Working Paper Series

17/2006

Scheduling patrol boats and crews for the Royal Australian Navy

Horn, M.E.T., Jiang, H. and Kilby, P.

CAMBRIDGE
Judge Business School

Author contact details are as follows:

Mark E T Horn
CSIRO Mathematical and Information Sciences
Canberra, Australia
mark.horn@csiro.au

Houyuan Jiang
Judge Business School
University of Cambridge
h.jiang@jbs.cam.ac.uk

Philip Kilby
Research School of Information Sciences and Engineering
Australian National University
philip.kilby@anu.edu.au

This paper is forthcoming in *Journal of the Operational Research Society*.

Please address enquiries about the series to:

Research Support Manager
Judge Business School
Trumpington Street
Cambridge CB2 1AG, UK
Tel: 01223 760546 Fax: 01223 339701
E-mail: research-support@jbs.cam.ac.uk

Scheduling patrol boats and crews for the Royal Australian Navy

For *Journal of the Operational Research Society*

Mark E. T. Horn[1*], Houyuan Jiang[2], and Philip Kilby[3]

Last update: 6 March 2006

1. CSIRO Mathematical and Information Sciences, GPO Box 664, Canberra ACT 2601, Australia. mark.horn@csiro.au
2. Judge Business School, University of Cambridge, Trumpington Street, Cambridge CB2 1AG, UK. h.jiang@jbs.cam.ac.uk
3. Research School of Information Sciences and Engineering, Australian National University, Canberra ACT 0200, Australia. philip.kilby@anu.edu.au

* Corresponding author. Tel: +61-2-6216-7054, Fax: +61-2-6216-7111.

Abstract

The Royal Australian Navy's Patrol Boat Force carries out essential tasks in the surveillance, policing and defence of Australia's coastal waters. To help the Navy make efficient use of a new generation of boats, the authors have developed optimisation procedures to schedule the activities of the boats and their crews. The procedures, called CBM, use simulated annealing and specialised heuristic techniques within a multi-stage problem-solving framework. Tests show that CBM is reliable in terms of solution quality, and flexible with respect to the range of scheduling conditions applied. CBM has proved valuable to the Navy as an investigatory tool, and it is planned that it should be adapted for operational use, as part of a decision support system to aid in the ongoing management of patrol boat operations.

Keywords: heuristics, military, multi-objective, optimization, planning, scheduling, metaheuristics, simulated annealing, penalty methods.

Introduction

The Royal Australian Navy's Patrol Boat Force carries out essential tasks in the surveillance, policing and defence of Australia's coastal waters. This is a substantial responsibility, given that Australia, an "island continent", has a coastline more than 25,000 km. long. The Patrol Boat Force "carries out surveillance, interception, investigation, apprehension and the escort to port of vessels suspected of illegal fisheries, quarantine, customs or immigration offences." (http://www.minister.defence.gov.au/Hilltpl.cfm?CurrentId=3079, accessed 3 March 2006).

The Patrol Boat Force currently comprises 15 Fremantle Class vessels, capable of speeds up to 30 knots, with a complement of 24 crew members per boat (http://www.navy.gov.au/fleet/patrol.html, accessed 3 March 2006). These boats came into service in the early 1980s, and although they are still very effective, their maintenance is becoming difficult to sustain. The Australian Government announced a replacement program in July 2001 (http://www.minister.defence.gov.au/2001/240.doc, accessed 3 March 2006). The new vessels, known as Armidale Class, will be larger than their predecessors (e.g. 56 metres instead of 42 metres in length), with greater geographical range and more powerful weapons and communications systems.

In planning the replacement program the Navy took the opportunity to review the procedures used in scheduling the patrol boats' activities. The existing practice was to define timings for those activities assuming a permanent "marriage" of each patrol boat with a particular crew. This was good for the sailors' morale, but with substantial portions of the year given over to maintenance of the boats and to leave and on-shore training for crews, it placed limits on the time a boat or a crew could spend actually at sea. Consequently there was good reason to consider an alternative arrangement known as *multi-crewing*, under which more than one crew can be assigned to each boat over time, and vice versa. With more crews than boats, this should permit more intensive use of the boats. The potential for greater flexibility and efficiency was intuitively apparent to the Navy, but it was apparent that care would be needed to assure satisfactory working conditions for the crews, and the simple spreadsheet-based scheduling methods in current use gave no support for exploring the combinatorial complexities posed by multi-crewing.

To overcome these deficiencies the Navy sought assistance from the Australian Defence Science and Technology Organisation (DSTO), and the Commonwealth Scientific and Industrial Research Organisation (CSIRO), Australia's national research agency. CSIRO's rôle was to analyse the Navy's scheduling requirements and to develop procedures to meet those requirements, while quality control tasks were carried out by DSTO.

The project was initiated after the tenderers for the replacement program had specified maintenance requirements for the new boats. The Navy wished to explore the resource implications of the tenderers' proposals, through questions such as the following.

- Given a set of tasks to be performed by a certain number of patrol boats, how many crews would be needed to support those tasks, and could a satisfactory schedule of leave and training times be provided for the crews?

- What overheads would be involved in multi-crewing, in terms of travel and other costs incurred by change-overs between crews?
- What would be the smallest number of boats that must be purchased in order to ensure support for the expected tasks?

To help the Navy to find answers to questions like these, our approach has been to emulate the scheduling process. Although a voluminous literature exists on scheduling and staff rostering (Ernst, Jiang et al. 2004), the problem considered here is distinctive in that it requires (in effect) two conventional problems to be solved simultaneously. In particular, a solution will comprise schedules for both boats and crews, linked by way of a common set of activities (i.e. the deployments of the boats). This involves a departure even from recent research on integrated vehicle and crew scheduling (Haase et al. 2001, Cordeau et al. 2001, Freling et al. 2003), in that it requires the timing of all activities.

To address these conditions we have developed a specialized software package to plan the efficient use of crews and boats, called CBM ("Crews, Boats and Missions"). The package uses a range of optimization techniques, including steepest-descent improvement, simulated annealing, and specialised constructional heuristics. In a parallel investigation, we have applied Integer Linear Programming techniques (ILP) to the scheduling problem, in an effort to provide a benchmark for the evaluation of CBM.

In the next Section of this paper we outline the scheduling conditions addressed in the project. In subsequent Sections we describe the problem-solving framework adopted for CBM, the handling of scheduling conditions within that framework, and key aspects of the CBM annealing procedure. We then describe the decomposition and formulations used in the ILP algorithms, and report the outcomes of computational tests. We conclude with a discussion of our experience in developing and applying the procedures.

Scheduling conditions

The scheduling problem addressed here follows from a set of decisions taken by the Navy with respect to workload and resource allocations over a given period. Those decisions define, in aggregate, the activities to be performed, and the numbers of crews and boats available to perform the activities. The main scheduling tasks then are to establish timings for all the activities, and to assign each activity to a specific boat and a specific crew. Associated tasks are to assign each crew to a “home port”, and to determine the location of each boat when not deployed at sea.

The scheduling period of primary interest to the Navy is a calendar year, and the "atomic" time-unit is a week. We explain the various conditions below; further details are given in Tables 1-3, and under the heading "Handling scheduling conditions", below.

Activities of boats

The main activities of a patrol boat are its maritime *deployments*. A *mission* involves the deployment together of one or more boats at the same time: it comprises a certain number of boats (typically one or two), and has a certain duration (typically between one and eight weeks). The geographical attributes of a mission are considered irrelevant in the scheduling context. One type of mission with specific constraints is a *workup*, which provides “hands-on” training for the crews of all the participating boats, except for a *consort boat*. The Navy specifies the work of the Patrol Boat Force as a set of *mission groups*, each group comprising a specified number of missions of a particular type that are to be planned during a specified time-window. For example, a mission group might be defined as six fisheries-patrol missions, each of three weeks duration and requiring a single vessel, spread evenly in the period from March to November.

When not deployed at sea each boat is located at a port (Darwin or Cairns). With multi-crewing, a boat returns after each deployment to the home-port of the crew assigned to that deployment; this port is also the place where any maintenance activity scheduled before the next deployment will be carried out. The time taken by a boat to reach or return from a deployment is considered negligible for scheduling purposes. Apart from planned deployments, provision must be made for unexpected contingencies, called *surge* conditions. For this purpose the schedule must reserve a certain minimum number of “uncommitted” boats in port at all times.

Maintenance is pre-specified as a set of activities to be performed on each boat, each with a particular duration and a time-window within which the activity must be performed. The scheduling process must then define the time and the port at which each maintenance activity is to be performed. Maintenance activities on a given boat can overlap in time, and no restriction is placed on the location of a maintenance activity except for a limit on the total number of boats in maintenance at each port at any given time.

Activities of crews

The main activities of crews are *deployments* at sea, *on-shore training* and *leave*. An *operational cycle* is a period between successive deployments of a crew. Because we assume that all activities are undertaken simultaneously by all members of a crew, we need not consider crew members as individuals. Furthermore, we make no prima facie distinction between crews: a scheduling problem merely specifies a certain number of crews based at each port.

Detailed requirements are defined for the three kinds of crew activity mentioned above. In the first place, each crew must be assigned to no more than one workup mission during the year, and the crew should be assigned to same boat for the workup and for the crew's next deployment after the workup. Secondly, each crew must have at least a week of on-shore training and other naval activities after each of its deployments. Thirdly, an annual quota of seven weeks' leave is defined for each crew, allocated if possible in blocks of three or four weeks; furthermore, the leave-blocks are required to be spread approximately evenly through the year, coinciding with school holidays as far as possible.

Workload and handovers

Several conditions are designed to equalize workloads over time and amongst crews and boats. These include objectives to make the timing of missions as nearly equal as possible within the time-windows of the mission-groups, and to equalize the operational cycles of all crews. Similarly, upper and lower limits are placed on the total deployment-time in each month, and on the total length of a crew’s deployments during the year.

Some further conditions are associated with the practice of multi-crewing. A handover occurs when a boat's current crew is replaced by another crew; the handover is *remote* if the two crews have different home ports. Associated objectives are to minimize the total numbers of both handovers and of remote handovers. The first of these is concerned with the extra work incurred when one crew replaces another on a boat, the second with the cost of flying a crew to a port other than its home port. In addition, limits are placed on the number of boats to which any crew can be assigned during a year, and on the number of crews assigned to any boat.

A notation for the scheduling problem is given in Table 1. This provides the basis for the definition of objectives and constraints, as set out in Tables 2 and 3 respectively.

Table 1: Notation for problem specifications

Table 2: Objectives

Table 3: Constraints

CBM: scheduling framework

Because of the evident impracticality of handling all the conditions simultaneously, we adopted a multi-stage scheduling scheme for CBM. The three main modules are outlined below, with reference to the conditions identified in Tables 2 and 3.

1. *CBM-FAS*. A Fleet Activity Schedule (FAS) defines timings for all maintenance activities and missions. The CBM-FAS module uses a steepest-descent improvement heuristic, with the search neighborhood comprising all re-timings of all maintenance activities and missions, subject to the time-window constraints. The primary criterion is objective O1 (spread missions evenly).
2. *CBM-COP*. A Combined Operations Plan (COP) assigns timed activities to individual boats and crews. The CBM-COP module instantiates the deployments of each mission, and arbitrarily assigns the deployment and maintenance activities (with timings inherited from the CBM-FAS stage) to crews, boats and ports. A simulated annealing procedure is then used to improve the assignment, with a search neighborhood comprising re-assignments of deployments with respect to boats or crews, exchanges of assignments in these respects, and changes to the maintenance timings defined originally by CBM-FAS. The objective is primarily concerned with crew equity considerations (Objectives O2 and O3b) and handover reduction (Objectives O5 and O6).
3. *CBM-LAT*. A Leave and Training Plan (LAT) allocates leave and training activities to each crew, using a specialized heuristic procedure in which the focus is on crew satisfaction criteria (Objectives O3c and O3d).

This scheme involves a progression from broad to more detailed scheduling decisions: there is obviously a danger that sub-optimal decisions are locked-in at too early a stage, but this is avoided in practice by the ability to generate a range of solutions by iterating through the system. In particular, alternative plans can be generated through the use of randomised starting plans (e.g. multiple FASes, then multiple COPs for each FAS), which can be built easily with constraints relaxed under the soft constraint mechanism discussed in the following Section. The most extensive search is reserved for the COP stage, as reflected in the computational behaviour of the procedures; for example, in the computational tests described later in this paper, the execution times for CBM-FAS, CBM-COP and CBM-LAT were approximately in proportions 1 : 1000 : 0.1.

CBM: handling scheduling conditions

In CBM we treat most of the scheduling constraints as "soft", by way of penalty components in the objective function. That is, the total cost to be minimized in each case is a weighted sum of simple costs associated with objectives, and penalty costs arising from violations of constraints (see Tables 2 and 3, column 4). The penalty mechanism includes several constraints that are "logically hard"; for example C8 and C16 say that a crew or a boat can do only one thing at a time, and incur penalties based on the total extent of any temporal overlaps in these respects. Other hard constraints (e.g. C1 and C2) are handled procedurally.

The constraint mechanism includes internal scaling factors as well as externally-defined weighting factors, both types of factor being applied within the optimization procedures as multipliers to the “raw” costs defined in Tables 2 and 3. The scaling factors convert the cost components to a common measure (boat-weeks), and so facilitate adaptation to problems of different size. The weighting factors are parameters indicating the relative importance of the various cost components. We determined their relative magnitudes by a process of trial and error, with the aim of assuring satisfaction of the hard constraints, while retaining sensitivity to the soft constraints (the set of values adopted in practice range from 1 to 150).

An interesting point regarding relations between the scheduling conditions and the framework outlined earlier concerns the use of overlapping criteria, where a single underlying intention may be reflected in several different constraints and objectives. Near-redundancy of this kind in the scheduling requirements mostly reflected specific client requirements, but it often also served a strategic purpose, where criteria with aggregate impacts in early stages of the problem-solving framework may help to ensure satisfaction of detailed conditions applied later on, as indicated below.

a. O3a penalizes missions scheduled during school holidays. Optimizing the FAS in this respect helps to ensure that CBM-LAT can obtain a good result with respect to O3b (regarding leave for crews during school holidays).

b. C3 states that the total number of boats deployed and under maintenance in each week must not exceed the total number of boats. Satisfying this condition in the FAS ensures prima facie feasibility in the COP with respect to C16 (two activities cannot occur at the same time).

c. C6a states that the total number of boats under maintenance in any given week must not exceed the total maintenance capacity of all ports. Satisfying this condition in the FAS ensures prima facie feasibility in the COP with respect to C6 (individual port capacities should not be exceeded).

d. O1 favours spreading missions evenly over each mission-group's time-window. Optimizing the FAS in this respect helps to ensure acceptable results with respect to COP conditions such as O2 (regular operational cycles for each crew).

e. Optimizing the COP with respect to O2 (regular operational cycles for each crew), O3b (at least one non-deployment period coinciding with school holidays), C9 (limiting each crew's total deployment activity), and C13 (at least one week to separate each successive pair of missions) helps to ensure acceptable results with respect to the LAT criteria concerned with crew activity timings.

CBM: annealing procedure

Simulated annealing is well-established as a technique that emulates phase-change processes in physics, gradually "cooling" a solution until it reaches a stable (and potentially optimal) form (Laarhoven and Aarts 1987). We chose it for the CBM-COP module because of its suitability to complex problems, notably its ability to escape from local optima through the use of protracted cooling schedules. We use the following notation to describe the simulated annealing scheme used in CBM-COP.

t – Current temperature.

cop – Current plan.

cop^* – Least-cost plan found so far.

$z(cop)$ – Cost of current plan.

`MaxTemp` – Initial temperature (≈ 25.0).

`MinTemp` – Terminal temperature (≈ 10.0).

`ItersPerTemp` – Number of tests to be performed at *t* (≈ 2000).

`CoolRatio` – Factor to be applied to *t* after `ItersPerTemp` tests (≈ 0.9995).

The annealing procedure was implemented initially in the following "standard" form.

1. Construct the initial *cop* with arbitrary assignments of deployments to crews and boats, subject to the activity timings given in the FAS. Define *cop** = *cop* and *t* = `MaxTemp`.
2. Do the following `ItersPerTemp` times.
   Obtain a modified plan *cop'* by random selection in the neighbourhood of *cop*, that is, by considering all available deployment-moves, deployment-swaps, and (subject to time-windows) all maintenance re-timings. Calculate the resulting increase in cost δ ($\delta = z(cop') - z(cop)$), then:

2a. If $\delta \leq 0$, accept the change ($cop \leftarrow cop'$), and if $z(cop') \leq z(cop^*)$, $cop^* \leftarrow cop'$.

2b. If $\delta > 0$, accept the change ($cop \leftarrow cop'$) with probability $e^{-\delta/t}$.

3. Apply cooling to *t* ($t \leftarrow t$. `CoolRatio`). Then if $t \geq$ `MinTemp`, repeat step 2; otherwise stop.

By itself this scheme did not yield the consistent downward trend in solution cost predicted by classical annealing theory; instead, the cost tended to fall rapidly in the initial stages of cooling, thereafter fluctuating without substantial further improvement, even with lengthy cooling schedules (e.g. with `CoolRatio` ≈ 0.9999 and `ItersPerTemp` ≈ 20000). This difficulty may be explained in broad terms by the large number of different cost components, and the mathematically irregular linkages between them.

To achieve more consistent and reliable performance we extended the annealing framework in several respects, with the general aim of focussing the procedure on "productive" regions of the search space (see Dowsland 1993). The extensions are outlined below.

a. The most important extension is concerned with control over the overall trend of the search. We found that after a *cop** was reached (i.e. after $cop^* \leftarrow cop'$ in step 2a), the value of $z(cop)$ sometimes would follow an upward trend, and then remain substantially higher than $z(cop^*)$ until the end of the run. To prevent such divergences from the (presumed) region of "good" solutions, we monitor the length of any sequence of changes to *cop* without improvement in *cop**, and restore $cop \leftarrow cop^*$ when the length reaches a predefined limit `ForceDownFrequency` (≈ 200).

b. Several extensions were designed to speed up optimisation. The implementation of the neighbourhood-scan operation in step 2 is very efficient, allowing the testing of changes at a rate of more than 50,000 per second (for hardware details see "Computational tests", below). Even so, the basic procedure was wasteful in that it accepted very few changes once a "middle range" of temperatures is reached. To curtail futile searching, we define several additional parameters which can take precedence over the `ItersPerTemp` and `MinTemp` limits in steps 2 and 3 under certain conditions:
   - `MaxUphillPerTemp` (≈ 500): limit on the number of non-improving changes made at any given temperature. After this limit is reached, only improving moves are accepted at the current temperature.
   - `MaxRejectRun` (≈ 100): limit on the number of consecutive rejections allowed at temperature *t*. When this limit is reached, the annealing procedure proceeds immediately to the next temperature.
   - `MaxFlatToFreeze` (≈ 500000): maximum number of changes to test without improvement to the least-cost plan. When this limit is reached, the annealing procedure terminates without waiting for the temperature to reach `MinTemperature`.

c. We allow control over the annealing procedure's view of the search neighbourhood, by means of parameters specifying the aggregate distribution of change-types to be tested. The aim here is to focus attention on the more fruitful types of changes, which (as we found) were swaps and, to a lesser extent, moves. The parameters are `MissSwapPortion` (≈ 0.5) and `MaintShiftPortion` (≈ 0.2), specifying the numbers of mission-swaps and maintenance-shifts respectively, in proportion to total changes. For example, if these parameters have values 0.5 and 0.2 respectively, the overall proportions of swaps, moves and shifts tested during the annealing run are approximately 50%, 30% and 20%.

ILP: decomposition and formulations

Our initial intention in developing a solution procedure based on integer linear programming (ILP) was to obtain exact benchmarks against which to assess solutions obtained from CBM. It was soon apparent however that available ILP codes could not handle all the scheduling criteria in a single step. We therefore devised instead a decomposed approach, still with the intent of “benchmarking” CBM. As in CBM, the respective stages are run sequentially, with the outputs of each stage fed as inputs to the next stage. The stages are similar to those in CBM, but with a less unified coverage of the tasks covered by CBM-COP and CBM-LAT.

Table 4: ILP modules

As indicated in Table 4, there are four ILP modules, concerned respectively with the construction of a FAS (ILP-FAS, like CBM-FAS), a Crew Activity Plan (ILP-CAP), a Boat Activity Plan (ILP-BAP), and Port Assignment Plan (ILP-PAP). The ILP implementation is quite complex, as indicated by the numbers of constraint-sets used (for details see Horn et al. 2003). It is however limited in its treatment of operational cycle lengths (O2 is not applied), and of the "sequence-based" criteria concerned with handovers, remote handovers and workup requirements (O5, O6 and C21 respectively). Of these, the treatment of O5 is of particular interest. Our investigations showed that an exact formulation of O5 in the ILP-BAP stage would pose a multiple travelling salesman problem, and hence a large and intractable ILP. To avoid this, we addressed handovers via a proxy measure analogous to that used in C18. With nc_b denoting the number of crews using a boat b to perform deployments during the planning period, the objective is to minimize the sum of nc_b taken over all boats. We used a similar proxy to handle remote handovers (O6) in ILP-PAP.

Computational tests

We ran tests of the CBM and ILP procedures for a scenario representing expected working conditions, with variants defined in terms of the size of the Patrol Boat Force. The basic scenario assumes 15 crews and 2 ports, and the work to be performed during a year comprises 22 maintenances (two maintenances on most boats), 11 mission-groups, 78 missions, and 101 deployments. The variants are referred to here as B-10, B-11 and so on, indicating numbers of boats ranging from 10 to 15. Of these, B-10 indicates a limit of feasibility (i.e. no feasible solution has been found at this level of multi-crewing), while at the other extreme, B-15, with one boat for each crew, is a single-crewing arrangement.

The CBM tests were run on a Dell Optiplex PC with a 1.7GHz Pentium-4 processor. The ILP tests were run on a 500MHz Dec Alpha, using CPLEX (CPLEX 2001). A non-zero *optimality gap* parameter was used when applying CPLEX to ILP-FAS, ILP-CAP and ILP-BAP, in order to avoid excessive running times; consequently the solutions found by the various ILP modules are in some cases sub-optimal.

Table 5 shows the results from the CBM and ILP procedures for each of the test problems. The CBM results in each case refer to a least-cost scheduling plan P^*, chosen from 100 full solutions generated for the problem; in particular, 10 COPs obtained for each of 10 FASes yielded 100 different plans, of which P^* was defined as a matched pair (FAS,COP). The criterion for selecting P^* was the total weighted cost Z (see below).

The first part of the table shows unweighted costs for the main objectives and constraints, as defined in the rightmost columns of Tables 2 and 3. Aggregate weighted costs are shown in the rows labeled Z_{FAS}, Z_{COP} and Z ($Z = Z_{\text{FAS}} + Z_{\text{COP}}$). For uniform comparison, all the costs have been calculated using the CBM code. In all our tests, the scheduling plan with the least value of Z also has the least Z_{COP}, presumably due to substantive dependencies between a COP and its underlying FAS. The last row of the table shows the total elapsed CPU time for each scheduling run (for CBM this is the time to generate all 100 solutions in each run).

Table 5: Computational tests of CBM and ILP

As the results show, CBM produced solutions for every problem, although the COP produced for B-10 included serious infeasibilities (i.e. violations of C3 and C16). The ILP procedures found a FAS for every problem except B-10, and a COP for every problem except B-10 and B-11. The availability of good plans that use fewer boats than crews indicates the feasibility of multi-crewing as a scheduling strategy.

While the FASes produced by the ILP procedure were of better quality than those from CBM, the COPs were consistently better from CBM than from ILP. These differences can be attributed mainly to differences in problem decomposition: CBM has effectively two stages so far as the main scheduling decisions are concerned, and is therefore more likely than ILP (with four stages) to obtain "globally good" results. Also, as indicated previously, several criteria are not addressed directly in the ILP models.

Conclusion

A major emphasis of the work reported here has been to model a problem in terms that are faithful to the nuances and complexities of actual conditions, for example with respect to the assurance of good working conditions for the Patrol Boat crews. This approach has yielded a set of criteria that are more extensive and irregular than those reported in the classical scheduling literature. In CBM however we have succeeded in addressing those criteria with an effectiveness that is indicated both by the computational tests reported above and by informal discussion with Navy clients.

In logistical terms, the research shows that revisions to the crewing and maintenance régimes of the Patrol Boat Force can lead to substantial improvements in efficiency while maintaining crew satisfaction, and ensuring at the same time that the patrol boats can meet their on-going and emerging commitments. In addition, following its use of CBM in the patrol boat replacement program, the Royal Australian Navy has confirmed the value of CBM in relation to key questions concerning the numbers of boats required, and maintenance and crewing régimes.

It is intended in future to adapt CBM for operational use, as part of a decision support system to aid in the ongoing management of patrol boat operations. Although this will require some revisions to scheduling conditions (e.g. allowing some decisions to be made in advance) and assumptions (e.g. regarding temporal scope and granularity), the solution framework described here appears to be sufficiently robust to accommodate such changes in a straightforward way.

Finally, the project raises an interesting methodological point regarding relations between exact and approximate methods in practice. The modeling of clients' requirements is more faithful in CBM than in ILP, and the solutions obtained with CBM are clearly of superior quality to those from ILP. The lesson here is that although heuristics are commonly regarded as inferior to exact methods, in practice they can be matched more precisely to the situation at hand, without sacrifice of solution quality.

Acknowledgments

Commander Gerry Kelly of the Royal Australian Navy, and Dr. Ping Cao of the Australian Defence Science and Technology Organization, were sponsors of the project discussed in this paper. The authors thank them for their support and encouragement.

References

CPLEX (2001). CPLEX version 7.0, Users' manual for callable library. CPLEX Optimization Inc.

Cordeau J-F, Stojković G, Soumis F and Desrosiers J (2001). *Benders decomposition for simultaneous aircraft routing and crew scheduling*. Transportation Science, 35(4):375–388.

Dowsland KA (1993). Simulated annealing. In: CR Reeves (ed). Modern heuristic techniques for combinatorial problems. Wiley: New York, pp. 20-69.

Ernst AT, Jiang H, Krishnamoorthy M, and Sier D (2004). Staff scheduling and rostering: A review of applications, methods and models. European Journal of Operational Research 153(1): 3-27.

Freling R, Huisman D, and Wagelmans APM (2003). *Models and algorithms for integration of vehicle and crew scheduling*. Journal of Scheduling, 6(1):63–85.

Haase K, Desaulniers G, and Desrosiers J (2001). *Simultaneous vehicle and crew scheduling in urban mass transit systems*. Transportation Science, 35(3):286–303.

Horn MET, Jiang H and Kilby PJ (2003). Scheduling the Australian Patrol Boat Force. CSIRO Mathematical and Information Sciences Technical Report, CMIS 03/169.

Laarhoven van PJM and Aarts EHL (1987). Simulated Annealing: Theory and Application. Reidel: Dordrecht.

Table 1: Notation for problem specifications

Symbol	Definition	Constraint
B	The set of boats specified for the problem, distinct with respect to their maintenance requirements.	
C	The set of crews. In a problem specification this is merely implicit in $ncrews(P)$.	
P	The set of ports at which crews are based and where maintenance can be performed.	
MA	The set of maintenances to be performed on B.	C1
MG	The set of mission-groups specified for the problem.	C2
$nboats$	The number of boats in the patrol-boat fleet, $nboats = \|B\|$	
$nports$	The number of ports, $nports = \|P\|$.	
$ncrews_p$	The number of crews for which p is the home port.	C7c
$ncrews(P)$	The total number of crews to be assigned to boats, $ncrews(P) = \|C\| = \sum_{\forall p \in P} ncrews_p$	C7b
$mcap_p$	The maintenance capacity of port p, that is, the number of boats on which maintenance can be carried out at p at any given time.	C6b
$mcap(P)$	The total maintenance capacity of all ports, $mcap(P) = \sum_{\forall p \in P} mcap_p$	C6a
$wdep(MG)$	The total duration of deployments implied by MG, in boat-weeks.	
$ndeps(MG)$	The total number of deployments implied by MG.	
`MinAvail`	Minimum workload per month, as a ratio of the average monthly workload.	C4
`MaxAvail`	Maximum workload per month, as a ratio of the average monthly workload.	C4
`MinSurge`	Minimum number of boats to be reserved at any time to meet surge conditions.	C5
`MinDeployed`	Minimum number of annual deployment weeks for any crew.	C9
`MaxDeployed`	Maximum number of annual deployment weeks for any crew.	C9
`LeaveYearly`	Total annual leave for each crew, in weeks.	C10a
`MinLeaveBlock`	Smallest allowable leave-block for a crew, in weeks.	C10b
`MaxLeaveBlock`	Largest allowable leave-block for a crew, in weeks.	C10b
`MinTrainBlock`	On-shore training period to be allocated to each crew between deployments, in weeks.	C13
`MaxCrewsPerBoat`	Maximum number of different crews that can be assigned to any boat during the year.	C18
`MaxBoatsPerCrew`	Maximum number of different boats to which any crew can be assigned during the year.	C19

Table 2: Objectives

Objective	CBM module	Description	Cost
O1	CBM-FAS	The missions of each mission-group *mg* should have their timings spread evenly across the time-window of *mg*.	Sum over all mission-groups $mg \in MG$: sum over all missions $m \in mg$: absolute difference between timing of *m* and a regular timing with respect to time-window of *MG*.
O2	CBM-COP	The durations of all operational cycles should be as nearly equal as possible.	Sum over all operational cycles *c*: absolute difference between duration of *c* and "ideal" duration, defined as 52 . *ncrews(P)* / *ndeps*(*MG*).
O3a	CBM-FAS	Deployments should not be made during school holidays, as far as possible.	Sum of deployment-weeks occurring during periods other than school holidays.
O3b	CBM-COP	Each crew should have at least one week of leave during school holidays.	Number of crews for which a holiday-leave match is infeasible. Such a match requires at least one non-deployment interval with size > `MinTrainBlock + MinLeaveBlock`, and with at least one week coinciding with a school holiday.
O3c	CBM-LAT	The leave-blocks allocated to each crew should be spread as evenly as possible through the year.	*Not explicit.*
O3d	CBM-LAT	The leave-blocks allocated to each crew should be as large as possible (see also constraint 10b).	*Not explicit.*
O5	CBM-COP	Handovers of boats are to be minimized. In the sequence of deployments assigned to a boat, a handover occurs when there is a change of crews between one deployment and its successor.	Total number of handovers.
O6	CBM-COP	Remote handovers are to be minimized. A remote handover occurs when a boat has a deployment of a crew *a* with home port p_a, followed by a deployment of another crew *b* with home port $p_b \neq p_a$.	Total number of remote handovers.

Table 3: Constraints

Constraint	CBM module	Description	CBM penalty cost
C1	CBM-FAS, CBM-COP	All specified maintenance activities (*MA*) must be included in the scheduling plan, subject to their respective time-window limits.	*Hard constraint*
C2	CBM-FAS, CBM-COP	All deployments implied by the mission-groups (*MG*) must be included in the scheduling plan, subject to their respective time-window limits.	*Hard constraint*
C3	CBM-FAS	The total of deployments and boats under maintenance in each week must not exceed *nboats*.	Sum over all weeks: sum of boats deployed and in maintenance, in excess of *nboats*.
C4	CBM-FAS	The total boat-weeks deployed during each month must be within lower and upper bounds defined for that month. The bounds for month *m* are defined as ratios `MinAvail`, `MaxAvail`, applied to the deployment-density quota, which is obtained as $wdep(MG)\ .\ d_m / 365$, where d_m is the number of days in month *m*.	Sum over all months: sum of deployments in violation of the defined bounds.
C5	CBM-FAS	The number of boats available for surge in each week is defined as *nboats* minus the number of boats assigned to deployments in that week; this number should be no less than `MinSurge`.	Sum over all weeks: number of boats in deficit of surge requirement.
C6a	CBM-FAS	The total number of boats under maintenance at any given time must not exceed *mcap(P).*	Sum over all weeks: number of boats in maintenance in excess of *mcap(P).*
C6b	CBM-COP	The number of boats under maintenance at each port *p* at any given time must not exceed $mcap_p$.	Sum over all weeks: sum over all ports $p \in P$: number of boats in maintenance at *p* in excess of *mcap(p).*
C7a	CBM-COP	All deployments must be assigned to crews.	*Hard constraint*
C7b	CBM-COP	The number of crews in the scheduling plan must be *ncrews(P).*	*Hard constraint*
C7c	CBM-COP	The number of crews based at port *p* must be $ncrews_p$.	*Hard constraint*
C8	CBM-COP	A crew can perform no more than one activity (i.e. a deployment, leave or training) at a time.	Sum over all crews $c \in C$: sum over all deployments *d* of *c*: number of weeks during which *d* overlaps with other deployments of *c.*
C9	CBM-COP	The total time during which each crew is deployed during a year is subject to lower and upper limits `MinDeployed`, `MaxDeployed`.	Sum over all crews $c \in C$: extent to which total deployment-weeks of *c* lie below or above the acceptable range.
C10a	CBM-LAT	Each crew must have `LeaveYearly` weeks of leave during the year.	*Hard constraint*
O10b	CBM-LAT	The duration of the leave-blocks allocated to crews is subject to lower and upper limits `MinLeaveBlock`, `MaxLeaveBlock`.	*Hard constraint*

Constraint	CBM module	Description	CBM penalty cost
C12	CBM-COP	Each crew should be assigned to no more than one workup deployment during the year.	Number of instances of a boat $b \in B$ assigned to more than one workup mission.
C13	CBM-COP, CBM-LAT	Each crew should have a period of training during each operational cycle, of duration `MinTrainBlock`.	Sum over all boats $b \in B$: number of instances of operational cycles separated by periods smaller than `MinTrainBlock`.
C15	CBM-COP	All deployments must be assigned to boats.	*Hard constraint*
C16	CBM-COP	No more than one activity (i.e. a deployment, or a set of maintenances) can be performed on a given boat at any given time.	Sum over all boats $b \in B$: sum over all deployments d of b: number of weeks during which d overlaps with other deployments of b.
C18	CBM-COP	The number of different crews assigned to deployments carried out by each boat must not exceed `MaxCrewsPerBoat`.	Sum over all boats $b \in B$: extent to which the number of crews assigned to b exceeds `MaxCrewsPerBoat`.
C19	CBM-COP	The number of different boats on which a crew may be deployed must not exceed `MaxBoatsPerCrew`.	Sum over all crews $c \in C$: extent to which the number of boats assigned to c exceeds `MaxBoatsPerCrew`.
C20	CBM-COP	Each maintenance activity of a boat must be done at the home port of the crew assigned to the boat for the deployment immediately preceding or following the maintenance.	*Hard constraint*
C21	CBM-COP	After a workup deployment of a given crew on a given boat, the boat's next deployment (if any) should be conducted by the same crew.	Sum over all boats $b \in B$: number of cases in which different crews are assigned to a workup mission and its successor.

Table 4: ILP modules

Module name	Output	Function	Number of constraint-sets
ILP-FAS	Fleet Activity Schedule	Like CBM-FAS, defines timings for all missions in each mission group and all maintenance services for all boats.	12
ILP-CAP	Crew Activity Plan	Assigns missions, leave, trainings, and standby periods to crews so as to maximize quality of life for crews in terms of leave and holiday timings.	22
ILP-BAP	Boat Activity Plan	Assigns boats to deployments so as to minimize, in principle, the number of handovers between crews.	6
ILP-PAP	Port Assignment Plan	Assigns crews to ports so as to minimize, in principle, the number of remote handovers.	5

Table 5: Computational tests of CBM and ILP

	B-10		B-11		B-12		B-13		B-14		B-15	
	CBM	ILP	CBM	ILP	CBM	ILP	CBM	ILP	CBM	ILP	CBM	ILP
O1	171.6	-	176.5	80.15	103.9	84.64	111.7	78.21	83.6	79.33	98.2	74.46
O2	210.0	-	184.0	-	200.0	242.0	171.0	237.0	193	232.0	165	211.0
O3a	112	-	110	110	107	119	115	112.0	117	111	106	109
O3b	0	-	1	-	0	0	0	0	0	1	0	0
O5	42	-	29	-	27	57	14	37	12	18	2	8
O6	0	-	1	-	0	4	1	5	0	1	0	0
C3	1	-	0	0	0	0	0	0	0	0	0	0
C4	0	-	0	0	0	0	0	0	0.29	0	0	0
C5	0	-	0	0	0	0	0	0	0	0	0	0
C6a	0	-	0	0	0	0	0	0	0	0	0	0
C6b	0	-	0	-	0	0	0	0	0	0	0	0
C8	0	-	0	-	0	0	0	0	0	0	0	0
C9	0	-	0	-	1	0	0	0	0	0	0	0
C12	1	-	2	-	0	0	1	0	2	0	0	0
C13	15	-	15	-	7	0	6	0	2	0	4	0
C16	2	-	0	-	0	0	0	0	0	0	0	0
C17	0	-	0	-	0	0	0	0	0	0	0	0
C18	0	-	0	-	0	0	0	0	0	0	0	0
C19	0	-	0	-	0	0	0	0	0	0	0	0
C21	0	-	0	-	0	7	0	5	0	1	0	1
Z_{FAS}	303.6	-	286.5	190.15	210.9	203.7	226.7	190.2	203.5	190.3	204.2	183.5
Z_{COP}	964.0	-	599.0	-	469.0	1124.0	374.0	909.0	317.0	435.0	219.0	337.0
Z	1267.6	-	885.5	-	679.9	1327.7	600.7	1099.2	520.5	625.3	423.2	520.5
CPU time (minutes)	535	*	512	*	468	140	474	70	415	255	415	55

* For B-10 and B-11, execution of the ILP procedures was terminated after running approximately a week.

Table captions

Table 1: Notation for problem specifications

Table 2: Objectives

Table 3: Constraints

Table 4: ILP modules

Table 5: Computational tests of CBM and ILP