSA is evidenced when tracking aircraft shortages. We present a comparison of average, maximum and end carryovers in Table 5. Average carryover reports the daily number of unassigned missions, but an average of 1 does not mean a new shortage arises each day. Rather, it indicates that when a backlog arises it may propagate for many days until there is an opportunity to catch up on mission fulfillment. Maximum carryover tracks how large the backlog queue of missions grows over a 100-day simulation. End carryover measures the size of the backlog queue on completion of day 100. LBSA has on average of $\sim$1 carryover per day, including on day 100, in all 30 scenarios generated, with more than 8 carryovers on some days. UBSA almost always has sufficient aircraft to complete all missions assigned each day, with carryovers rarely exceeding two missions on any day during the entire time horizon, and ending with all missions complete on day 100. Some LBSA carryover at day 100 is due to horizon truncation. Extending the horizon would allow all missions to be eventually executed. Carryover remained small and stable across replications, with no evidence of accumulation over the 100-day horizon under either sampling method.

Table 6 compares the 5 different allocations: the current allocation, and for both UBSA and LBSA the optimized compliant and unconstrained redeployments of aircraft. The numbers in parenthesis indicate the change in the number of aircraft suggested by the new allocations compared to the current allocation. The recommended redeployment of aircraft is quite consistent both between UBSA and LBSA, under both compliant and unconstrained optimization, reflecting that bases with large fleet concentration areas attract additional aircraft under both sampling approaches.

Base and Aircraft Type	Current	Compliant		Unconstrained	
		UBSA	LBSA	UBSA	LBSA
KNTD C-40A	0	2 (+2)	2 (+2)	2 (+2)	2 (+2)
KNTD C-130T	5	5	6 (+1)	7 (+2)	8 (+3)
KNBG C-40A	0	0	0	1 (+1)	1 (+1)
KNBG C-130T	4	2 (-2)	2 (-2)	1 (-3)	1 (-3)
KNFW C-40A	3	2 (-1)	2 (-1)	1 (-2)	1 (-2)
KNFW C-130T	0	2 (+2)	2 (+2)	2 (+2)	2 (+2)
KNIP C-40A	3	0 (-3)	0 (-3)	0 (-3)	0 (-3)
KNIP C-130T	5	2 (-3)	2 (-3)	2 (-3)	2 (-3)
KADW C-40A	0	0	0	0	0
KADW C-130T	5	2 (-3)	2 (-3)	0 (-5)	0 (-5)
KWRI C-40A	0	0	0	1 (+1)	1 (+1)
KWRI C-130T	5	2 (-3)	2 (-3)	0 (-5)	0 (-5)
KNUW C-40A	3	2 (-1)	2 (-1)	2 (-1)	2 (-1)
KNUW C-130T	0	2 (+2)	2 (+2)	3 (+3)	3 (+3)
KNZY C-40A	3	4 (+1)	4 (+1)	4 (+1)	4 (+1)
KNZY C-130T	0	4 (+4)	3 (+3)	5 (+5)	4 (+4)
KNTU C-40A	3	5 (+2)	5 (+2)	6 (+3)	6 (+3)
KNTU C-130T	0	3 (+3)	3 (+3)	4 (+4)	4 (+4)
PHNG C-40A	2	2	2	0 (-2)	0 (-2)
PHNG C-130T	0	0	0	0	0

Table 6: Comparison of current, compliant, and unconstrained aircraft allocation under UBSA and LBSA.

Strategic Insights: If NALO wishes to move forward with a reallocation plan, the constraints they impose on the reallocation have an important impact on the resulting cost reduction. A reallocation plan takes significant time, money, and logistical planning to execute. This could likely lead to only moving a small number of aircraft rather than an overhaul of the entire 41 aircraft fleet that we analyze using the new unconstrained and compliant allocations under both sampling method approaches.

The model suggests that increasing the number of C-40As at bases KNTD, KNZY, or KNTU, and increasing the number of C-130Ts at bases KNZY or KNTU would yield the largest reduction in cost. When reallocating aircraft, we observe C-40As would most effectively be moved out of base KNIP, and C-130Ts be moved out of bases KADW or KWRI.

Sampling Approach Takeaways: UBSA and LBSA both provide insight into the potential cost benefit of optimized employment, and optimized redeployed aircraft allocations. However, a few related observations lead us to favor UBSA. First, applying LBSA to reallocation of aircraft is not straightforward, because the historical data cannot be directly applied to estimate the modes and maxima of triangular distributions, unlike UBSA where the binomial distribution probability parameter remains applicable for scaling. Second, the level of C-130T availability generated by LBSA is unreastically low, due to the conservative assumptions underlying the construction of the triangular distributions. Third, LBSA has significant carryover, due to its limited aircraft availability. Accordingly, we proceed with only UBSA in our sensitivity analysis.

4.2 Sensitivity Analysis

We examine the impact of decreasing and increasing aircraft supply. We extend our computational study by running comparative experiments: reducing the C-130Ts by 1 to 15 aircraft, reducing C-40s by 1 to 17 aircraft, and increasing C-40s by 1 to 15 aircraft. Given the already low utilization of C-130Ts and their higher cost, we do not increase their presence in the fleet. Each reduction or increase is executed in isolation, with the same setup as our initial computational experiments, with 30 scenarios generated for 100 days each. We compare results for compliant and unconstrained redeployments.

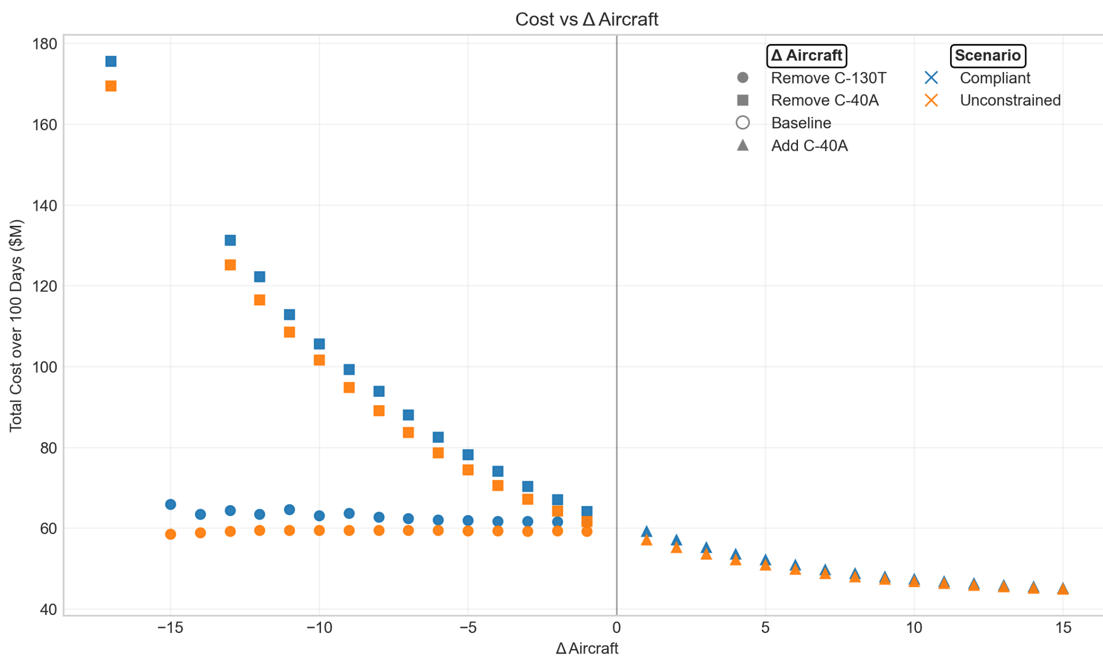

Figure 1: Total cost vs. Δ aircraft (compliant vs. unconstrained).

Figure 1 shows the total cost of missions for 100 days, averaged over 30 scenarios, for varying fleet compositions. We measure the flight-hour cost of executed missions. When aircraft are removed, the assignment of missions to aircraft becomes less flexible, so cost increases, and vice versa. Removing C-130Ts has little impact, because their historical assignment rate is roughly one-third that of C-40As. As long as a few remain available, missions can be satisfied. This is not true for C-40As, whose removal increases cost significantly and whose addition decreases it. Note: we do not model changes in the distribution of mission execution time that could arise from fleet composition changes; however, the underlying binomial parameter derived from Table 3 would be affected by such changes.

Figures 2 and 3 show the average and maximum carryover for 100 days, averaged over 30 scenarios, for varying fleet compositions. Adding aircraft has no impact on average carryover because that was already near 0 in the baseline, but maximum carryover noticeably decreases

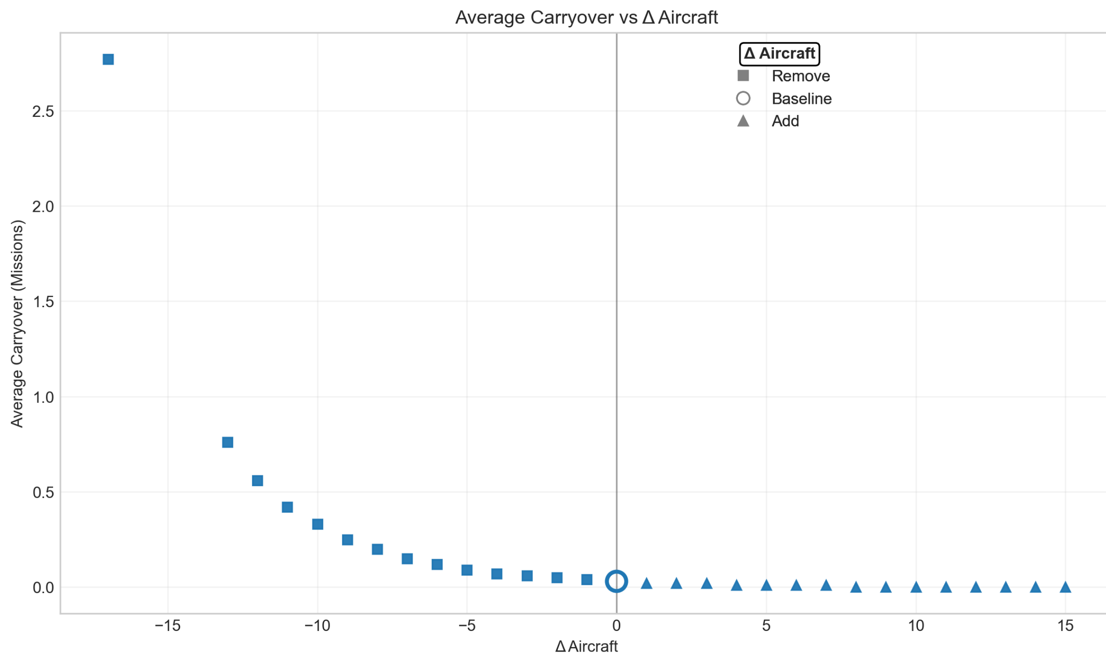

Figure 2: Average carryovers vs. Δ aircraft.

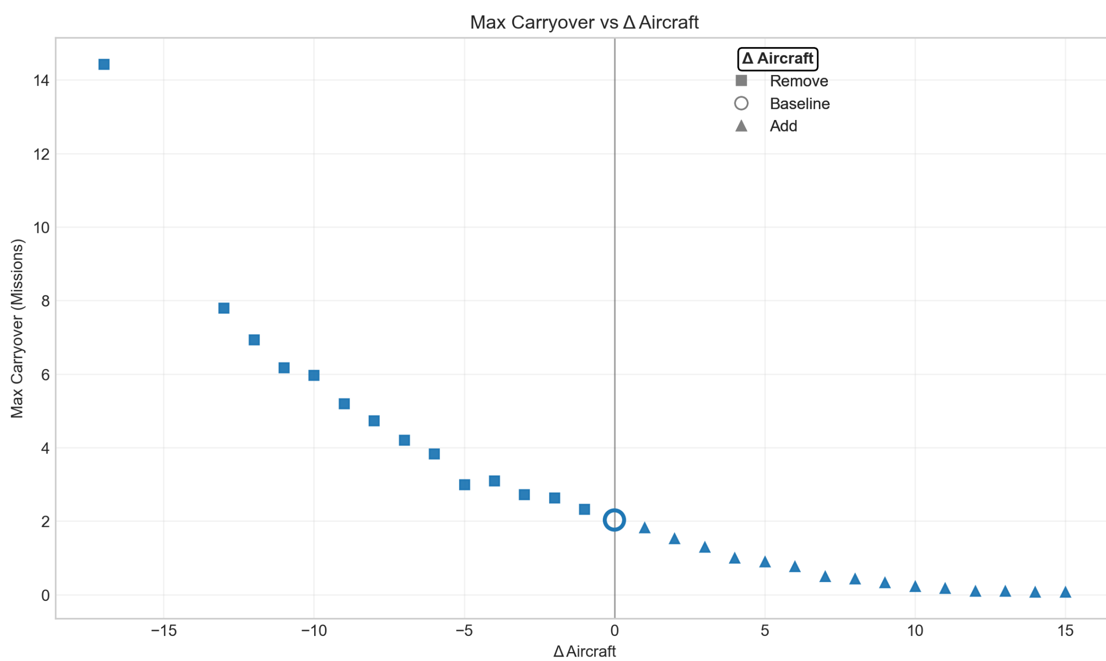

Figure 3: Maximum carryovers vs. Δ aircraft.

until ∼9 aircraft are added. Removing aircraft, independent of type, increases both average carryover and maximum carryover, as expected.

5 CONCLUSION

This work provides a framework to improve NALO's use of C-40A and C-130T aircraft that fulfill NUFEA missions, by exploring potential reallocation of aircraft across bases. We implement two variations of a MIP whose objective is minimizing cost through assignment of available aircraft to missions on a given day. The employment model uses the aircraft given a fixed allocation to bases, initially the current allocation, whereas the deployment model flexibly pools and reallocates the aircraft to bases. Deployment reallocations are then aggregated across the time horizon, by base and aircraft type, to guide two types of reallocation plans: one is compliant with administrative constraints, and the other is unconstrained. These new allocations are subjected to simulation runs of the employment model for comparison against the current allocation.

NALO's historical data does not capture aircraft availability, so we can only infer bounds based on tracked usage. To address this data gap, we introduce two statistical sampling approaches: UBSA uses a binomial distribution to sample daily available aircraft, with number of trials set to the number of aircraft of each type stationed at each base, and the probability of success set to $p \sim 0.382$, the historical daily proportion of aircraft not flying a mission, equating to a daily average availability of 15.6 aircraft. LBSA uses a triangular distribution with a minimum of 0, and a mode and maximum equal to the average and maximum historical numbers, respectively, of each aircraft type at each base on a given day, equating to daily average availability of 9.5 aircraft. The triangular distribution parameters are adjusted when considering compliant and unconstrained reallocations, while preserving approximately the same 9.5 daily average availability.

We simulate 100 days of missions, and replicate 30 times for robustness. The cost of the mission assignments is \$105 million under the current allocation. We apply both UBSA and LBSA, and find similar C-40A availability ($\sim$3 on average per day) under both approaches yields similar cost estimates, even though total aircraft availability is significantly higher under UBSA. Employment optimization improves total cost by $\sim$35%, and deployment reallocation demonstrates potential for an additional $\sim$5%-10% improvement compared to the historical baseline cost, all told on the order of a hundred million dollars annually.

We recommend that if NALO moves forward with a reallocation of the aircraft in their current fleet, they should station more aircraft of either type at the bases closer to the higher U.S. Navy fleet concentration areas. These bases include, for the C-40A, KNTU, KNZY, and KNTD; and for the C-130T, KNZY or KNTU. When deciding from which bases to reallocate C-40A and C-130Ts, we recommend they reallocate C-40As from KNIP, and C-130Ts from either KADW or KWRI.

In future work, our models can be extended to include factors like crew availabilities, scheduled maintenance cycles, cargo types (passengers or pallets of cargo), and cargo capacities for both the C-40A and C-130T. The deployment model could be expanded to include administrative constraints for the reallocation of the aircraft. Furthermore, it could be extended to a multi-day model in order to directly optimize deployment, including potential expansion of the fleet, and handling carryovers. The employment model could be refined to deliver a real-time scheduling tool for NALO.

NOTES

[a]These aircraft do not have identical capabilities. It would be straightforward to restrict certain missions from being flown on a particular aircraft type, but historical data does not capture such information. NALO verified that most missions can be carried out on either type, so this assumption is acceptable and does not materially affect our strategic cost comparisons.

[b]The cost depends only on the historical data, because both sampling methods use the same 30 mission sets as input, ensuring a paired and directly comparable baseline.

ACKNOWLEDGMENTS

This research was sponsored by NALO through a grant from the Naval Research Program. The thoughtful feedback we received in the peer-review process led to significant improvements, and we are grateful to the editors and reviewers for their contributions.

AUTHOR STATEMENT

REFERENCES

Armacost, Andrew P. et al. (2004). "UPS optimizes its air network". In: *Interfaces* 34.1, pp. 15–25. DOI: `10.1287/inte.1030.0060`.

Barnhart, Cynthia and Amy Cohn (2004). "Airline schedule planning: Accomplishments and opportunities". In: *Manufacturing & Service Operations Management* 6.1, pp. 3–22. DOI: `10.1287/msom.1030.0018`.

Brown, Gerald G. et al. (2013). "Optimizing intratheater military airlift in Iraq and Afghanistan". In: *Military Operations Research* 18.3, pp. 35–52. ISSN: 10825983, 21632758. URL: `http://www.jstor.org/stable/24838479` (visited on 05/20/2025).

Cummins, Lane and Tony Wilborn (2009). *Cost-benefit analysis of the Department of the Navy's transition from C-9 aircraft to C-40 aircraft for logistic support aircraft*. URL: `https://hdl.handle.net/10945/10395` (visited on 05/20/2025).

Garg, Ankur et al. (2024). “Integrated commercial and operations planning model for schedule design, aircraft rotation and crew scheduling in airlines”. In: *Networks* 83.4, pp. 653–672. DOI: `10.1002/net.22211`.

Homsi, Gabriel et al. (2021). “The assignment and loading transportation problem”. In: *European Journal of Operational Research* 289.3, pp. 999–1007. ISSN: 0377-2217. DOI: `10.1016/j.ejor.2019.07.039`.

Macaluso, Giovanni (2024). *Optimal allocation of aircraft to support fleet essential airlift.* URL: `https://hdl.handle.net/10945/73173` (visited on 09/02/2025).

ACRONYMS

DOD Department of Defense. 2

FLSW Fleet Logistics Support Wing. 1–3, 6

IR individual ratio. 13

KADW Joint Base Andrews, Maryland. 2, 17, 18, 22

KNBG Naval Air Station Joint Reserve Base New Orleans, Louisiana. 2, 17

KNFW Naval Air Station Joint Reserve Base Fort Worth, Texas. 2, 17

KNIP Naval Air Station Jacksonville, Florida. 2, 8, 17, 18, 22

KNTD Naval Air Station Point Magu, California. 2, 17, 18, 22

KNTU Naval Air Station Oceana, Virginia. 2, 9, 17, 18, 22

KNUW Naval Air Station Whidbey Island, Washington. 2, 17

KNZY Naval Air Station North Island Coronado, California. 2, 17, 18, 22

KWRI Joint Base McGuire-Dix-Lahkehurst, New Jersey. 2, 17, 18, 22

LBSA lower-bound sampling approach. 8–11, 14–18, 21, 22

MIP mixed-integer program. 3, 4, 21

NALO Navy Air Logistics Office. 1–3, 5, 7, 9, 10, 13, 14, 18, 21–23

NUFEA Navy Unique Fleet Essential Airlift. 2, 21

PHNG Marine Corps Air Station Kaneohe Bay, Hawaii. 2, 17

UBSA upper-bound sampling approach. 7, 8, 10, 11, 14–18, 21, 22

VR fixed-wing aircraft transportation squadron. 2