A CONSTRAINT-BASED APPROACH TO SHIP MAINTENANCE FOR THE IRISH NAVY

George Boyle

Department of Computer Science

University College Cork

Dr. James Little

Department of Computer Science

University College Cork

Dr. Joseph Manning

Department of Computer Science

University College Cork

Dr. Roman van der Krogt

Department of Computer Science

University College Cork

Abstract

The Irish Naval Service performs an annual maintenance on each of its ships, known as a "refit". During a refit, the ship is taken out of normal service and maintenance activities are performed for twenty working days. The ship must be available for patrol at the end of the refit period, so timely completion is essential.

The officer in charge of the dockyard must organise the team of workers, coordinate with the ship's staff and other naval units, and enlist the services of outside contractors when necessary. Naval refits are characterised by constraints that reflect the confined working environment of the ship, which presents numerous mechanical and safety challenges. In the extreme case, there are tasks that require an entire area of the ship to be cleared on health and safety grounds. The nature of such tasks means that delays can have significant knock-on effects. Furthermore, many of the estimates of task duration, particularly of engine repair work, cannot be fully confirmed until the engine has been dismantled and a thorough inspection conducted. These facts create considerable uncertainty in the extent of work required, and so the ability to quickly react to changes and reschedule is paramount, in order that the ship be ready to return to patrol duty.

In this paper we present a scheduling model based on constraint programming that deals with the issues of space and safety, while giving particular focus to the aspect of coping with unexpected changes. It has undergone initial evaluation on a real scenario at the Irish Naval Service dockyard.

Keywords

Constraint Programming, Navy, Maintenance Scheduling

Introduction

The Irish Naval Services has a maintenance/refit policy across all their ships, based on a 28 day period (20 working days) every year. By the time the ship arrives in dock, a clear picture of the work required to be done has been established. By 4 months previous, all requests for work on the ship have been closed and by 3 months, an initial maintenance schedule has been drawn up and will not be altered, unless for emergencies. The reason for having the work known well in advance is for orders for parts and equipment to be placed with procurement and delivery made in time. There are two types of maintenance task, those which are present because of a required service level (alternator servicing every 6000 hours, checking of anchor cables, etc.) and those which are less predictable and occur over time due to wear and tear (wiring renewal, pipe replacement). A team of specialist fitters, riggers, electricians and plumbers are employed at the dockyard to carry out the majority of tasks associated with the maintenance. However, many of the tasks can also be carried out by the ship staff in repairing and maintaining regular items, such as painting and cleaning. It is the tasks which require the use of dockyard resources that we will be concerned with in this paper. A good schedule is one in which all the tasks are completed in the 20 working day time window without the need for unplanned outsourcing.

This paper presents an innovative mathematical formulation and an effective methodology to solve this real-life naval maintenance scheduling problem. We will discuss how this model has helped in the quality of scheduling in the Irish Naval Service, and opened up several possible improvements through the ability to quickly reschedule in reaction to unexpected events. The scope of the problem discussed here has been based on the maintenance of a single ship, but we believe that the model and methodology can be extended to the whole fleet.

Related Work

Preventive maintenance models go back a long time; Wagner et al [7] present several Mathematical Programming formulations to "accomplish the required maintenance on a timely basis and concomitantly to keep to a minimum the fluctuations in the aggregate amount of manpower required." The scheduling was restricted to small datasets and covered only one skill. Yet, the objectives are as relevant today, focussing on a variety of workforce smoothing measures within a fixed time period. In the area of naval ship maintenance, Deris et al [3] look at the strategic decisions of when ships are scheduled for maintenance, without going into the detailed scheduling of the tasks. In the wider field of maintenance scheduling, Ahire et al [1] present a preventive maintenance methodology, not specific to any industry, but being characterised by tasks requiring different skills and with the objective to carry these out in the minimum time. Their approach uses an evolutionary algorithm on datasets of up to 1000 tasks and 50 workers of different skills, but all from generated data. Their focus was on performance comparison with exhaustive enumeration, but only for small problems and with simulated annealing on larger problems. Our approach uses Constraint Programming (CP) with resource scheduling and domain specific constraints from a real world problem. Our solution also tries to maximise the total net savings within a time window (although we also minimise makespan as a secondary objective). The CP-based scheduling approach has certainly found success in many other similar domains [2,4,5]. However, it seems that the area of naval maintenance scheduling is one which has not been addressed previously.

Another related area to our work is in the scheduling of ship building [6], which "needs to consider the spatial resources as well as traditional ones like manpower and machines." However, these areas are typically in the fabrication units rather than on the ship. Also, in building a ship there are strict physical constraints related to the order it is built up and fitted out. In maintenance of an already built ship, in theory any part of the ship can be worked on at any time.

There are several features about our problem which differ from other preventive maintenance problems. Locational/space constraints are particularly significant in ships, where access and operations in a confined area mean that certain tasks cannot take place at the same time. The close proximity of operators also means that some tasks which involve welding or gasses need a whole area for themselves for health and safety reasons.

Background

Currently, the maintenance schedule is produced manually by the officer in charge of the Naval Dockyard (NDY scheduler) using Microsoft Excel and Project to collect and store the task data, and deliver a graphical schedule respectively. Since many of the restrictions around what can happen and when, are contained in his head, our first step was to extend the spreadsheet to be able to articulate these as well. Now, finish-before-start, finish-before-end and no-overlap constraints are recorded per task alongside other information such as expected durations, resource and location requirements. Also recorded are how many tasks can be worked on in a particular location at one time. Finally, optional resources levels and their durations for certain tasks are also recorded in the spreadsheet.

The Model

The model was built with IBM Optimisation Studio, using the OPL modelling language. The entities we are dealing with reflect tasks, resources and a discrete timeline. The decisions amount to when each task is to take place and how many labour resources to put on it. The constraints are divided into the following categories.

Resources Constraints

There are two types of labour/equipment resources identified, those which are spread across a number of tasks at the same time (support resources) and those which are dedicated to only one task for its entire duration. The former category includes cranes and foremen. A limit is imposed on how many tasks can be supervised simultaneously. The tasks are constrained generally in their durations, although several can be done in a variable amount of time depending on the resources chosen to be applied to it.

Space Constraints

The confinement on many ships can mean that it is sometimes difficult for two or more tasks to take place the same area, or use the same access routes. The NDY scheduler has already indicated which areas these are and hence which tasks are affected. For the same type of reason, tasks involving gas or welding, even in a large area, may require other tasks to be absent.

Temporal constraints

There are some cases where one task must follow another sequentially for logical reasons. Examples are “Deammunition” before “Magazine Service”, and “Remove Turbo” before “Rebalance Turbo”.

Other Constraints

The granularity of time is half a day since this is the minimum the NDY scheduler currently allocates any task to a person. Using a scheduling model, it is easy to change this, and future work could look at the possible merits of adjusting the granularity. Of particular significance to scheduling is the task of engine service which takes the full 20 days to complete and sets a lower bound on completion time.

Objective

The objective is to maximise the number of tasks carried out internally, before any work is outsourced within the scheduling window. Beyond that, it is to finish the tasks as early as possible.

Experiment And Results

The refit for the L.E. Roisin began on Monday 2nd May 2011, and had a maintenance list which consisted of 129 tasks. Of these, 56 were assigned to the naval dockyard (NDY), with the rest being carried out by the ship's staff (SS), a specialist weapons-electrical unit (WEU), and outside contractors (OC). The NDY resources available to carry these out are shown in tables 1, 2 and 3.

Job title	*Quantity*	*Utilisation (%)*
Foreman Fitter	1	91
Chargehand Fitter	1	51
General Fitter	14	85
Turner	2	33
Fuel-based Fitter	1	0
Fuel-pump Fitter	1	25
Chargehand Plumber	1	50
Plumber	2	98
Chargehand Electrician	1	0
Electrician	3	87
Fitter General Operative	4	54
Foreman Constructor	1	23
Chargehand Welder	1	0
Welder	5	17
Joiner	2	25
Chief Rigger	1	44
Rigger	4	41
Constructor General Operative	8	55
Chargehand General Operative	1	0

Table 1: Labour Availability and Utilisation

Machine	*Quantity*	*Utilisation (%)*
Large Crane	1	75
Small Crane	1	60
Telehandler	1	10
Fork Lift	1	3

Table 2: Machine Availability

Location	*Simultaneous Tasks*	*Utilisation (%)*
Forward Engine Room	5	83
Aft Engine Room	3	45
Fo'c'stle	6	21
Afterdeck	1	23
Flag Deck	2	8
Funnel Space	1	5
Quay Wall	4	66

Table 3: Space Restrictions

The initial aim of the experiments was to make a comparison with the existing method of scheduling, in order to validate our tool and build confidence in its results. The second aim was to explore any issues or opportunities which arose around having a validated scheduling model. The results produced a schedule in which all tasks could be accommodated within the 20 day working period, with only pre-planned tasks being outsourced. Tables 1, 2 and 3 show the summary utilisation for each resource, before any changes taking place. A full analysis of the utilisation on dockyard staff is shown in Table 1. Immediately apparent is that some of the resources are being highly used, while others are not. This was not surprising to the NDY scheduler, since it reflects the profile of this particular refit, in terms of the types of tasks included. For example, there was not much welding usually required for this type of ship. These unused operatives are redeployed around the dockyard (refit work comprises just over half of all work in the dockyard).

There were few differences between our original schedule and the manual schedule. For the most part, both resulted in schedules where tasks started as early as possible given the constraints. However, having an automated tool has the advantage of being able to be adjusted quickly while keeping all the constraints enforced. Once it was shown that the model could produce a reasonable and practical result, the NDY scheduler then decided to use our schedule. In this refit, we accomplished this by sending the output to the NDY scheduler in a format that could be read into Microsoft Project.

One of the places where the model proved most useful was in reacting to unexpected events, which arise frequently. Here, we present two cases encountered during the L.E. Roisin refit, which are exemplary of the type of changes that may arrive.

Scenario 1 involved planned maintenance on the magazine sprinkler system. The job was to carry out a 12-annual and 1-annual service, and would occupy both plumbers in the dockyard for 10 days. One day into the refit, it was discovered that the 12-annual service had already been carried out. Consequently, the task would now only take the plumbers just 3 days. Under normal circumstances, this would mean one of two things for the plumbers: they would either be re-assigned to non-refit work, or they would be slack. In either case, we suggest that the ability to quickly reschedule under the new circumstances would allow for better planning, giving advance notice of the plumbers' availability for non-refit work, or slack periods as the case may be. Figure 1(a) shows the utilisation profile for the plumbers before the change, compared with after (figure 1(b)).

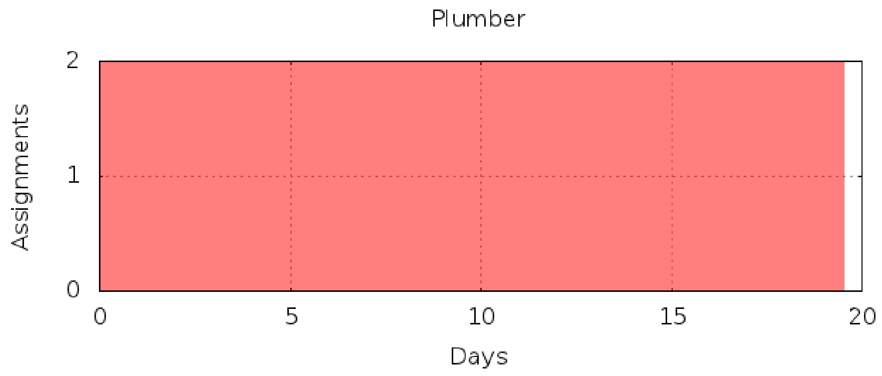

Figure 1(a): Plumber utilisation profile in original schedule

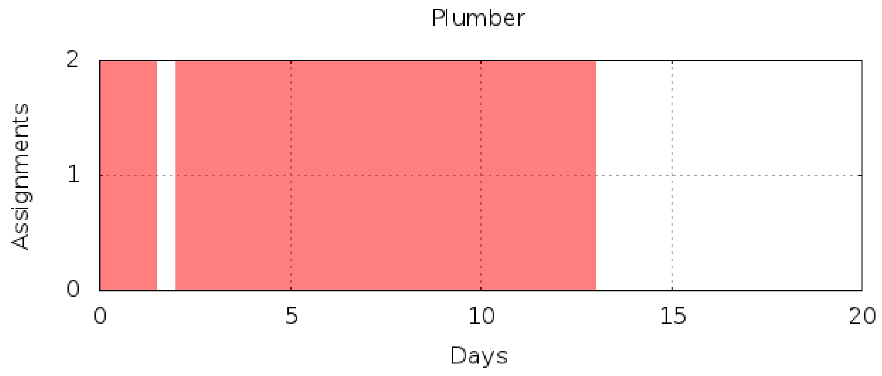

Figure 1(b): Plumber utilisation profile after reschedule

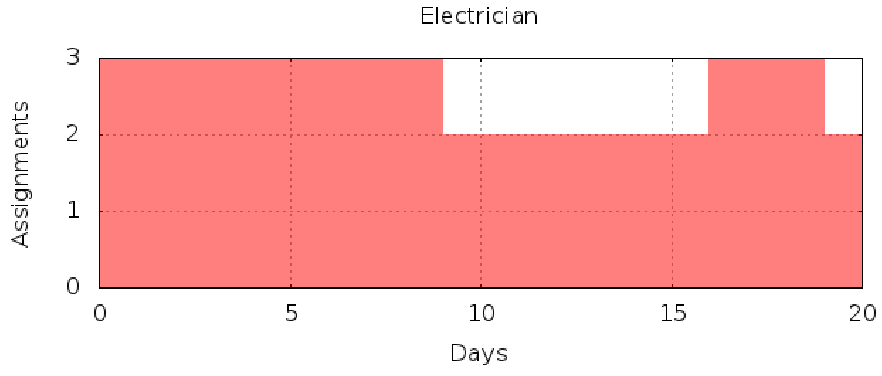

Figure 2(a): Electrician utilisation profile in original schedule

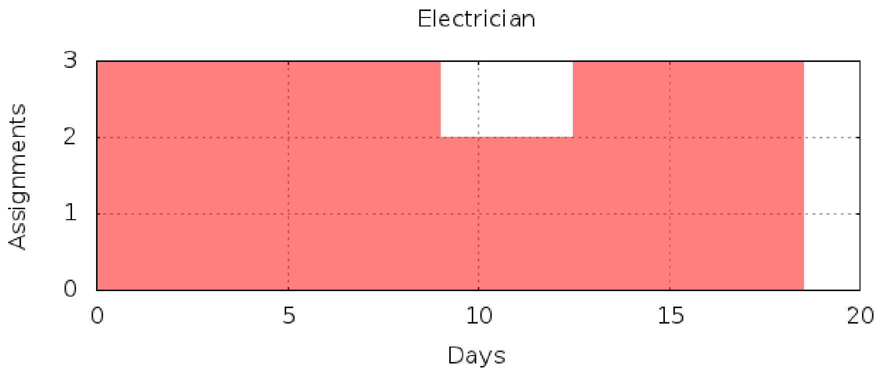

Figure 2(b): Electrician utilisation profile after reschedule

Here are the main consequences of the reschedule:

- By moving their other work around, the plumbers could finish over a week before they were originally due to finish
- The overall plumber utilisation went from 98% to 63%
- The rearrangement also allowed the electricians to finish their work 1.5 days early (shown in figure 2(a) and 2(b))
- The service of the Port Main Engine by the fitters was able to finish 1 day early

All this was possible due to a complex interaction of tasks and resources. The shortened length of the sprinkler task gave the other plumber tasks more freedom to be rearranged, which in turn, gave more flexibility in the use of the confined space, which then allowed more room for the electricians' and fitters' tasks to be rearranged. Spotting this and rescheduling for it manually are not easy; there are too many dependencies to adequately keep track of all of them. This is where the model can provide a real benefit.

The primary advantage of rescheduling in this case is that it shifted the free time for the plumbers to the end, meaning that they could be called upon if something else unexpected were to arise towards the end of the refit. If they had just been left slack, and a problem arose later, then the opportunity to front-load their work would already have passed. Furthermore, the ability to bring forward some of the electricians' and fitters' work would not have been obvious without the reschedule.

Scenario 2 concerned a job to service the one of the ship's main alternators. The dockyard must remove the alternator, which is sent away to outside contractors to be serviced, and is then replaced by the dockyard staff. On the eighth working day of the refit, it was discovered that the removal and replacement tasks would each take 4 days instead of 3, and would require an additional welder and general operative as support staff. This was due to a deck modification, unique to the LE Roisin, which was obstructing the path of the alternator's route off the ship. The decking needed to be removed, to make way for the alternator to pass out, and to be replaced afterwards.

This situation, involving task overruns, is naturally more problematic than the previous example. In general, the NDY scheduler deals with this by a right-shift of all tasks affected by the delay. If this causes the refit to exceed the required window, he must then take action by calling in staff from non-refit work, or using overtime. The model can quickly identify if the remaining tasks can be rearranged within the specified constraints without impacting the refit window, and if not, can minimise the additional resources required, resulting in less disruptions and easier planning. The consequences of the second change were not that significant in this case. The adjustment was able to be absorbed in the schedule with only minor impact on other tasks. However, this might not always be the case.

Conclusions

Rescheduling manually is difficult and time-consuming, particularly considering that the refit is already a busy time for the NDY scheduler. As we have seen, the dependence and overlap between tasks and resources, and possible knock-on effects is extensive. A manual reschedule is too large a job to be practical, especially in an environment subject to change at short notice. Our initial refit model has proved to give a very useful and flexible representation of the problem. This comes at the cost of expressing the tasks and resources in a format suitable for the model, though the NDY scheduler has stated that this is no more than the effort required for a manual schedule. Once the data has been correctly supplied, schedules can be produced virtually instantaneously. The main benefits are the ability to quickly react and reschedule when unexpected events arise, as well as the potential for exploring some hypothetical scenarios.

Future Work

The next phase of work with the Navy is due to begin in September 2011. It includes incorporating a measure of confidence in the estimates of workload. Because of the requirement for ships to be back on patrol on time, the NDY scheduler has a preference for schedules that present a low risk of overrun, given that all potential ship-stoppers must be completed, and any tasks that have been started must also be finished. It should be possible to use the confidence measure to assess the risk of overrun, providing a metric by which alternative schedules can be judged, and giving a parameter for optimisation when generating the schedules.

Acknowledgements

The authors would like to thank Lieutenant Commander Cian O'Mearain, Captain Bob Scarrot and Commodore Mark Mellett of the Irish Naval Services for the time and data provided in this research project. The funding was provided by Science Foundation Ireland reference 10/RFP/CMS2946.

References

[1] S. Ahire, G. Greenwood, A. Gupta and M. Terwilliger. (2000). Workforce-constrained preventive maintenance scheduling using evolution strategies, Decision Sciences, 31 (4):833—859

[2] T. Creemers, L.R. Giralt, J. Riera, C. Ferrarons, J. Roca and X. Corbella. (1995). Constraint-based maintenance scheduling on an electric power-distribution network, Proceedings of th Third International Conference on the Practical Applications of Prolog (PAP95), Paris, April 1995, 135—144.

[3] S.Deris, S. Omatu, H. Ohta, .S. Kutarc and P. A. Samat, (1999). Ship maintenance scheduling by genetic algorithm and constraint-based reasoning, European Journal of Operational Research 112 (3):489—502.

[4] R. van der Krogt, J. Little and H. Simonis, (2009). Scheduling in the real world: lessons learnt. Proceedings of the ICAPS 2009 Scheduling and Planning Applications Workshop, 92—98.

[5] R. van der Krogt, J. Little, K. Pulliam, S. Hanhilammi, and Y. Jin (2007). Scheduling for cellular manufacturing, Proceedings of the 13th International Conference on Principles and Practice of Constraint Programming, 105—117.

[6] J.K. Lee, K.J. Lee, H.K. Park, J.S. Hong and J.S. Lee (1997). Developing scheduling systems for Daewoo shipbuilding: DAS project, European Journal of Operational Research, 97, 380—395.

[7] H. Wagner, R.J. Giglio and R.G. Glazer (1964). Preventive maintenance scheduling by mathematical programming, Management Science, 10 (2):316—334.