

# Force Generation Plan Automation for the Royal Australian Navy

T. Weir

*Joint and Operations Analysis Division, Defence Science and Technology Organisation  
Email: [terence.weir@dsto.defence.gov.au](mailto:terence.weir@dsto.defence.gov.au)*

**Abstract:** The Royal Australian Navy's (RAN) Force Generation Plan (FGP) details the day-to-day operations of about 120 ships, crews, submarines, diving teams and air assets plus miscellaneous entities. The FGP is an annual directive concerned with scheduling the activities that are undertaken by the RAN units. Required activities are specified in an associated document called the Navy Activity Schedule (NAS). The NAS provides overarching guidance to RAN maritime operations planning staff on Navy's priorities and resource levels for operational commitments, preparedness levels and international and domestic engagement activities.

The FGP is manually generated as a result of interactions and conferences involving Navy operations' authorities. A scheduling tool, the Fleet Activity Management Tool (FAMT), allows schedulers to coordinate unit assignments to activities while keeping track of other factors such as requirements for units to undergo maintenance and personnel to have rest and recreation. The scheduler makes heuristic choices based on 'business' rules dictated by policy documents. However, while FAMT can help to generate the FGP, the actual scheduling of units to NAS activities requires an operator to make the assignment. Manual scheduling can be time consuming and is not guaranteed to cover all the required activities in the most efficient manner. Typically, much iteration of the schedules occurs.

This paper describes a simple decision support system for scheduling RAN fleet units using automation. The paper starts by describing the employment allocation problem and the constraints conditioning the generation of an appropriate overall schedule. Then it focuses on describing the heuristic approach of the solver using a simulated annealing algorithm, together with its flexible interaction with the planning tool FAMT.

The optimisation model developed here offers schedulers the advantage of considering many candidate schedules to determine the best schedule. Apart from the obvious opportunities for labour reduction, automation also offers the advantage of providing a consistent repeatable process with robust evaluation criteria as to the quality of a FGP. In comparison to an actual FGP, the FGP generator demonstrated an improvement in the allocation of units to the NAS by about 30% on key fleet management measures, with commensurate improvements to training opportunities and preparedness. Trial runs of the FGP generator consistently found satisfactory unit allocation schedules for difficult pathological scenarios.

This research shows that automation of the FGP can lead to substantial improvements in efficiency, while maintaining satisfactory unit programmes, and ensuring that RAN fleet units can meet their on-going and emerging commitments. The solver has been integrated into FAMT and is available to schedulers. The paper notes some avenues for improvement of both FAMT and the FGP solver and the method's applicability to wider defence resource allocation problems.

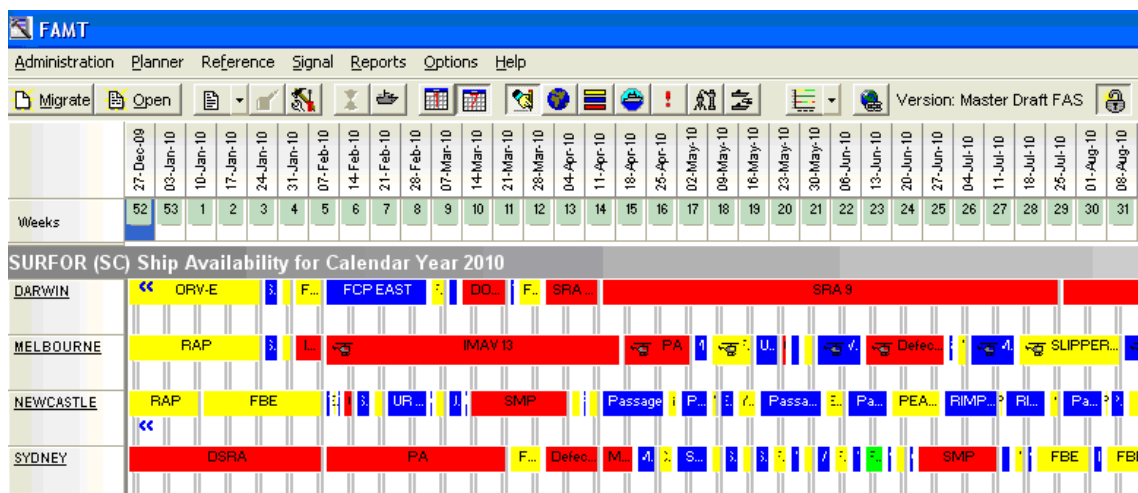
**Keywords:** *Scheduling, Optimisation, Heuristics, Simulated annealing, Naval units*

## 1. INTRODUCTION

The Royal Australian Navy's (RAN) Force Generation Plan (FGP) is a directive detailing the day-to-day operations of about 120 ships, submarines, diving teams and air units plus miscellaneous entities. Required activities to be conducted by RAN units are specified in an associated document called the Navy Activity Schedule (NAS). The NAS provides overarching guidance to RAN maritime operations planning staff on Navy's priorities and resource levels for operational commitments, preparedness levels and international and domestic engagement activities. The NAS is published approximately annually providing a rolling three year forecast of activity requirements; the commensurate FGP is released once or twice a year and then updated weekly.

Requirements for RAN fleet units to participate in activities originate from the Government, Chief of Defence Force, Commander Joint Operations (CJOPS), Chief of Navy, Commander Australian Fleet (COMAUSFLT), unit commanders and others. The Program of Major Service Activities (PMSA), managed by CJOPS, lists the significant activities, such as joint and combined exercises, that require RAN and other Australian Defence Force (ADF) units. Generally, resource demand exceeds availability. Director General Maritime Operations Division (DGMAROPS) is the authority for FGP development.

This paper describes a simple decision support system for scheduling RAN fleet units using automation. Although fleet schedulers use a customised scheduling program, the Fleet Activity Management Tool (FAMT) (Figure 1), for assigning fleet units to activities, it is not in itself an automated decision aid. Rather, the FGP is constructed manually following interactions and conferences involving COMAUSFLT and subordinate Force Commanders (commanders of groups of units) and DGMAROPS. FAMT allows the scheduler to coordinate unit assignments to activities while keeping track of other factors such as requirements for units to undergo maintenance and personnel to have rest and recreation. The scheduler makes heuristic choices based on 'business' rules dictated by policy documents. However, while FAMT can help to generate the FGP, the actual scheduling of units to NAS activities requires an operator to make the assignment. Manual scheduling can be time consuming and is not guaranteed to cover all the required activities in the most efficient manner. Typically, much iteration of the schedules occurs, but it is impossible for schedulers to consider all possible schedule combinations.



**Figure 1.** The FAMT interface in a weekly view. Units are listed in the leftmost column, and each unit's scheduled events are represented by horizontal bars. The type of event is color coded, e.g. blue for 'at sea'.

This paper describes the schedule allocation problem and the constraints conditioning the generation of an appropriate overall schedule. Changing operational requirements mean it is unlikely that automation will ever fully replace the manual scheduling processes. However, automation does offer some advantages in the ability to rapidly conduct 'what-if' analyses. This is useful in the schedule development process where schedule changes are frequently tested and an experimental environment is available in FAMT for options consideration. The paper goes on to describe the heuristic approach of the solver together with its flexible interaction with the planning tool FAMT. An experimental evaluation completes the analysis to demonstrate the solver's advantages over a purely manual approach.

## 2. SCHEDULING CONTEXT

DGMAROPS schedules fleet units on behalf of COMAUSFLT to undertake operational commitments in the Australian area and abroad. These commitments are a set of activities that seek the assignment of one or more fleet units. Activities, typically operations, exercises, test and trial periods and domestic engagement events are classified according to the objectives to which they contribute and are assigned priorities in accordance with Chief of Navy direction. Most significant activities are listed in the PMSA but some other activities are also required. These include maintenance activities and other events that are necessary to support the successful transition of units through their operational continuum. The FAMT interface dialogue for inserting an activity includes options for assigning activity classification, type and priority. The interface dialogue includes the unit demands for the activity and lists those units that are available for allocation.

Many groups of units in the RAN force structure, commanded by a force commander (see Figure 2), are semi-autonomous and in this sense, the FGP is a compilation of ‘mini’ activity schedules. Within each force command, the RAN force structure is divided into highly independent categories of units of the same type. For example, Surface Force consists of destroyers and frigates and afloat support and amphibious units. Within a category, unit operational capabilities are similar. Correspondingly, FAMT is structured in this way; i.e. schedulers will normally work with units from a single functional grouping.

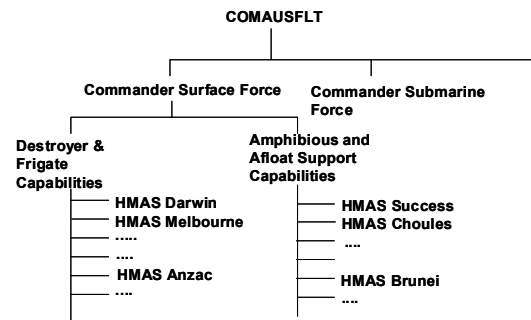


Figure 2. An extract of the RAN’s force structure

Furthermore, the activities undertaken by each functional grouping are different and depend on their capabilities and management arrangements. Therefore, this paper, and the FAMT software, describes scheduling the activities for a single functional category, such as Surface Force destroyers and frigates. However, the techniques discussed and demonstrated here are equally applicable to all other categories apart from the multi-crewed Armidale Class Patrol Boat (ACPB) group, which is supported by a separately developed stand-alone tool called NaMOS (Zadeh et al. 2009). Indeed, the FAMT software allows planning staff to choose any Force type (having, by definition, similar characteristics), other than ACPB, to schedule its units.

The following assumptions are fundamental to the approach taken in this paper:

1. All activities specify a fixed start time, duration, location (including exercise areas and geographical regions) and numbers of units and type required.
2. Activities are to be satisfied in order of their priority.
3. Operational activity for each platform is to be maintained in accordance with the business rules issued by COMAUSFLT.
4. A platform can contribute to at most one activity on any day.
5. Pre-scheduled activities are fixed in start time, duration and location and have known unit allocations that cannot be changed (although they may be incomplete).
6. The FGP is to satisfy prescribed preparedness levels for Navy unit classes.

Assumption 1 reflects current practice because most exercise commitments such as those specified for Navy in the PMSA are made years in advance. Assumptions 1, 2 and 3 constrain assignment of units by invoking trade-offs between meeting the NAS requirements and satisfying unit maintenance, preparedness and personnel requirements. Assumption 2 ensures that the highest priority activities are satisfied first. Assumption 4 ensures that units cannot be assigned to mutually exclusive activities. Assumption 5 results in known periods where schedules have to be created for each unit. Assumption 6 ensures that the FGP satisfies higher level requirements. Therefore, developing the FGP means nominating the fleet unit allocations to activities that meet unit demands while maintaining the required preparedness posture and meeting maintenance and personnel constraints. These factors will be fully described in Section 3.

## 3. SCHEDULING REQUIREMENTS

COMAUSFLT policy is to develop programmes for units that will efficiently and effectively achieve the fleet’s preparedness, training and unit and personnel sustainment requirements and those commitments directed by the Chief of Navy in the NAS.

Key FGP drivers are the conduct of operations, providing opportunities for career training and progression through collective and individual training, engaging with the community and meeting preparedness obligations. The latter, preparedness requirements, specify how many of a particular class of unit are to be available within particular timeframes at any one time.

The periods normally allotted to each unit in the FGP for major maintenance, trials, work-up and self-maintenance periods are governed by usage upkeep cycles for each class of unit and documented in maintenance availability plans produced by force commanders. These activities are scheduled well in advance, depend on the availability of repair facilities and contractors, and are generally assumed fixed in time and location. Units scheduled for major maintenance events are unavailable for other at-sea activities although their crews may be used. The availability plan can be generated in FAMT using a modified differential evolution algorithm (Storn and Price 1997). Details of this process are not reported here.

Other activities include those associated with training, leave and other individual unit events. These events prepare units to work up to required levels of preparedness and to participate in major activities such as an overseas deployment.

The PMSA consists of extended operations and exercises or deployments to undertake combined or national tasks. These activities have fixed start and completion times. Deployment preparations and post deployment activities are usually conducted in the unit's home port. Any period in which the unit operates away from its home port continuously for a period of more than sixteen weeks is considered a deployed period.

COMAUSFLT's policy objectives establish a set of criteria to provide for unit maintenance, training opportunities and activity requirements and to ensure units' crews receive reasonable time for rest and recreation in home port. These guidelines include constraints on maximum number of days that each unit may spend at sea in a two year period (OPTEMPO), a minimum amount of time in home port (PERSTEMPO), constraints on the length and frequency of unit deployments, the requirement that unit maintenance be conducted and the requirement for block leave periods to provide personnel with adequate rest and recreation. The policy objectives are used to derive ideal targets in terms of at-sea and in-port time for each unit.

Navy has an obligation to meet its preparedness requirements dictated by strategic level guidance. These require that Navy has units available in designated numbers to meet specified readiness time frames. The overall schedule must ensure that the readiness requirements are met each day. Each unit is allocated to a readiness time frame on each day depending on its employment and location. This is compared with a readiness requirement for a unit type for each day. The possibility exists that some of the activity requirements cannot be satisfied and/or that some of the daily readiness requirements may not be met because units are not available.

## **4. THE SCHEDULE OPTIMISATION PROCEDURE**

### **4.1. Optimisation Approach**

The integer linear programming (ILP) approach of Brown et al. (1990) is seminal in Navy ship scheduling. The problem presented there is similar to this one and we have used the Brown et al. approach in setting up the solution where possible. However there are differences between the approaches. Whereas Brown et al. considers only deployment activities, we consider all annual activities and preparedness requirements and develop a schedule for the complete planning period. The schedule cost structure is also different. Of most importance, the solution method is totally different in that it uses a heuristic approach.

Simulated Annealing (Kirkpatrick et al. 1983) is an applicable optimisation method for complex scheduling problems such as the problem outlined here, due to the flexibility that this search strategy offers in the face of multiple and complex criteria. Other approaches, such as the use of ILP, are often restricted to expert users and require data pre-processing whereas a heuristic approach offers flexibility in solution delivery. The aim here is to provide a deployable desktop solution for schedule automation using simulated annealing, integrated into a well used, familiar management information system.

### **4.2. Schedule Costs**

The approach taken is to view the FGP as an annual schedule of weekly activities for each platform. These are represented as a grid of units by time periods. Each grid element consists of an activity for a platform with start and end dates and location information and is able to be overlaid with NAS activity information. The fundamental activities are 'at Sea' or 'at Home Port'. These activities have a set of attributes such as

name, description, location, exercise or operation name (if applicable), start and end dates and the two personnel factor attributes OPTEMPO and PERSTEMPO representing time at sea and home port respectively.

In terms of schedule cost, each grid element receives scores based the nature of the activity being undertaken. For example, a seven day activity receives a score of four for attribute OPTEMPO and a score of three for attribute PERSTEMPO for operations where a platform spends four days at sea and returns to home port for the remainder of the week. Similarly a week long at ‘Home Port’ activity receives an OPTEMPO value of zero and a PERSTEMPO value of seven. These scores are summed over all periods for each ship. All units of measurement are in days.

The objective function measures how well the overall schedule meets the NAS requirements, preparedness requirements and COMAUSFLT policy objectives. It is a weighted sum of costs attributed to the non-achievement of the NAS, the non-achievement of preparedness requirements, plus the sum over all units of deviations from the ideal values of COMAUSFLT’s policy objectives. If the NAS requirements are not satisfied then penalties are applied based on the relative importance of the activity. Similarly, if the daily readiness requirements are not satisfied penalties are applied depending on the timeframes and unit quantities of the requirements. Penalties, based on the size of deviations, are also applied to the COMAUSFLT policy objectives. The optimisation procedure seeks to minimise the objective function; i.e. it minimises the deviations from a schedule that meets all requirements. Weights used in the objective function are chosen so as to maintain its components to the same order of magnitude. The weights can be adjusted in the software, but experiments showed that as long as the weights are chosen so that the costs of overachieving OPTEMPO and underachieving PERSTEMPO were higher than their complements, good results are achieved.

### 4.3. Schedule Optimisation

The schedule optimisation procedure first instantiates the known maintenance, holiday and exercise allocations for each unit and randomly assigns sea or home port activities for the remaining grid entries (weeks). The main optimisation technique used here is simulated annealing, based on a search neighbourhood of assigning units to activities and comprising re-assignments of sea and home port weeks, and exchanges of assignments in these respects. The optimisation technique essentially generates a new annual schedule for each unit at each iteration of the algorithm that is tested for improvement in terms of cost reduction. The neighbourhood function consists of reallocating units to activities, followed by a change of a weekly event that is not a pre-scheduled maintenance or other activity. The neighbourhood function contains an iterative procedure that identifies feasible unit allocations to the NAS requirements.

The optimisation procedure is aided by the generation of two lists at the commencement of the optimisation run. The Activity Clash List identifies those activities that are mutually exclusive because they have temporal overlaps or are in disjoint locations such that a vessel cannot finish one activity and transit to start the other (it is assumed that units take ten days to transit between disjoint locations). A second list, the Ship-Activity Tabu List, identifies those exercises that cannot be undertaken by a unit because it has preclusive commitments.

## 5. IMPLEMENTATION

The optimisation procedures discussed here were first integrated into FAMT in 2006 and updated in 2011. The interface for the optimiser component is a simple menu based system (Figure 3), allowing the user to change parameter settings, to retain elements of the current schedule or to create a completely new schedule. The progress of the optimisation procedure can be viewed and changes to the activity allocations can be witnessed ‘on the fly’.

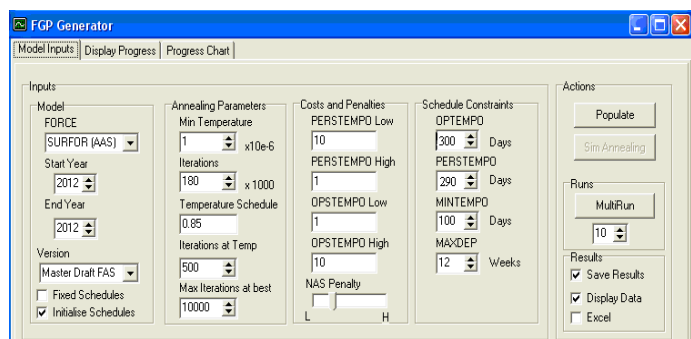


Figure 3. FAMT optimiser component

The output is supported with simple graphics of the optimisation run (Figure 4). Along the x-axis we show the iteration count for all cooling steps while y-axis shows the solution cost. The upper line presents the accepted solutions while the lower line shows the fitness of the currently best known solution. The chart in Figure 4 is

a typical result for a simulated annealing algorithm. The solver presents the user with the solution in a variety of graphical formats as in Figure 5 together with a summary of the values of key scheduling factors. The user may save the solution into the FAMT database. The solution can be amended if desired and re-saved. The solver can be run again on the saved schedule. In this way, the scheduler may preserve desired aspects of a particular solution while seeking an improved schedule.

Numerous tests of the optimisation procedure were run through its development with typical factors encountered in producing a FGP. The work to be performed during the year comprised up to 80 exercises with varying ship demands and maintenance periods of up to 52 weeks for each ship based on actual long term planning. A run of the procedure typically takes five minutes on a scheduler's workstation with 2GB of RAM running at 1.8 GHz.

This compares very favourably to the time it takes to develop a schedule through the normal manual processes. Tests of the solver against actual plans also revealed improved allocation of units of up to 30% on key measures including OPTEMPO, PERSTEMPO, NAS satisfaction and preparedness requirement satisfaction.

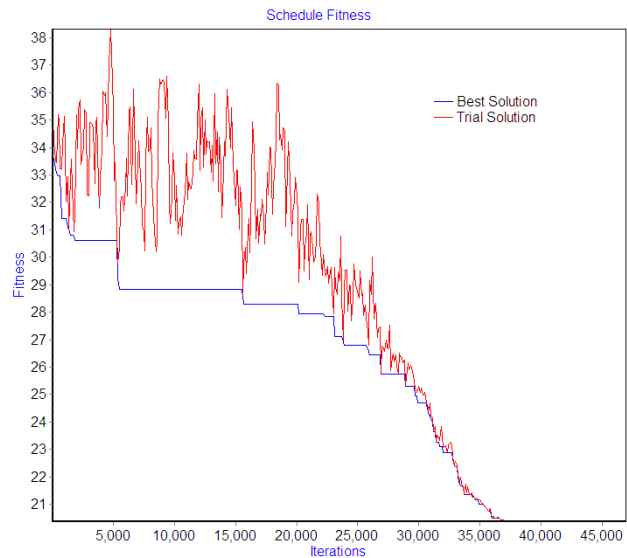


Figure 4. Output of the Annealing Procedure.

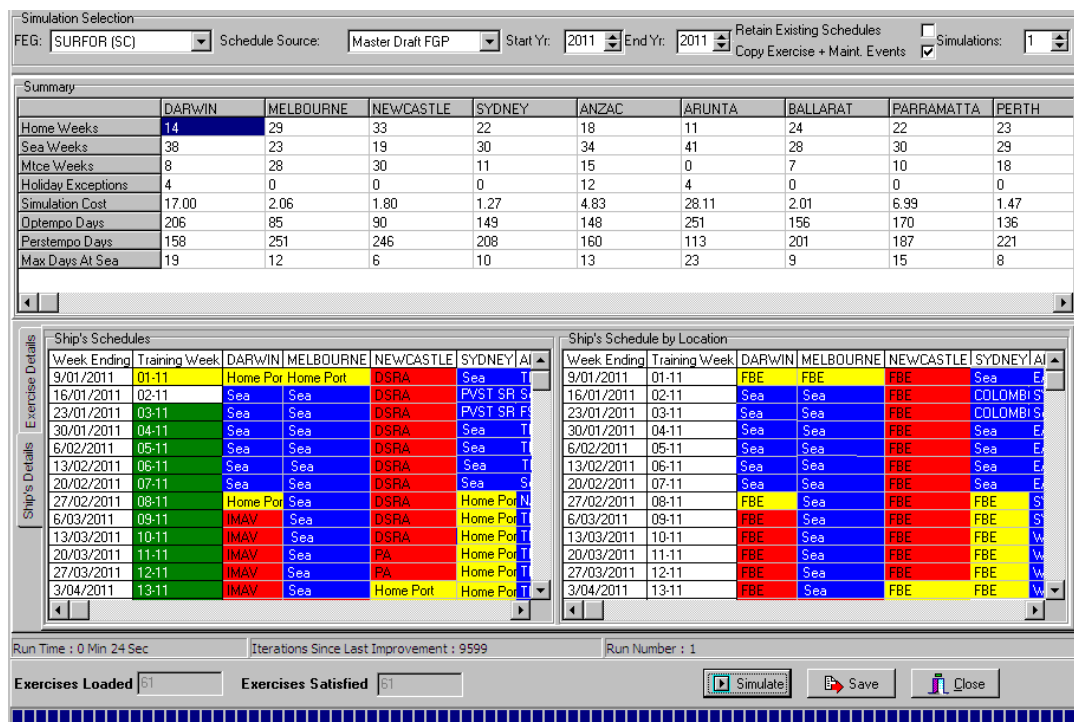


Figure 5. FAMT displaying a solution of the optimisation run. The upper table displays the key optimisation measures. All NAS requirements are achieved.

## 6. EXTENSIONS

The methodology employed here is robust enough to incorporate a range of improvements. A clear improvement in FAMT and the FGP automation procedure would be the inclusion of more detailed geographic information associated with the activities, so that they incorporate more precise transit times. Secondly, a more efficient algorithm could be designed for unit to activity allocation. Finally, further

guidance could be sourced from operations staff on conflicting business rules, ones that lack definition and those that are intuitive but are not formalised.

On the other hand, there are some easy-to-construct extensions. For example activity requirements could be expanded to include warfighting roles to be undertaken by units (unit warfighting capabilities differ across ship types). While already stated in the definition of an activity, the solver did not include role as a constraint because it is not a high priority requirement of schedulers. In the same way, the preparedness requirements could be stated and measured against each role. While the preparedness requirements are stated on a daily basis, the solver tests them weekly; this could be amended at the expense of longer run times.

Another easy extension would be to include fuel costs as part of the overall cost function. Schedule planning staff estimate fuel usage based on days at sea, but its rate of use is ship class dependent. If an upper bound for overall fuel use is stated for an individual Force grouping (e.g. frigates), then that value can be used as a target in the cost function.

It should be possible to integrate the NaMOS ACPB crew scheduling tool with the solver described here. This could be done by incorporating crewing in FGP solver or by modifying NaMOS to accept inputs directly from the NAS and adding a range of new constraints to the NaMOS engine.

A key driver of the NAS (and indeed ADF total activity) is the PMSA. The solution given here solves the optimisation problem of the allocation of Navy units to the NAS. This approach could be applied to the wider problem of optimising resource allocation to the PMSA, given available and required ADF units.

## 7. CONCLUSIONS

The scheduling problem addressed in this paper involved challenges with respect to the range and complexity of conditions addressed. It appears clear that the heuristic procedures are capable of achieving results of very good quality.

This research has shown that automation of the FGP can lead to substantial improvements in force generation planning efficiency, while satisfying the NAS, complying with policy guidance and ensuring that the fleet units can meet their on-going and emerging commitments. Improved allocations of units to the activities of the NAS offer units' crews greater opportunities for training and commensurate improvements in preparedness. The optimisation model developed here offers schedulers the advantage of considering a large number of candidate schedules to determine the best schedule. Trials using the automation procedures demonstrated significant improvements in key schedule measures over manual scheduling. Apart from the obvious opportunities for labour reduction, automation also offers the advantage of providing a consistent repeatable process with evaluation criteria reflecting documented policy guidance. Further input such as including more detailed geographic information data and extensions such as including additional tasking constraints will provide even better scheduling solutions.

## ACKNOWLEDGMENTS

Code integration of the solver into FAMT and its testing was conducted by Mr Sung Hong of Clearz Pty Ltd. Valuable comments by David Marlow and anonymous referees improved the quality of this paper.

## REFERENCES

- Brown, G.G., Goodman, C.E. and Wood, R.K. (1990) Annual scheduling of Atlantic fleet naval combatants. *Operations Research*, Vol 38, No. 2, pp 249-259
- Kirkpatrick, S., Gelatt, C.D. Jr. and Vecchi, M.P. (1983) Optimization by simulated annealing. *Science*, 220, 671-680
- Storn, R. and Price, K. (1997) Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces, *Journal of Global Optimization* 11: 341–359
- Zadeh, H.S., Storey, I. and Lenarcic, J. (2009) NaMOS; scheduling patrol boats and crews for the Royal Australian Navy. In: IEEE Aerospace Conference 2009, 1-12.