



# Scheduling United States Coast Guard helicopter deployment and maintenance at Clearwater Air Station, Florida

R.A. Hahn<sup>a</sup>, Alexandra M. Newman<sup>b,\*</sup>

<sup>a</sup>Office of Aeronautical Engineering (CG-41), United States Coast Guard Headquarters, 2100 2nd Street S.W., Washington, DC 20593, USA

<sup>b</sup>Division of Economics and Business, Colorado School of Mines, 1500 Illinois Street, Golden, CO 80401, USA

## Abstract

To mitigate the effects of a corrosive operating environment, the Coast Guard has planned an extensive preventative maintenance program for its Sikorsky HH60J helicopters based on helicopter flight time. We construct a mixed integer linear program that schedules the weeks during which each helicopter undergoes maintenance, as well as the weeks during which a helicopter conducts operations either at Clearwater Air Station, Florida or at one of two deployment sites. The schedules must consider different maintenance types, maintenance capacity and various operational requirements, e.g., the number of helicopters simultaneously patrolling a deployment site. Using data from the operations at Clearwater, we generate optimal schedules on a Unix workstation with modest computing power in less than 3 min. Planners have been using our mixed integer linear program since January, 2005 for scheduling guidance. © 2006 Elsevier Ltd. All rights reserved.

*Keywords:* Maintenance scheduling; Deployment scheduling; Military operations; United States Coast Guard; Integer programming applications

## 1. Introduction and background

The United States Coast Guard (USCG), which falls under the Department of Homeland Security, maintains three primary types of assets: cutters, i.e., vessels 65 ft or greater in length, boats, i.e., vessels less than 65 ft in length, and both fixed- and rotary-wing aircraft. There are over 200 aircraft based at 27 different air stations in the continental United States, Alaska, Hawaii and Puerto Rico. The largest air station is the Coast Guard Air Station (CGAS) in Clearwater, FL which currently maintains nine Sikorsky HH60J helicopters (see Fig. 1) and six C130 fixed-wing aircraft. The Coast Guard uses the HH60J helicopter for its medium-range missions up to 300 nautical miles offshore.

The search and rescue mission is one of the Coast Guard's oldest missions, and has a goal to "minimize the loss of life, injury, property damage or loss by rendering aid to persons in distress and property in the maritime environment" [1]. In 2003, the Coast Guard responded to over 30,000 incidents and saved over 5000 lives in search and rescue missions. Two other Coast Guard missions are particularly important at Clearwater—drug interdiction and alien migration interdiction. In 2004, the Coast Guard seized nearly 26,000 pounds of marijuana and over 240,000 pounds of cocaine with a combined worth of nearly \$8 billion, and interdicted nearly 11,000 illegal migrants [1].

\* Corresponding author. Tel.: +1 303 273 3688; fax: +1 303 273 3416.

E-mail addresses: [rhahn@comdt.uscg.mil](mailto:rhahn@comdt.uscg.mil) (R.A. Hahn), [newman@mines.edu](mailto:newman@mines.edu) (A.M. Newman).

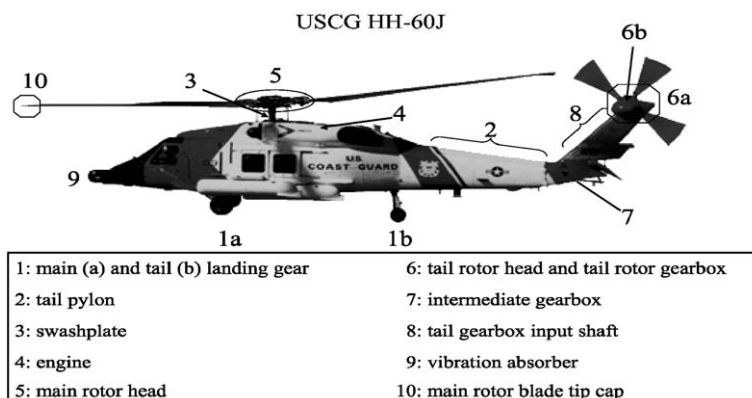


Fig. 1. A schematic of the Coast Guard's Sikorsky HH60J. The 10 parts labeled in the figure are the focus of the Coast Guard's heavy maintenance program for these aircraft. Part 2, the tail pylon, contains the tail rotor shaft. Parts 4 and 6–8, the engine, the tail rotor head and gearbox, the intermediate gearbox and the tail gearbox input shaft, respectively, are covered for protection by a cowling. There are two vibration absorbers (part 9), one in the nose of the aircraft and one in its cabin.

CGAS in Clearwater conducts these missions with aircraft stationed at Clearwater, the home base, and with aircraft stationed at several deployment sites in the Caribbean, i.e., Great Inagua and AUTEK, Andros Island. Aircraft at the deployment sites fly fixed-length patrols each day, whereas flying time at Clearwater varies daily. Because the Coast Guard's missions require that its personnel conduct extensive training exercises and operations using helicopters hovering at low altitudes over salt water and sometimes near sand, these aircraft participate in CGAS, Clearwater's HH60J heavy maintenance program; this program mitigates the effects of the corrosive and abrasive environment and requires that after a certain number of hours of flight time, an aircraft return to (if it is not already flying in) Clearwater for maintenance.

Heavy maintenance is an intensive, intrusive preventative maintenance regimen for the HH60J helicopters that is based primarily on flight hours accrued by each aircraft. The HH60J heavy maintenance division in Clearwater, staffed by approximately 30 specially-assigned Coast Guard mechanics, conducts some combination of four different inspection types on each helicopter after every 200 h of flying time. Each inspection type, i.e., a 200-, 400-, 600- and 800-h inspection type, possesses varying degrees of intrusiveness and entails inspection of different parts of the airframe (see Fig. 1). A 200-h heavy maintenance inspection type addresses the landing gear and wheels and the tail pylon. Additionally, mechanics check the swashplate, a component of the main rotor head that allows directional control, and the chip detection system; this electrical system is located in the engine, in the main gearbox, located directly below the main rotor head, and in the tail rotor gearbox, and detects metallic particles in the lubrication system—a possible indication of premature failure of the engine, for example. A 400-h heavy maintenance inspection type requires inspection of the engine (covered by a cowling in Fig. 1), its fuel filter, and its oil filter, inspection of the main rotor head and further inspection of the swashplate. Mechanics also check the tail rotor system, which consists of the tail rotor head, tail rotor gearbox, intermediate gearbox, tail gearbox input shaft and tail rotor blades. A 600-h heavy maintenance inspection type requires that vibration absorbers, the main rotor head gearbox (underneath the swashplate) and the main rotor blade tip caps be checked. Finally, the 800-h heavy maintenance inspection type requires extensive disassembly and inspection of the main rotor head. There are 12 inspection combinations of varying durations that are conducted in sequence over a 2400-h cycle. Table 1 shows which inspection types constitute each of the 12 inspection combinations, and the duration of each inspection combination.

When a heavy maintenance inspection combination on a helicopter comes due, the helicopter can no longer be flown, i.e., it is “grounded,” until the inspection combination is completed. If maintenance capacity is tight, aircraft queues form at the maintenance facility, which diminishes the number of operational aircraft and impacts Clearwater's ability to conduct its missions. Heavy maintenance inspection schedules must therefore consider: (i) the inspection combinations that a helicopter must sequentially undergo during the 2400-h maintenance cycle and (ii) the number of required operational aircraft at each deployment location and at Clearwater. Because there is no obvious way to synchronize required flight time and maintenance for all aircraft across both deployment sites and at the home base,

Table 1  
Inspection combination characteristics

Inspection combination	Corresponding inspection types				Duration of inspection combination (weeks)
1	200				1
2	200	400			2
3	200		600		2
4	200	400		800	3
5	200				1
6	200	400	600		2
7	200				1
8	200	400		800	3
9	200		600		2
10	200	400			2
11	200				1
12	200	400	600	800	3

Each inspection combination consists of between one and four inspection types of varying intrusiveness, and lasts between one and three weeks.

manual efforts are time-consuming and result in inconsistent, and sometimes undesirable, solutions. We develop a mixed integer linear programming model to provide a deployment and maintenance schedule over a 12-week horizon to ensure that helicopters are operationally available subject to their maintenance requirements.

Helicopters undergo inspections other than those in the heavy maintenance program. These other inspections are either (i) unplanned or (ii) non-intrusive inspections of helicopter components. In either case, these other maintenance activities do not affect heavy maintenance inspection capacity, as they are conducted utilizing dedicated ground crew maintenance personnel that work around the clock. In this paper, we address only heavy maintenance inspections, and specifically, those at CGAS, Clearwater.

In the following section, we review related literature, and in Section 3, we explain manual scheduling methods and introduce our optimization model. In Section 4, we discuss details designed to increase model tractability. Section 5 describes the model's implementation at CGAS, Clearwater and the Coast Guard's evaluation of the automatically-generated schedules. Section 6 presents numerical results, and we conclude with a summary and suggestions for future research.

## 2. Literature review

There exist both civilian and military applications of optimization models for maintenance scheduling of ships and aircraft, though we restrict our literature review primarily to aircraft maintenance scheduling. In a commercial application for Air Canada, Boere [2] presents an optimization model to minimize the number of flying hours lost due to premature maintenance checks. Principle constraints include maintenance facility capacity and worker availability, and the model considers multiple aircraft types. Because of the size of the model and the computing power at the time, the author uses simulation to generate maintenance schedules. The model does not consider commercial routing of aircraft. Martin et al. [3] address the problem of scheduling fractionally owned aircraft for commercial use while considering, inter alia, maintenance requirements, e.g., the number of hours flown since an aircraft was last maintained. Their model minimizes operating costs (including repositioning and outside contracting when an aircraft is not available) while meeting routing (demand) and downtime (crew rest) restrictions. Maintenance constraints require that neither the allowable number of flight hours, nor the allowable amount of elapsed time, nor the allowable number of take-offs, landings, engine starts, etc. is exceeded between a maintenance. Because of the many routing and crew pairing considerations, the model does not include detailed maintenance scheduling. Cohn and Barnhart [4] simultaneously solve the problem of pairing crews to flights and assigning aircraft to flights, where the latter problem only makes the assignment if the appropriate maintenance can be performed during an aircraft's flight schedule. Similar to Martin et al., their model size is extremely large to account for all the crew pairings and routings, and, therefore, the authors do not schedule maintenance in detail.

Military models are, in general, less concerned with the complex routing and crew pairing considerations that occur in commercial applications. Brown et al. [5] formulate a model to assign Navy ships, submarines and aircraft to

schedules consisting of naval exercises and maintenance. Their generalized set partitioning model minimizes penalties for assigning too many or too few combatants to exercises and/or maintenance. Ayik [6] selects deployment and maintenance schedules for Navy ships to minimize both gaps in coverage at a deployment site as well as the number of weeks of early and late deployments. Constraints include shipyard capacity and bounds on the allowable amount of time a ship may deploy early or late. Albright [7] groups locally performed maintenance tasks for Navy helicopters to reduce maintenance costs and eliminate some aircraft downtime associated with flying to a central maintenance depot. However, this model does not consider scheduling individual aircraft for maintenance. Meeks [8] and Baker [9] both schedule Navy aircraft for depot-level maintenance and other major overhauls while minimizing the out-of-service time for each aircraft subject to constraints on the amount and duration of various maintenance types. These models consider maintenance in detail, but do not address aircraft deployment schedules.

Deployment and maintenance scheduling models for the Coast Guard also exist. Sibre [10] and Brown et al. [11] both schedule Coast Guard cutters (i.e., ships). Sibre poses a fixed charge quadratic assignment model to schedule missions for Coast Guard cutters while considering the order in which missions are scheduled, the physical limitations of the cutters and mission requirements. The model seeks to minimize penalties for undesirable actions such as long transitions between missions. He solves the model for sample problems containing, e.g., two ships and a seven-week horizon. Maintenance time is accounted for in the schedule, but no detailed sequencing of maintenance types is done nor is maintenance capacity addressed. Brown et al. consider deployment needs and operational policies, equitable distribution of cutters to deployments and required maintenance time. To factor in unplanned events, the authors generate alternate schedules that differ as little as possible from current schedules. The model covers a 12-week horizon and contains about the same number of assets as ours; however, the only requirement for maintenance scheduling is that there exists enough time during the planning horizon to perform maintenance. In other words, like Sibre, their model does not consider detailed maintenance scheduling.

### 3. Mixed integer linear programming scheduling model

Prior to 2005, CGAS, Clearwater planners were using a spreadsheet to manually schedule heavy maintenance inspections and deployments. The system provided a visual presentation of the schedule and flight time accrued on each HH60J based on each helicopter's operational activities and maintenance status. Manually generating three-month schedules required 4–8 h. And, because schedules must be updated 5–10 times over the course of the planning horizon due to unplanned failures, planners spent up to nearly two weeks manually generating schedules over a 12-week planning horizon. In addition, because of the number of goals and the subjectivity involved in determining the importance of these goals, different planners generated qualitatively different schedules, and no planner could consider all appealing attributes of a schedule during its generation.

In the spring of 2004, CGAS, Clearwater requested assistance from the Coast Guard Aircraft Repair and Supply Center to provide an alternative to CGAS, Clearwater's heavy maintenance and deployment scheduling system. In turn, Aircraft Repair and Supply Center planners solicited current Coast Guard graduate school students to work with CGAS, Clearwater planners to develop an appropriate optimization model.

Our maintenance scheduling model requires primarily that we satisfy a large set of constraints that negotiates between detailed operational and maintenance requirements. We categorize the constraints as follows: (i) general operational constraints, (ii) operational constraints specific to Clearwater, the home base, (iii) maintenance scheduling constraints and (iv) counting constraints.

Operational constraints ensure that helicopters are logically deployed to meet mission requirements. That is, the requisite number of helicopters must be flying at the home base and at the several deployment sites each week. Beyond this, we seek a schedule that: (i) precludes aircraft from being unnecessarily repositioned between deployment sites and the home base between inspection combinations, (ii) prevents a deployed aircraft from returning to Clearwater until it has fewer than a certain number of flight time hours before its next inspection combination and (iii) promotes an equitable distribution of flight hours among all helicopters over the planning horizon. These constraints make the schedule understandable and easy to execute, and prevent any single aircraft from being overtaxed.

Operational constraints specific to Clearwater address the uniqueness of the way in which operations are conducted at the home base as opposed to at the deployment sites. A helicopter assigned to Clearwater accrues a variable number of flight hours due to the sporadic nature of operations there. Therefore, we enforce a minimum and a maximum number of hours that must be flown in Clearwater each week, and we place lower and upper bounds on the number of flight

hours accrued per week on each helicopter. We also place a lower bound on all flight hours in Clearwater over the model's time horizon with a view to meeting a yearly flight time goal. The air station is funded by hours flown, so it is imperative to ensure that during each time horizon over the course of a year, Clearwater is on track to meet this target. By contrast, flight time of helicopters on deployment, i.e., those with a home base of Clearwater but stationed at AUTECH or Great Inagua, is predictable and therefore fixed for each week of the planning horizon because these aircraft fly specific patrols each day.

Maintenance scheduling constraints address both the capacity at the maintenance facility and the time at which a helicopter comes due for maintenance (and can no longer be flown). Maintenance capacity allows only two helicopters to undergo a heavy maintenance inspection combination during any week. To prevent more inspection combinations than necessary, a helicopter is not scheduled for maintenance until it has no more flight time available; then it is forced to begin maintenance. We include constraints to mitigate end effects, i.e., an end to the schedule with many helicopters due for maintenance just beyond the planning horizon.

Counting constraints track the amount of time between inspection combinations. There are 12 different inspection combinations, and an inspection combination must be conducted every 200 h. Therefore, each inspection combination occurs every 2400 h, and the time between inspection combinations is decremented according to a helicopter's flight time.

Forcing a helicopter to remain at the same deployment site until it returns to Clearwater (shortly before maintenance) and requiring a hard minimum number of hours a helicopter must fly in Clearwater each week are the most difficult constraints to satisfy. Therefore, we construct an objective function to minimize the use of elastic penalty variables associated with a helicopter leaving a deployment location, and with violating constraints on the desired number of hours a helicopter flies in a given week at Clearwater.

The time fidelity of the 12-week schedule is weekly. Initial conditions account for the following three pieces of information: (i) the number of hours at the beginning of the horizon until a particular helicopter is due for a certain inspection combination, (ii) the location at which a helicopter is stationed at the start of the horizon and (iii) for each helicopter undergoing maintenance at the beginning of the horizon, whether or not the helicopter was undergoing maintenance in any contiguous weeks prior to the start of the horizon. The parameters in all three cases replace initial variable values. In (i) and (ii), values for only one period prior to the horizon are necessary, while for (iii), for each helicopter undergoing a particular inspection combination lasting  $t$  weeks, we require  $t - 1$  weeks of historical information to provide the number of weeks a helicopter must remain in maintenance after the first week of the horizon. We obtain all initial condition parameters from the actual schedule followed during the prior 12-week horizon. For ease of exposition, we omit these initial condition parameters, and the corresponding constraint exceptions, in the formulation below.

The associated notation and formulation follow:

*Indices:*

- $h$  helicopter,
- $i$  inspection combination,
- $l$  location ( $o$  = home base),
- $t$  time period.

*Sets:*

- $\mathcal{H}$  set of all helicopters (i.e., those with Clearwater as their home base),
- $\mathcal{I}$  set of heavy maintenance inspection combinations,
- $\mathcal{L}$  set of helicopter operational locations.

*Parameters:*

- $S_{it't}$  1 if inspection combination  $i$  starts in week  $t'$  and is ongoing in week  $t$ , 0 otherwise,
- $T$  number of weeks in the planning horizon,
- $\underline{R}_l, \bar{R}_l$  required minimum and maximum number of helicopters at location  $l$ ,
- $H$  2400 h, i.e., number of hours between a particular inspection combination,
- $U$  number of hours before an inspection combination below which an aircraft may not be deployed,

- $\underline{B}, \bar{B}$  minimum and maximum total number of hours a helicopter can fly over the horizon,
- $\underline{M}, \bar{M}$  minimum and maximum number of hours a helicopter can fly while at Clearwater each week,
- $\underline{G}, \bar{G}$  minimum and maximum number of flight hours flown each week at Clearwater,
- $\hat{G}$  minimum number of hours flown by all helicopters over the planning horizon,
- $\hat{M}$  target number of weekly flight hours for each helicopter stationed at Clearwater,
- $F_l$  number of flight hours required in a week of a helicopter that is deployed at site  $l \neq o$ ,
- $\hat{B}_i$  number of accumulated flight hours below which a helicopter must enter inspection combination  $i$  at the end of the horizon to avoid end effects,
- $\tilde{G}$  upper bound on the number of flight hours on each helicopter in Clearwater during the last week of the horizon,
- $\alpha$  scale factor for the penalty related to a helicopter's movements.

Decision variables:

- $z_{hit}$  1 if helicopter  $h$  begins inspection combination  $i$  at the beginning of week  $t$ , 0 otherwise,
- $y_{hlt}$  1 if helicopter  $h$  is operating at location  $l$  during week  $t$ , 0 otherwise,
- $w_{hit}$  1 if helicopter  $h$  is forced to return to home base in week  $t$  for impending inspection combination  $i$ , 0 otherwise,
- $e_{hlt}$  penalty variable that equals 1 if helicopter  $h$  is in location  $l$  in week  $t$  and in location  $l' \neq l$  in week  $t + 1$ , 0 otherwise,
- $x_{hit}$  number of hours of flight time remaining at the beginning of week  $t$  on helicopter  $h$  before inspection combination  $i$  is due,
- $c_{ht}$  number of hours helicopter  $h$  flies in week  $t$  at Clearwater,
- $g_{ht}$  shortfall between desired and actual number of hours flown at Clearwater by helicopter  $h$  during week  $t$ .

Formulation:

$$(P) : \min \sum_{h \in \mathcal{H}} \sum_{l \in \mathcal{L}} \sum_{t \in 1 \dots T-1} \alpha(T-t+1)e_{h,l,t+1} + \sum_{h \in \mathcal{H}} \sum_{t \in 1 \dots T} (T-t+1)g_{ht}$$

s.t.

$$\sum_{i \in \mathcal{I}} \sum_{t' \in 1 \dots t} S_{it'} z_{hit'} + \sum_{l \in \mathcal{L}} y_{hlt} = 1 \quad \forall h \in \mathcal{H}, t \in 1 \dots T, \tag{1}$$

$$\underline{R}_l \leq \sum_{h \in \mathcal{H}} y_{hlt} \leq \bar{R}_l \quad \forall t \in 1 \dots T, l \in \mathcal{L}, \tag{2}$$

$$y_{hlt} \leq y_{h,l,t+1} + \sum_{i \in \mathcal{I}} w_{h,i,t+1} \quad \forall h \in \mathcal{H}, l \in \mathcal{L} - \{o\}, t \in 1 \dots T-1, \tag{3}$$

$$y_{hlt} \leq y_{h,l,t+1} + e_{h,l,t+1} \quad \forall h \in \mathcal{H}, l \in \mathcal{L} - \{o\}, t \in 1 \dots T-1, \tag{4}$$

$$H(1 - w_{hit}) + (U - 1)w_{hit} \geq x_{hit} \quad \forall h \in \mathcal{H}, i \in \mathcal{I}, t \in 1 \dots T, \tag{5}$$

$$U y_{hlt} \leq x_{hit} \quad \forall h \in \mathcal{H}, l \in \mathcal{L} - \{o\}, i \in \mathcal{I}, t \in 1 \dots T, \tag{6}$$

$$\underline{B} \leq \sum_{t \in 1 \dots T} c_{ht} + \sum_{l \in \mathcal{L} - \{o\}} \sum_{t \in 1 \dots T} F_l y_{hlt} \leq \bar{B} \quad \forall h \in \mathcal{H}, \tag{7}$$

$$\underline{G} \leq \sum_h c_{ht} \leq \bar{G} \quad \forall t \in 1 \dots T, \tag{8}$$

$$\underline{M}y_{hot} \leq c_{ht} \leq \bar{M} \left( 1 - \sum_{l \in \mathcal{L} - \{o\}} y_{hlt} - \sum_{t' \in 1 \dots t} S_{i't'}z_{hit'} \right) \quad \forall h \in \mathcal{H}, i \in \mathcal{I}, t \in 1 \dots T, \quad (9)$$

$$c_{ht} + g_{ht} \geq \hat{M} \quad \forall h \in \mathcal{H}, t \in 1 \dots T, \quad (10)$$

$$g_{ht} \geq \hat{M} \left( \sum_{l \in \mathcal{L} - \{o\}} y_{hlt} + \sum_{i \in \mathcal{I}} \sum_{t' \in 1 \dots t} S_{i't'}z_{hit'} \right) \quad \forall h \in \mathcal{H}, t \in 1 \dots T, \quad (11)$$

$$\sum_h \sum_t c_{ht} \geq \hat{G}, \quad (12)$$

$$\sum_{t' \in 1 \dots t} (S_{i't'}z_{hit'} + S_{i't't'}z_{h'i't'}) \leq 2 \quad \forall h, h' \in \mathcal{H}, i, i' \in \mathcal{I}, t \in 1 \dots T \ni h \neq h', \quad (13)$$

$$Hz_{hit} + x_{hit} \leq H \quad \forall h \in \mathcal{H}, i \in \mathcal{I}, t \in 1 \dots T, \quad (14)$$

$$z_{hit} + x_{hit} \geq 1 \quad \forall h \in \mathcal{H}, i \in \mathcal{I}, t \in 1 \dots T, \quad (15)$$

$$x_{hiT} \geq \hat{B}_i(1 - z_{hiT}) \quad \forall h \in \mathcal{H}, i \in \mathcal{I}, \quad (16)$$

$$c_{hT} \leq \tilde{G} \quad \forall h \in \mathcal{H}, \quad (17)$$

$$x_{h,i,t+1} = Hz_{hit} + x_{hit} - \sum_{l \in \mathcal{L} - \{o\}} F_l y_{hlt} - c_{ht} \quad \forall h \in \mathcal{H}, i \in \mathcal{I}, t \in 1 \dots T - 1, \quad (18)$$

$$z_{hit}, y_{hlt}, w_{hit}, e_{hlt} \text{ binary} \quad \forall h, i, l, t, \quad x_{hit}, c_{ht}, g_{ht} \geq 0 \text{ and integer} \quad \forall h, i, t. \quad (19)$$

Constraints (1) ensure that each helicopter is scheduled either for a heavy maintenance inspection combination or to a location each week. Constraints (2)–(7) constitute the general operational constraints and, respectively: (2) ensure that the required number of helicopters is stationed at each location each week, (3) require that a helicopter either be at the same deployment site as in the prior week or that it return to Clearwater for maintenance, (4) penalizes moves from deployment sites, (5) keep a deployed helicopter from returning to Clearwater until it possesses below a certain number of flight hours (usually 50) before its next inspection combination, (6) prevent helicopters with too few hours of flight time remaining before an inspection combination from deploying and (7) seek to distribute flying time equally amongst all aircraft. Constraints (8)–(11) are operational constraints specific to Clearwater. Constraints (8) bound the total flight time in Clearwater between a minimum and a maximum value for each week in the horizon; constraints (9) force a helicopter to fly between a hard minimum and maximum number of hours each week if it is in Clearwater while preventing it from accruing flight time in Clearwater if it is either deployed or in maintenance; constraints (10) penalize a helicopter for deviating from a target number of weekly flight hours in Clearwater; constraints (11) strengthen the formulation by forcing helicopters not in Clearwater to incur the penalty associated with missing target flight time there; finally, constraints (12) place a lower bound on the total flight hours in Clearwater over the planning horizon with a view to satisfying a long-term flight hour goal. The following three constraints address the way in which maintenance is conducted. Constraints (13) ensure that no more than two helicopters are simultaneously undergoing maintenance. Constraints (14) do not allow a helicopter to begin an inspection combination until it is due. Conversely, constraints (15) force a helicopter to begin maintenance when maintenance is due. Constraints (16) and (17) help to prevent end effects, specifically, a situation in which many helicopters come due for an inspection combination just beyond the end of the horizon. Respectively, the constraints force all helicopters to have a certain number of flight hours still remaining (at the beginning of the last week of the horizon) before each inspection combination, and, to complement this constraint, ensure that each helicopter flies no more than a moderate number of hours during the last week of the horizon. Constraints (18), the counting constraints, reset the flight hour clock after a helicopter enters maintenance, and decrement the flight hour clock according to a helicopter’s flight time. Finally, nonnegativity and integrality of variables are enforced, as appropriate. The objective function minimizes weighted violations of constraints (4) and (10).

Although we impose constraints (16) and (17) in  $(P)$  to avoid end effects, there are other approaches to ameliorate this issue. Grinold [12] describes and evaluates four methods: truncation, salvage and primal and dual equilibrium. However, he cites feasibility issues, as well as complexity of implementation, as drawbacks to using the four methods he describes. There may also be a need to identify auxiliary parameters, making the implementation of such techniques difficult for a lay user of our models. Walker [13] shows that certain infinite-horizon linear programs have an equivalent finite-horizon approximation and provides a method of determining near-optimal solutions for both linear and integer programs with infinite horizons. However, this procedure comes at the expense of increased computation time and is not easily implementable in our planning context. In the Numerical Results section, we contrast the solution time and solution quality of models with a time horizon of  $T$  and the above-mentioned end effect-eliminating constraints with a time horizon of  $T + n$  (where  $n$  is sufficiently large to eliminate end effects) omitting the end effect-eliminating constraints.

#### 4. Increasing model tractability

$(P)$  is an NP-complete mixed integer linear program. Let us consider only a small subset of the constraints, specifically, constraint sets (2) and (7). For the special case in which  $\underline{G} = \bar{G} = 0$  (in (8)), we see that  $c_{ht} = 0$ . And, letting  $\underline{R}_l = \bar{R}_l = 1$  and  $\underline{B} = 0$ , and considering only a single time period (allowing us to drop the  $t$  subscript), constraint sets (2) and (7) reduce, respectively, to:  $\sum_h y_{hl} = 1 \forall l \in \mathcal{L}$  and  $\sum_l F_l y_{hl} \leq \bar{B} \forall h \in \mathcal{H}$ . Let the  $h$  index correspond to bins, the number of bins be the cardinality of  $\mathcal{H}$ , the  $l$  index correspond to items and the size of each item be  $F_l$ . Then, our problem reduces to the bin packing problem which is NP-complete [14, p. 226], so we conclude that our problem is at least as hard.

Because of the model's theoretical complexity and the fact that its instances are large, they are not easy to solve. We provide details in this section regarding how we have carefully formulated the problem in order to expedite solution time.

Without a weighting scheme, the objective would possess symmetry in that the coefficients on all terms would equal one, making it difficult for the algorithm to identify clearly dominating or dominated solutions. Such seemingly-identical solutions may, in actuality, not be equally desirable. In fact, empirical studies showed that solutions with equal weights on both objective function terms possessed an undesirably large amount of violation of helicopters remaining at a single deployment site until they return to home base for an impending inspection. Weighting the corresponding penalty by a factor of five improved solution quality. An additional consideration is a solution with slightly more total violation, but in which the majority of the violation occurs later in the horizon. Compare this solution with one in which less violation is incurred, but the violation occurs earlier in the horizon. Arguably, as the model may need to be reoptimized due to unplanned aircraft failures, penalties in the elasticized constraints early in the horizon are likely to be realized with greater probability than those later in the horizon. Therefore, postponing violations *to some extent* is arguably more desirable than incurring such violations early in the horizon. To this end, we weight both objective function terms decreasing by time such that violations earliest in the horizon are given the greatest weight. That is, the greatest weight is equal to the number of time periods in the time horizon,  $T$ , and is the weight on the terms when  $t = 1$ . Those violations one time period later are given one unit less weight, i.e.,  $T - t + 1 = T - 1$  when  $t = 2$ , etc. This not only breaks the symmetry in the model, but also produces desirable schedules. Others (see, e.g., Brown et al. [20]) have also found such discounting useful in improving integer programming model tractability.

Careful construction of two sets of constraints helps to provide a more tractable model, and also enhances the quality of solutions that the model yields. A cleaner formulation of constraints (3) and (4) uses the following single constraint set to penalize the movement of a deployed aircraft if it is not at the same deployment site in the next week as it was in the previous week, and if it is not returning to Clearwater for maintenance:

$$y_{hlt} \leq y_{h,l,t+1} + e_{h,l,t+1} + \sum_{i \in \mathcal{J}} w_{h,i,t+1} \quad \forall h \in \mathcal{H}, \quad l \in \mathcal{L} - \{o\}, \quad t \in 1 \dots T - 1. \quad (20)$$

However, constraints (20) permit helicopters without penalty to transition from Clearwater to a deployment site, if only for a short duration, and then to return to Clearwater for maintenance. The undesired result is helicopters with relatively few (but still a permissible number of) flight hours that deploy. In order to dissuade the model from producing solutions with this characteristic, we implement constraints (4) which explicitly penalize any movement of aircraft from a deployment site. Constraints (3) prevent deployed aircraft from moving until they return to home base for an



impending inspection combination. Not only do these constraints produce better solutions, but they increase model tractability in two ways: (i) they strengthen the formulation by providing a tighter upper bound on  $y_{hlt}$  (which, in  $(P)$ , is the minimum of the sum of two terms rather than the sum of three terms) and (ii) they differentiate the quality of solutions that would be equally good using constraints (20) by favoring fewer movements of helicopters to deployment sites.

The second set of carefully formulated constraints consists of constraints (10) and (11). These constraints penalize a helicopter's violation of the lower bound on the target number of weekly flight hours at Clearwater,  $\hat{M}$ , regardless of whether the helicopter is stationed at Clearwater or on deployment or in maintenance. A more intuitive constraint set would penalize for each week only those helicopters stationed at Clearwater:

$$c_{ht} + g_{ht} \geq \hat{M} \left( 1 - \sum_{l \in \mathcal{L} - \{o\}} y_{hlt} - \sum_{i \in \mathcal{I}} \sum_{t' \in 1 \dots t} S_{it't} z_{hit'} \right) \quad \forall h \in \mathcal{H}, t \in 1 \dots T. \quad (21)$$

Let us examine the mathematical structure of  $(P)$  versus that of  $(P)$  replacing (10) and (11) with (21) by considering simplified versions of these two problems,  $(\mathcal{P}^{\mathcal{S}})$  and  $(\mathcal{P}^{\mathcal{V}})$ , respectively, where, in our simplification, we: (i) drop all but one variable index ( $h = 1, 2$ ), (ii) examine only the relevant terms in the objective and constraints (8), (10) and (11) (or (21), rather than (10) and (11)) and (iii) let  $\zeta_h$  be a binary variable that combines the terms  $\sum_{l \in \mathcal{L} - \{o\}} y_{hlt}$  and  $\sum_{i \in \mathcal{I}} \sum_{t' \in 1 \dots t} S_{it't} z_{hit'}$ :

$$(\mathcal{P}^{\mathcal{S}}) : \min g_1 + g_2, \quad (22)$$

$$\text{s.t. } c_1 \leq \bar{M}(1 - \zeta_1), \quad (23)$$

$$c_1 + g_1 \geq \hat{M}, \quad (24)$$

$$c_2 \leq \bar{M}(1 - \zeta_2), \quad (25)$$

$$c_2 + g_2 \geq \hat{M}, \quad (26)$$

$$\underline{G} \leq c_1 + c_2 \leq \bar{G}, \quad (27)$$

$$g_1 \geq \hat{M}\zeta_1, \quad (28)$$

$$g_2 \geq \hat{M}\zeta_2, \quad (29)$$

$$c_1, c_2, g_1, g_2 \geq 0, \quad \zeta_1, \zeta_2 \text{ binary.} \quad (30)$$

$$(\mathcal{P}^{\mathcal{V}}) : \min g_1 + g_2, \quad (31)$$

$$\text{s.t. } c_1 \leq \bar{M}(1 - \zeta_1), \quad (32)$$

$$c_1 + g_1 \geq \hat{M}(1 - \zeta_1), \quad (33)$$

$$c_2 \leq \bar{M}(1 - \zeta_2), \quad (34)$$

$$c_2 + g_2 \geq \hat{M}(1 - \zeta_2), \quad (35)$$

$$\underline{G} \leq c_1 + c_2 \leq \bar{G}, \quad (36)$$

$$c_1, c_2, g_1, g_2 \geq 0, \quad \zeta_1, \zeta_2 \text{ binary.} \quad (37)$$

We now demonstrate that a lower bound on the objective function value of  $(\mathcal{P}^{\mathcal{S}})$  is always at least as tight as that of  $(\mathcal{P}^{\mathcal{V}})$ . If we add constraints (24) and (26) in  $(\mathcal{P}^{\mathcal{S}})$  and take the resulting difference from constraint (27), we obtain

$$\begin{cases} c_1 + c_2 + g_1 + g_2 \geq 2\hat{M}, \\ c_1 + c_2 \leq \bar{G} \\ \Rightarrow g_1 + g_2 \geq 2\hat{M} - \bar{G}. \end{cases} \quad (38)$$

If we add constraints (33) and (35) in  $(\mathcal{P}^{\mathcal{N}})$  and take the resulting difference from constraint (36), we obtain

$$\begin{cases} c_1 + c_2 + g_1 + g_2 \geq 2\hat{M} - \hat{M}\zeta_1 - \hat{M}\zeta_2, \\ c_1 + c_2 \leq \bar{G} \\ \Rightarrow g_1 + g_2 \geq 2\hat{M} - \bar{G} - \hat{M}\zeta_1 - \hat{M}\zeta_2. \end{cases} \quad (39)$$

Comparing the last lines of (38) and (39) we conclude that the two formulations  $(\mathcal{P}^{\mathcal{S}})$  and  $(\mathcal{P}^{\mathcal{N}})$  are equivalent in terms of the lower bounds that they provide on the objective only when the values of the binary variables are equal to zero. Because the binary variable values are not necessarily equal to zero,  $(\mathcal{P}^{\mathcal{S}})$  provides a tighter bound. This bound relates directly to the LP relaxation bound; in  $(\mathcal{P}^{\mathcal{N}})$ , fractional values for  $\zeta_h^*$  occur in the LP relaxation to lower the objective function value. The values of these same variables do not affect the bound in  $(\mathcal{P}^{\mathcal{S}})$ ; hence, the values have less incentive to be positive (fractional) which not only provides a tighter LP relaxation bound, but also leaves fewer variables available for branching.

Let us consider the following numerical example: let  $\bar{M} = 10$ ,  $\hat{M} = 8$ ,  $\underline{G} = 2$  and  $\bar{G} = 5$ . In  $(\mathcal{P}^{\mathcal{S}})$ , the optimal LP relaxation objective is equal to the optimal integer programming objective of 11. The corresponding solution, which is the same in either case, is  $c_1^* = 5$ ,  $g_1^* = 3$ ,  $g_2^* = 8$  and the other optimal variable values (most notably, the  $\zeta_h^*$ ) equal to zero. In  $(\mathcal{P}^{\mathcal{N}})$ , the optimal LP relaxation objective is 0 with a corresponding solution of  $\zeta_1^* = 1$ ,  $\zeta_2^* = 0.8$ ,  $c_2^* = 2$  and the other optimal variable values equal to zero. However, the optimal integer programming objective is 3 with a corresponding solution of  $\zeta_1^* = 1$ ,  $c_2^* = 5$ ,  $g_2^* = 3$  and the other optimal variable values equal to zero.

The poor LP relaxation value in  $(\mathcal{P}^{\mathcal{N}})$  relative to that in  $(\mathcal{P}^{\mathcal{S}})$  arises because, in the former case,  $g_1 + g_2 \geq 11 - 8\zeta_1 - 8\zeta_2$ . By inspection, we can see that setting both  $\zeta_1$  and  $\zeta_2$  equal to 1 yields a contradiction so that in any integer feasible solution, we must have  $\zeta_1 + \zeta_2 \leq 1$ , which implies that  $g_1 + g_2 \geq 3$ . However, in the LP relaxation, the binary variables  $\zeta_1$  and  $\zeta_2$  can assume fractional values such that  $g_1$  and  $g_2$  (and hence the objective) can assume values of zero. By contrast, in  $(\mathcal{P}^{\mathcal{S}})$ , the constraint  $g_1 + g_2 \geq 11$  holds independent of the values of the binary variables and indeed must be satisfied not only in the LP relaxation solution but for any integer feasible solution as well.

For sake of performance comparison in the Numerical Results section, we label the problem just described  $(\hat{P})$ . That is,  $(\hat{P})$  is identical to  $(P)$  with the following three exceptions: (i) the objective function coefficients in  $(\hat{P})$  all have a weighting of one, (ii) constraints (3) and (4) in  $(P)$  are replaced with (20) in  $(\hat{P})$  and (iii) constraints (10) and (11) in  $(P)$  are replaced with (21) in  $(\hat{P})$ . We provide details of how we tune the branch-and-bound algorithm for best performance of  $(P)$  and of  $(\hat{P})$  in the Numerical Results section. We discuss in the next section the Coast Guard’s evaluation of the model and resulting schedules and specifically of the quality of schedules and the ease with which the model and the corresponding output are implemented.

### 5. Model implementation and impact at CGAS, Clearwater

In the fall of 2004, the then-Coast Guard graduate student and coauthor of this paper demonstrated the model to HH60J maintenance managers. Specifically, the interested parties included the following officers: (i) the HH60J product line manager who is responsible for all HH60J depot-level maintenance and spare part support, (ii) the HH60J systems manager at Coast Guard headquarters in Washington DC in the Office of Aeronautical Engineering, who manages all HH60J maintenance budgeting and directly supports the product line manager, (iii) the CGAS, Clearwater engineering officer, who is responsible for all aircraft maintenance at Clearwater and (iv) the HH60J product line engineering cell leader, who is a technical expert in HH60J maintenance. The entire group was satisfied with the model, which is currently in its beta-implementation at Clearwater and is being considered for implementation at other Coast Guard HH60J air stations. The model is expected to have the greatest impact at Clearwater, which maintains the most HH60J helicopters.

A survey spreadsheet, which acts as a front-end to a model file written in the algebraic modeling language AMPL [15,16], allows the HH60J maintenance scheduling officer to enter data such as the next inspection combination that a helicopter must undergo at the start of the horizon, and the maximum number of hours a helicopter may fly until it must undergo that inspection combination (see Fig. 2). A script written in Visual Basic [17] takes the information from the spreadsheet and creates an AMPL data file formatted to run with the model. The CPLEX solver, Version 9.0 [18] installed at the Aircraft Repair and Supply Center in Elizabeth City, NC, runs the model. A Visual Basic script processes

Question	Input to the Optimization Model	
What is the current / next inspection on the helicopter from the beginning of Mar-28-2005?		
Helicopter	Inspection Type	
6003	1	
6013	1	
6015	9	
6016	4	
6017	8	
6018	7	
6019	2	
6025	3	
6040	3	
How many hours remaining to that current / next inspection on the helicopter at the beginning of Mar-28-2005?		
Helicopter	Inspection Type	Hours remaining
6003	1	185
6013	1	30
6015	9	0
6016	4	180
6017	8	0
6018	7	115
6019	2	50
6025	3	0
6040	3	75

Fig. 2. Snapshot of one worksheet of the input survey. An Excel workbook allows planners to avoid interface directly with AMPL. This end-user survey is currently being implemented in conjunction with the optimization model at CGAS in Clearwater, FL.

The Long Term Heavy Maintenance and Deployment Schedule for 12 weeks

	Week 1	Week 2	Week 3	Week 4	Week 5	Week 6	Week 7	Week 8	Week 9	Week 10	Week 11	Week 12
6010	24	16	HM-4	HM-4	HM-4	21	21	21	21	21	21	21
6015	HM-1	16	16	18	16	16	16	16	16	16	14	15
6016	16	16	16	16	16	HM-8	HM-8	HM-8	16	16	16	15
6017	21	21	21	21	21	21	16	16	12	HM-7	13	15
6018	18	16	24	HM-7	16	16	16	16	16	18	18	18
6019	18	18	18	18	18	18	18	18	18	22	16	HM-7
6033	16	10	16	"	HM-2	HM-2	21	21	21	21	21	21
6038	21	21	21	21	21	16	16	13	HM-1	16	16	15
6039	16	16	16	16	16	16	16	16	16	16	15	HM-2

Fig. 3. Sample output for a 12-week horizon. Helicopters are scheduled both to operational duties at Clearwater and deployment sites, and for heavy maintenance. Aircraft identification numbers are shown in the first column and week numbers are shown in the first row. Cells in the three lightest shades of gray correspond to a weekly aircraft assignment to one of the two deployment sites or to the home base for the number of hours shown in that cell. The darkest cells containing "HM" followed by a number indicate that a helicopter is undergoing heavy maintenance, and, specifically, the inspection combination number following the "HM" (see Table 1).

the results, which are depicted in a Gantt chart (see Fig. 3) in which a row corresponds to a helicopter (identified by its tail number) and a column corresponds to a week in the planning horizon. Each cell gives the (color-coded) location at which a helicopter is to conduct its operations, along with the number of flight hours that week. Alternately, the cell gives the inspection number corresponding to the inspection combination that the helicopter is undergoing. Fig. 4 illustrates the entire software system.

All helicopters have a secret cushion of flight time, which is 10% of the number of flight time hours before a particular inspection combination comes due. So, for example, a helicopter may actually be flown 20 h over the 200-h limit before any 200-h inspection type, i.e., before inspection combinations 1, 2, . . . , 12, absolutely must be conducted. However, using this cushion is costly. The extra flight time creates wear and tear on the helicopters. Inspection combinations without using this cushion of flight time are scheduled to coincide with the replacement of aircraft components or consumable parts just after their useful hours have expired. Using the extra cushion of hours desynchronizes inspection combinations and the end of life of various components or consumable parts, causing them to be replaced early,

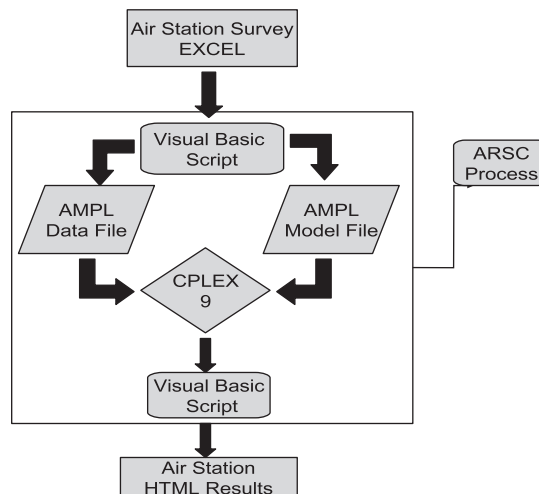


Fig. 4. Process flowchart. This flowchart depicts the way in which the input survey, optimization model and output presentation are linked to provide a workable, easy-to-use system for the Aircraft Repair and Supply Center (ARSC).

thus potentially wasting some component or part life. We do not factor this time cushion into the model, not only because it is costly but also because it is to be used only in extenuating circumstances. Were this cushion factored into the model, there would be no room for contingency planning. However, planners may and do alter the schedules using this cushion, e.g., in the case of unforeseen events. In such cases, the models are used as guidance, rather than strictly.

In January 2005, Air Station, Clearwater transferred one of its helicopters to another air station which had lost an aircraft. This presented the first opportunity for the model's use, in which planners ran the model using the same total yearly flight time requirements and only nine, rather than 10, helicopters to assess whether or not a feasible deployment and maintenance schedule existed with this reduced set of aircraft. The model showed that a heavy maintenance schedule was feasible only when the minimum number of hours a helicopter had to fly while at Clearwater each week was zero. The adjustment of this parameter was necessary to account for queuing of helicopters for maintenance. And although planners obtained a feasible schedule for the then-current horizon, they were aware that queuing for maintenance would ultimately lead to an inability to meet flying time requirements. By using the model to help to make this case, the operational commander reduced the total number of flying hours per year in Clearwater by about half of a single helicopter's annual flight time to compensate for the reduced number of aircraft.

Since this time, planners at the Aviation Logistics Division at the Coast Guard Aircraft Repair and Supply Center have begun extensive beta testing of the model with a view to working toward its full implementation. To this end, planners have correspondingly run numerous scenarios, which maintenance planners are using to help guide their scheduling decisions. Model instances generally solve to optimality in a matter of minutes and are completely justifiable and reproducible. In general, planners are pleased with the model's performance, output and utility. The model provides schedules with the following favorable characteristics that are not present in manually generated schedules: (i) the number of flight hours remaining until a helicopter is due for its next inspection combination when it returns from a deployment site is relatively consistent among aircraft and weeks; (ii) aircraft on deployment generally remain at that single deployment location for the majority of their flight hours between inspection combinations; (iii) end effects are addressed in that helicopter maintenance is evenly spaced throughout the horizon making bunching of maintenance unlikely in the subsequent horizon; (iv) specific flight time hours for helicopters in Clearwater are included in the model's decisions, rather than left to be estimated in real time and (v) because the model attempts to keep already-deployed aircraft at a site as long as possible, the model does not generally assign aircraft to a deployment site directly following maintenance. This last characteristic is particularly important because, due to the intrusive nature of the inspections, small problems may arise in the immediate postmaintenance weeks. If an aircraft is assigned to Clearwater directly following maintenance, experts can rectify the problems quickly with parts on-site. If a deployed aircraft becomes unoperational, it is difficult to fly the required hours on the fixed deployment patrols.

Additionally, the fact that the model yields a full, workable 12-week schedule in minutes, versus the 4–8 h required to manually generate a schedule, is a huge benefit that more than offsets a planner's effort of potentially making small changes to the computer-generated schedule in the case of unforeseen events. In an operational environment where helicopters are being repaired, parts are being ordered, failures are occurring, search and rescue help is being requested and drug interdiction is being conducted, among other daily activities, finding a quiet space to solve a brain teaser is difficult. Now, by entering some simple data into a spreadsheet, planners receive a viable schedule with little effort. They can then determine whether or not they have the time to make modifications or to take the schedule as is. Currently, the Aviation Logistics Division at the Coast Guard Aircraft Repair and Supply Center is working toward a full implementation of the model. Current efforts involve checking the validity of the model outputs, but, in general, planners are pleased with the schedules the model is producing.

## 6. Numerical results

We tested our model using 10 scenarios from the CGAS in Clearwater, FL. Because of the proprietary nature of the data, we discuss them only generally. The data sets differ in the following parameters: (i) in the three initial condition parameters (specified in Section 3 before the formulation), (ii) in the minimum number of hours a helicopter can fly at Clearwater each week ( $\underline{M}$ ), (iii) in the maximum number of flight hours flown each week at Clearwater ( $\bar{M}$ ), (iv) in the minimum number of hours flown by all helicopters over the planning horizon ( $\hat{M}$ ), (v) in the number of flight hours required in a week of a deployed helicopter ( $F_i$ ) and (vi) in the number of hours before an inspection combination below which an aircraft may not be deployed ( $U$ ). Whereas the initial condition parameters can differ substantially from one horizon to another, the other parameters differ between scenarios by generally 5% to 10% at most. Additionally, the second five data sets contain longer inspection types. These data sets are not as likely to correspond to current operations at Clearwater because there has been a trend toward shorter inspection types.

All 10 data sets require that either 9 or 10 helicopters be scheduled to some subset of 12 inspection combinations over a 12-week horizon. These are the largest data sets for HH60J aircraft scheduling applicable to any air station in the Coast Guard. No air station possesses more HH60J helicopters, and no air station is responsible for more deployment sites. A 12-week horizon is the maximum length planners are willing to consider due to the stochastic nature of events. Twelve weeks also corresponds to the maximum amount of time a helicopter can be stationed at a deployment site and is therefore, from the Coast Guard's perspective, a natural planning horizon length. Because the second five data sets contain longer inspection types, and because three of these instances contain 10 helicopters, these data sets, specifically, 7, 9 and 10, present the most difficult, yet realistic, scenarios that currently exist. The model instances passed to the solver, i.e., after the modeling language presolve, contain about 1000 binary variables and 15,000 constraints.

We conduct our numerical experiments using the AMPL programming language and the CPLEX solver, Version 9.0, and use the following CPLEX parameter settings: (i) the application of the relaxation induced neighborhood search heuristic every 120 nodes (for both ( $P$ ) and ( $\hat{P}$ )) [21], (ii) the emphasis of the algorithm on aggressively improving the lower bound (for ( $P$ )), and the emphasis of the algorithm on finding feasible solutions (for ( $\hat{P}$ )) and (iii) aggressively using a technique near the beginning of the search to reduce problem size by fixing variable values and eliminating constraints (for ( $\hat{P}$ )). We have carefully tuned these parameters to afford the best performance individually for ( $P$ ) and ( $\hat{P}$ ). We run all model instances on a Sun Ultra 10 computer with 256 MB RAM.

Table 2 gives the numerical results. On average, CPLEX returns an optimality gap of 27% for scenarios run for 1 h using the model without the enhancements, ( $\hat{P}$ ). Using the model stated in Section 3, ( $P$ ), CPLEX finds an optimal solution in an average of 79 s. Though we expect that it should be more difficult, on average, to obtain optimal solutions for the second five scenarios in which inspection types are longer, and specifically, for scenarios 7, 9 and 10 in which there are 10 helicopters (rather than nine) our improved model ( $P$ ) solves these five scenarios to optimality only 20 s slower, on average, than the more realistic first five scenarios. At any rate, we can reasonably conclude that we can solve real instances of ( $P$ ) to optimality within an acceptable amount of run time.

Because of the nature of the modifications we make, it is difficult to assess the quality of the objectives for instances of ( $\hat{P}$ ) relative to the strength of the lower bound. Specifically, the changes we make in the constraint sets cause penalties that would otherwise assume a value of zero to be sunk, thus artificially raising the objective function value of ( $P$ ). Even correcting for this still leaves qualitatively different solutions because of the weighting factor we employ on the first term in the objective of ( $P$ ). However, additional runs of greater length or with parameter settings that work well

Table 2  
Numerical results for 10 scenarios

Scenario number	Performance after 1 h	
	$(\hat{P})$ (%)	$(P)$ (s)
1	7	118
2	32	20
3	2	84
4	40	51
5	55	69
6	14	103
7	42	40
8	2	83
9	22	171
10	58	47

We compare the performance of the model without the enhancements,  $(\hat{P})$ , with that of the model stated in Section 3,  $(P)$ , after at most one hour of run time. Times reported in seconds signify that the model instance solved to optimality before the hour time limit, while percentages give the remaining gap after an hour if the model instance did not solve to optimality.

on specific instances suggest that the gap is a function in some cases of a poor integer solution, in others of a weak lower bound and sometimes of both.

In addition to solving  $(P)$  as stated in Section 3 for a 12-week planning horizon using end effect-eliminating constraints (16) and (17), we also solve the same model for a 20-week horizon without these two constraints and compare the quality of the solutions for a 12-week horizon. We deem a 20-week horizon sufficiently long to eliminate end effects over 12 weeks because the longest inspection combination lasts only three weeks. Using this easy-to-implement truncation method, we can determine whether our stated model adequately addresses the finite-horizon nature of our problem. When compared with the 12-week schedules that  $(P)$  generates, we observe that the truncated 20-week schedules exhibit poorer quality solutions in the following respects: (i) the number of hours flown on each aircraft over the horizon is not as equitably distributed and (ii) the target number of weekly flight hours for each helicopter stationed in Clearwater is not met to as great an extent. Furthermore, it is psychologically easier and more credible as a whole for planners to look at a good 12-week schedule than a 20-week schedule with poor advice near the end of the horizon. These undesirable characteristics, combined with the facts that planners are happy with the 12-week schedules  $(P)$  produces and run time to obtain a solution within 5% of optimality averages nearly 8 h for the 20-week model instances, lead us to conclude that  $(P)$  addresses end effect issues adequately for our planning purposes, and better than an obvious and easily implementable alternative.

By conducting numerical tests on  $(\hat{P})$  in which we impose each of the three modifications to  $(P)$  independently of each other, we are able to judge the relative impact of each. Specifically, imposing only the second modification discussed in Section 4, that of separating constraints (20) into constraints (3) and (4), results in model instances that are solvable to optimality in hundreds of seconds, on average, and has the greatest impact on model performance. Making only the first modification, that of replacing constraints (21) with constraints (10) and (11), results in model instances with gaps of a few percent after an hour of run time and has a moderate impact on model tractability. The modification with the least impact is that of placing weights on the terms in the objective function. This modification results in model instances with generally relatively large gaps (about 20%) after an hour of run time. Regardless of their individual contributions to enhancing overall model tractability, all modifications do have a significant impact on the performance of  $(P)$  and, when used in conjunction with each other, result in model instances that solve to optimality in less than 3 min.

## 7. Conclusions and future work

We have developed and implemented a mixed integer linear programming model for scheduling Coast Guard helicopter deployment and maintenance. The scheduling model distinguishes itself in that detailed maintenance must be combined with operations scheduling at a home base and at several deployment sites. We specifically address concerns such as the movement of helicopters from deployment sites between inspection combinations, the relationship

between the amount of flight time since or before an inspection combination and the location of an aircraft and end effects. Coast Guard schedulers are pleased with the optimization model, which is incorporated into a decentralized, easy-to-use, scheduling system. Current Coast Guard efforts are directed at evaluating the model, which appears to produce schedules that exceed both solution time and quality level expectations. Due to the ease with which model instances can be solved, planners might be interested in alternate optimal solutions. Planners might also be interested in modifying the model to account for limited maintenance capacity in which two helicopters can only simultaneously be undergoing an inspection combination if at least one of the two combinations is the least intrusive type. The addition of such constraints increases model size, and finding methods to increase tractability of these models may be worthwhile. Another option that may be attractive to planners in the future is incorporating the idea of *persistence* [19] into the schedules to minimize the disturbance to the present schedule should unforeseen events be so disruptive as to require rerunning the optimization model for the current planning horizon.

## Acknowledgments

We acknowledge the crucial implementation assistance of Jay Dhariwal, Supply Chain Analyst, at the Aviation Logistics Analysis Division at the United States Coast Guard Aircraft Repair and Supply Center. We also acknowledge the guidance provided by several then-current Maintenance Scheduling Officers, Lieutenant Eric Carter and Lieutenant Craig Massello, both at the CGAS in Clearwater, FL. Finally, we acknowledge the comments of two anonymous referees the help of Professor David Morton of the University of Texas at Austin regarding the proof of NP-completeness, and the help of Dr. Ed Klotz of ILOG in identifying the strength of our ultimate mathematical formulation.

## References

- [1] United States Coast Guard. ([www.uscg.mil/USCG.shtm](http://www.uscg.mil/USCG.shtm)) [accessed February 14, 2005].
- [2] Boere N. Air Canada saves with aircraft maintenance scheduling. *Interfaces* 1977;7(3):1–13.
- [3] Martin C, Jones D, Keskinocak P. Optimizing on-demand aircraft schedules for fractional aircraft operators. *Interfaces* 2003;33(5):22–35.
- [4] Cohn A, Barnhart C. Improving crew scheduling by incorporating key maintenance routing decisions. *Operations Research* 2003;51(3): 387–96.
- [5] Brown G, Goodman C, Wood R. Annual scheduling of Atlantic fleet naval combatants. *Operations Research* 1990;38(2):249–59.
- [6] Ayik M. Optimal long-term aircraft carrier deployment planning with synchronous depot level maintenance scheduling. MS thesis in operations research, Naval Postgraduate School, Monterey, CA, March 1998.
- [7] Albright M. An optimization-based decision support model for the navy H-60 helicopter preventative maintenance program. MS thesis in operations research, Naval Postgraduate School, Monterey, CA, September 1998.
- [8] Meeks B. Optimally scheduling EA-6B depot maintenance. MS thesis in operations research, Naval Postgraduate School, Monterey, CA, September 1999.
- [9] Baker R. Optimally scheduling EA-6B depot maintenance and aircraft modification kit procurement. MS thesis in operations research, Naval Postgraduate School, Monterey, CA, September 2000.
- [10] Sibre C. A quadratic assignment/linear programming approach to ship scheduling for the U.S. Coast Guard. MS thesis in operations research, Naval Postgraduate School, Monterey, CA, 1977.
- [11] Brown G, Dell R, Farmer R. Scheduling Coast Guard District cutters. *Interfaces* 1996;26(2):59–72.
- [12] Grinold R. Model building techniques for the correction of end effects in multistage convex programs. *Operations Research* 1983;31(2): 407–31.
- [13] Walker S. Evaluating end effects for linear and integer programs using infinite horizon linear programming. PhD dissertation in operations research, Naval Postgraduate School, Monterey, CA, 1995.
- [14] Garey M, Johnson D. *Computers and intractability: a guide to the theory of NP-completeness*. San Francisco, CA: W.H. Freeman; 1979.
- [15] Fourer R, Gay D, Kernighan BW. *AMPL: a modeling language for mathematical programming*. Pacific Grove, CA: Thompson Learning; 2003.
- [16] AMPL Optimization LLC. *AMPL*, version 10.6.16. 2001. ([www.ampl.com](http://www.ampl.com)).
- [17] Microsoft Corporation. *Visual Basic for applications*, version 6.3. 2003. ([www.microsoft.com](http://www.microsoft.com)).
- [18] ILOG Corporation. *CPLEX*, version 9.0. 2003. ([www.ilog.com](http://www.ilog.com)).
- [19] Brown G, Dell R, Wood R. Optimization and persistence. *Interfaces* 1997;27(5):15–27.
- [20] Brown G, Dell R, Newman A. Optimizing military capital planning. *Interfaces* 2004;34(6):415–25.
- [21] Danna E, Rothberg E, Le Pape C. Exploring relaxation induced neighborhoods to improve MIP solutions. *Mathematical Programming* 2005;102(1):71–91.