Balancing Physician Workloads Under Uncertain Admissions

T. Adams
Department of Engineering Science
University of Auckland
New Zealand
thomas.adams@auckland.ac.nz

Abstract

The number of patients that physicians care for depends on how many new admissions arrive during the shifts that they work. Therefore the workloads of the physicians depend on their roster and the distribution of patients among physicians can be improved by changing the roster. This paper discusses a method for creating a roster that balances the workloads of the physicians given a number of scenarios of patient admissions, and evaluates the value of such an approach.

Key words: Rostering, Healthcare Analytics, Sample Average Approximation, YPP.

1 Introduction

In hospitals the rostering of physicians is a continuous concern, particularly in departments that are staffed 24/7. Constructing rosters manually can be very time consuming since the rules around physicians shifts and work hours are often complex and dictated by collective agreements or unions. To avoid creating rosters by hand, mathematical programs (usually mixed inter programs (MIPs)) can be formulated using the rules for the roster as constraints which automatically generate rosters.

In a health care setting it is common for these MIPs to use soft constraints as the objective function (Beaulieu et al. 2000; Ferrand et al. 2011; Bard, Shu, and Leykum 2014; Santos et al. 2014). Soft constraints is the term used when the deviation from a constraint is used in the objective function rather than including the constraint itself. They are frequently used when including all of the constraints makes the problem infeasible, or in this context when there is no roster that can conform to all of the rules.

In situations where all of the rules can be met it is not necessarily obvious what the objective should be when creating a roster. In this work we assess the quality of a roster by determining its effect on the number of patients each physician in the roster is responsible for, otherwise known as their workload. We focus on the physicians' workloads because the General Medicine (GM) departments in the hospitals in the

Waitemata District Health Board (WDHB) have identified balancing the workloads of their physicians as a priority.

The concern of the GM departments stems from the concept of continuity of care. Continuity of care refers to a patient having a continuous relationship with their health care provider, in this case a physician. Fragmentation of care occurs when a patient is transferred from one physician to another, this has been shown to impact both patient satisfaction and length of stay (Hjortdahl and Laerum 1992; Epstein et al. 2010). Some fragmentation is unavoidable if a patient needs to be seen by one specialist and then another, however fragmentation can also occur when one physician is caring for many more patients than another and transfers some of their patients to the other physician. It is this fragmentation that the GM departments believe can be reduced by balancing the physicians' workloads. In particular they are interested in minimising the largest difference in workloads between any two physicians at any point in time.

The workloads of the physicians depend on when they work admitting shifts and how many patients arrive during their admitting shifts. An admitting shift is one during which the physician who is working it admits new patients to a ward, and the patient becomes part of the physician's workload. Not all shifts that the physicians work are responsible for admitting new patients. The physicians' workloads can therefore be influenced by changing their roster, and a roster that minimises the largest difference in workloads can be found if the following are known: the rules that the roster must adhere to; the number of patients admitted and discharged each day; and a mapping between the shifts and the physicians workloads (i.e. which shifts are admitting shifts).

The above approach will produce a roster that is optimal for the admissions data it is given, however it may not be optimal (or even close it it) if the admissions were different. One way of taking the uncertainty of patient admissions into account is by considering different scenarios for admissions, where each scenario has a different number of patients arriving on each day. An example of the cumulative number of patients in GM for 20 scenarios (light gray) generated from a single year of historical data (black) is given in Figure 1.

The construction of such scenarios is not the focus of this work. Instead we address finding a roster that minimises the difference in workload, given a set of scenarios.

2 Method

The problem is modelled as a two-stage stochastic program, the generic form of which is given below: The first stage contains the decision variables that determine the roster and the constraints that the roster must obey. In the second stage the roster is evaluated on each of the patient admission scenarios and the workloads of the physicians are calculated to find the maximum difference.

The background, formulation and a method for solving the single scenario version of this problem are the subject of (Adams et al. 2017), and a manuscript currently in revision for Operations Research in Health Care.

Given a single scenario of patient admissions the objective of the model is to minimise the largest difference in workloads at any point in the planning horizon. When more than one scenario is considered we can instead minimise the expected

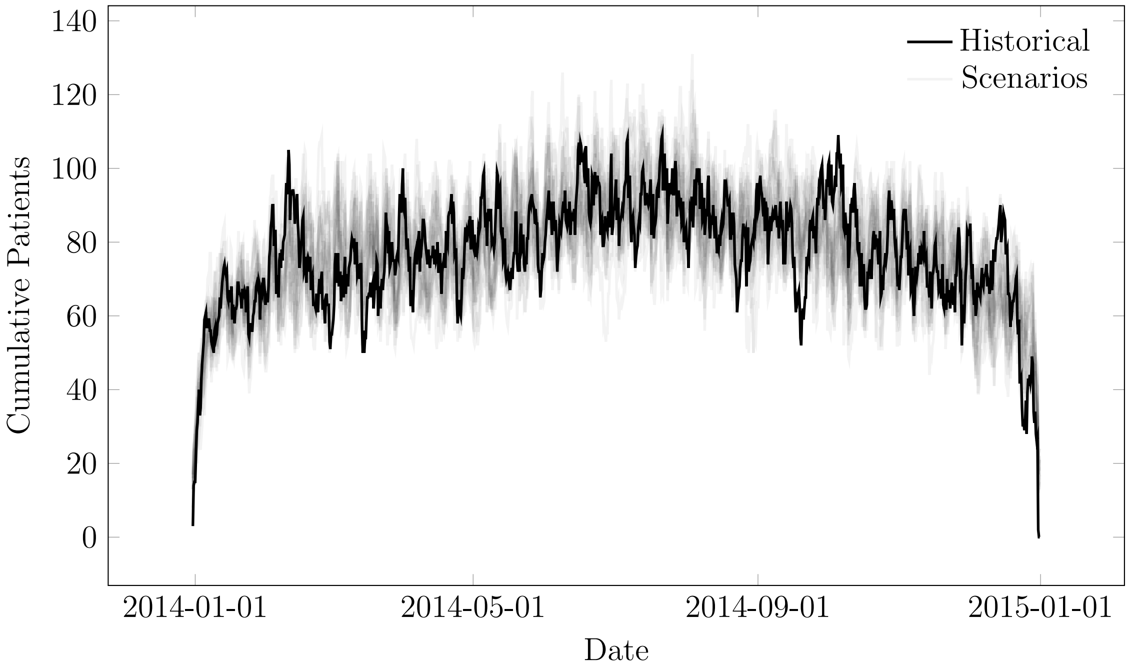

Figure 1: Cumulative Patients from Historical Data and 20 Generated Scenarios

value of the largest difference in workloads across the scenarios. If however the hospital is more worried about the worst scenarios a measure called conditional value at risk (CVaR) can be used. Given a distribution of values, CVaR at the β level is the expected value of the worst β% of the values. Therefore we can instead find the roster that minimises the expected value of the largest difference in workloads in the worst β% of scenarios.

To solve the two-stage stochastic program with a given number of scenarios a Benders decomposition is used within a branch-and-cut framework. In the benders decomposition the master problem is the first stage problem of finding a roster that conforms to the rules (and satisfies any Benders cuts that have been added already). The sub-problems are evaluating the roster on each of the scenarios and calculating the maximum difference in workloads. A single branch-and-bound tree is maintained for the master problem, and when an integer solution is found at a node it is used to generate benders cuts via the sub-problems.

2.1 The Sample Average Approximation Method

The sample average approximation (SAA) method works by solving a reduced version of the stochastic optimisation problem repeatedly. The problem is reduced by only considering a sample of the scenarios, rather than all of them at once.

First M samples, each with N scenarios, are generated. Each sample is used to construct a two-stage stochastic optimisation problem that has all of the first stage variables and constraints of the actual problem, but only considers the scenarios in M for the second stage. By solving these M problems we obtain M objective values and candidate solutions. The mean of these objective values provides a statistical estimate of a lower bound on the actual optimal objective function value (Mak, Morton, and Wood 1999; Norkin, Pflug, and Ruszczyski 1998).

The objective function value of the M solutions is evaluated for an independent sample N', which can be larger than N as we do not have to solve a problem with N' scenarios. The solution with the best objective value is then evaluated on another independent sample of size N' to estimate the upper bound.

An estimate of the optimality gap, and its variance, can then be formed (Kleywegt, Shapiro, and Homem-de Mello 2002). If the gap (or its variance) are above the desired threshold the process is restarted and at least one of M, N, N' is increased. For the purposes of this paper M remained constant at 10, N was increased from 5 to 75 in multiples of 5, and N' was set to $5N$. In addition 95% confidence of a 1% optimality gap was used as the threshold. This means that when the algorithm terminates we have found a solution that we are 95% sure is within 1% of the optimal solution, or have reached the limiting sample size of 75. This means during the SAA algorithm a maximum of 1500 scenarios could be used.

3 Results

This section compares the performance of four rosters. The first roster labeled 'Proposed' is a roster that was proposed by a rostering group at WDHB. The second roster was created by solving the SAA problem with the objective to minimise the expected value of the maximum difference in workloads, and is labeled 'Expected Value'. The third roster, labeled 'CVaR' was created by solving the SAA problem with the objective to minimise the expected value of the maximum difference in workloads in the worst 10% of scenarios, or CVaR @ 10%. The final roster was created by averaging the 1500 possible scenarios used to solve the SAA problems for the second and third rosters to create an 'expected scenario', and then finding the roster that minimised the largest difference in workloads for that scenario. This was called the 'Expected Scenario' roster.

Figure 2 shows how the estimates of the upper and lower bounds for the SAA problems progressed. For both the upper and lower bounds 95% confidence estimates are displayed, this means that we are 95% confident that the actual upper bound is less than the one displayed, and 95% confident that the lower bound is more than the one displayed. Note that only 12 iterations were required for the expected value objective, this is because at this point the 95% estimate for the optimality gap was below 1% and the algorithm was terminated. On the other hand, the CVaR objective reached the limiting sample size of 75 without reducing the estimate of the optimality gap below 1%. In general it can be seen that the bounds were further apart for the CVaR problem and particularly in the first few iterations. This may be due to the CVaR problem only considering the objective function value in the worst 10% of scenarios. The sample sizes therefore need to be larger before an accurate approximation of the objective function is obtained.

The four rosters were evaluated on 500 scenarios to generate a distribution of maximum workload differences for each roster, histograms of these distributions are given in Figure 3, and summary statistics of the distributions are given in Table 1. The 500 scenarios used in this testing were independent from any of the scenarios used to create the Expected Value, CVaR, and Expected Scenario rosters. The distributions for the CVaR and Expected Scenario rosters are identical because solving the respective problems produced the same roster.

The difference in the mean maximum workload difference between the expected value roster and the expected scenario roster was very small. This means that benefit

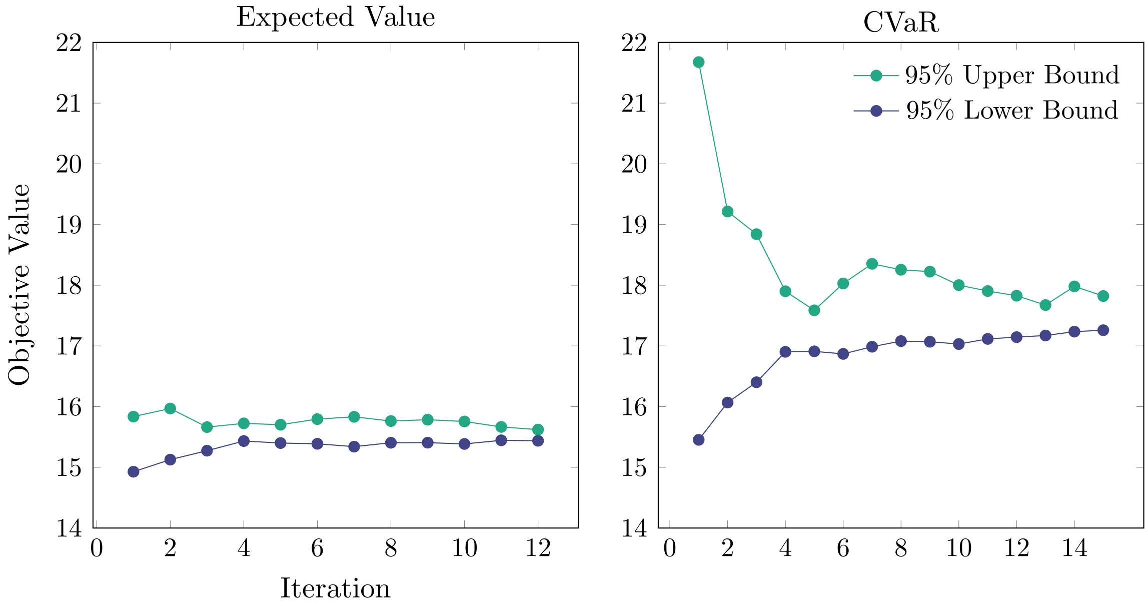

Figure 2: 95% Upper and Lower Bound Estimates for SAA Problems

of solving the SAA problem was only a 0.38% reduction in the mean maximum workload difference. In addition using the CVaR objective produced a roster that performed just 0.17% better than the expected value roster in the worst 10% of scenarios, which suggests that the expected value roster already performs very well in these scenarios. The CVaR roster did however perform 5.95% better in the very worst case.

Table 1: Summary Statistics of Maximum Workload Difference Distribution on 500 Test Scenarios

Roster	Minimum	Mean	Maximum	CVaR @ 10%
Proposed	15.32	17.80	21.59	20.01
Expected Value	13.42	15.59	20.25	17.65
CVaR	13.50	15.65	19.11	17.62
Expected Scenario	13.50	15.65	19.11	17.62

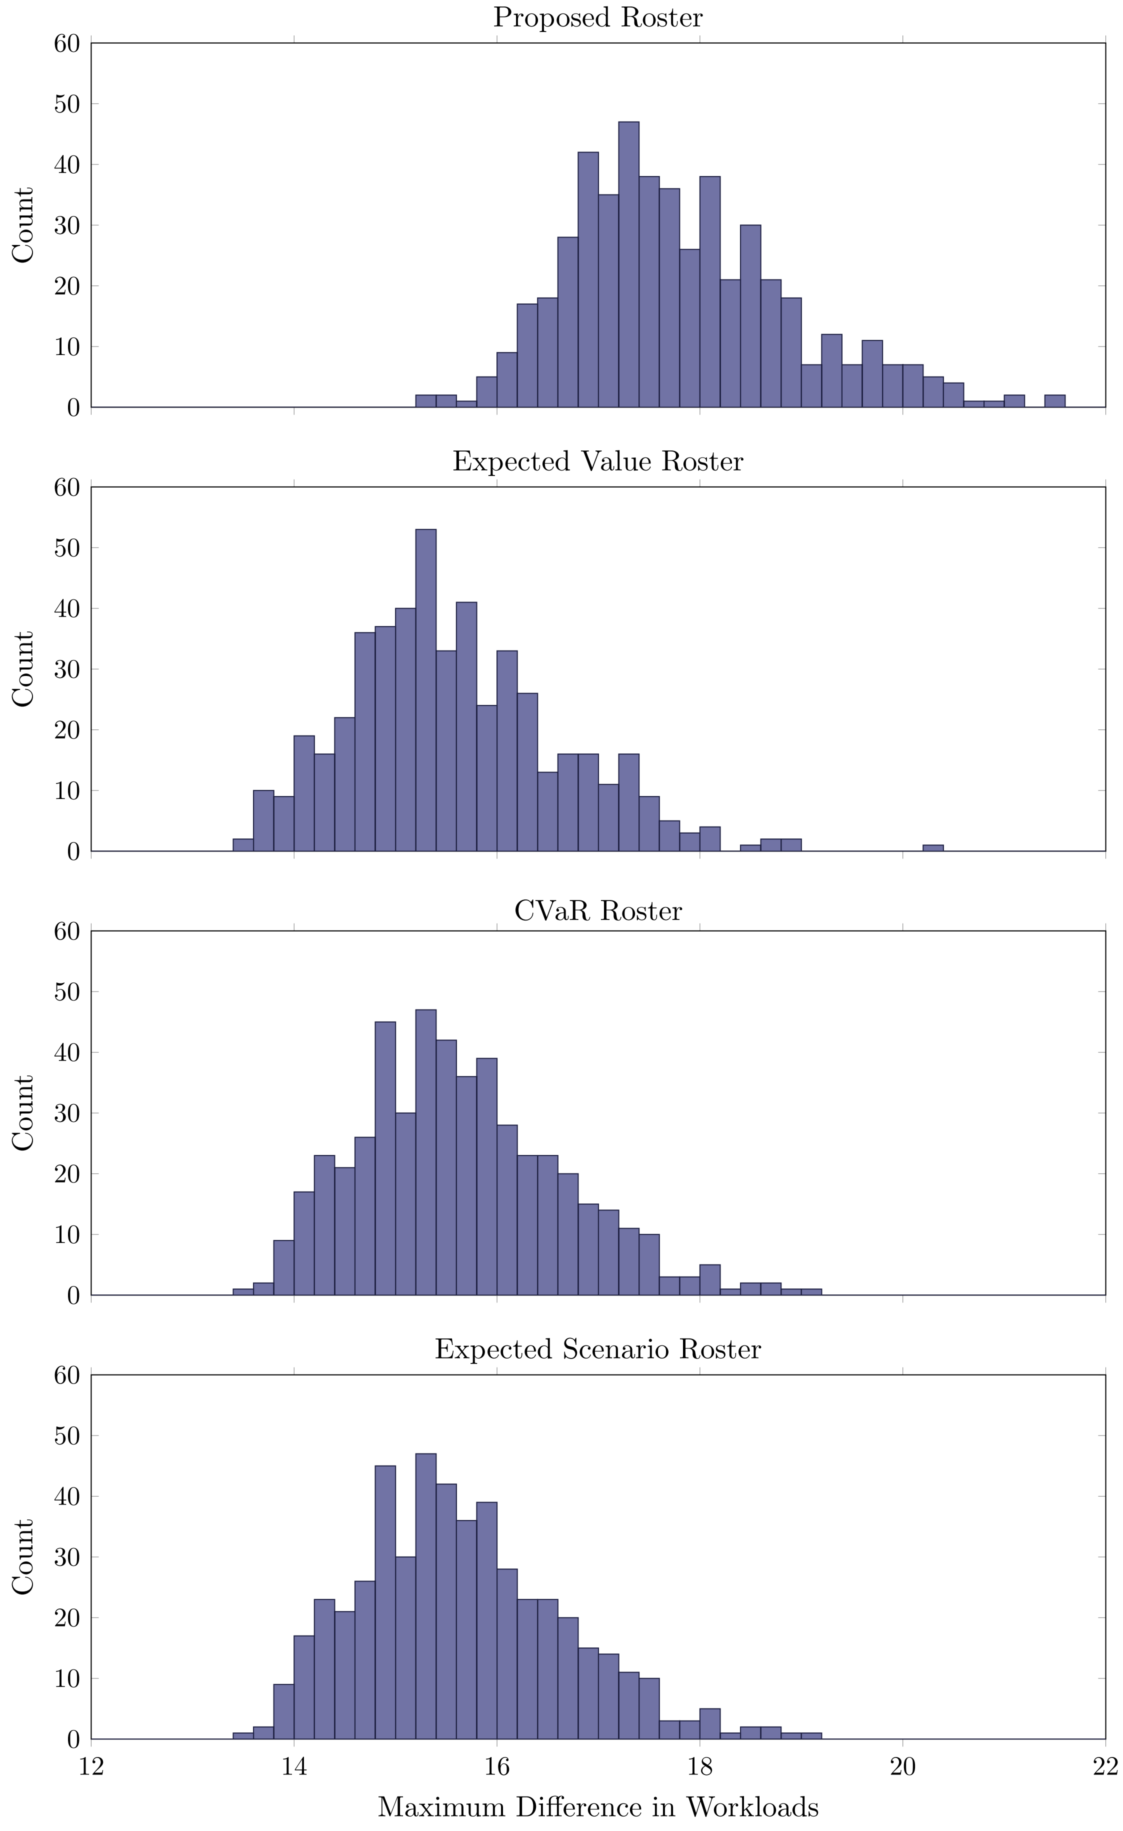

Figure 3: Histograms of Maximum Workload Difference Over 500 Scenarios

4 Conclusions

A method for creating a roster that balances the workloads of physicians given uncertainty in the number of admissions was presented. An SAA problem was solved which takes this uncertainty into account by using different scenarios of admissions. Compared to a roster that was created by averaging the uncertainty, the roster from the SAA problem provided little benefit. Further research is required to determine whether this is due to the nature of the uncertainty represented in the scenarios, or a characteristic of the type of rostering problem examined, or other factors.

Two objective functions were compared for the SAA problem: the first minimised the expected value of the maximum difference in workloads; the second minimised the expected value of the maximum difference in workloads in the worst 10% of scenarios. Very little difference was observed between the two approaches, once again further research is needed to determine the reason for this.

References

Adams, Thomas, Michael O'Sullivan, Jonathan Christiansen, Paul Muir, and Cameron Walker. 2017. "Rostering general medicine physicians to balance workload across inpatient wards: a case study." *BMJ Innovations* 3 (2): 84–90 (April).

Bard, Jonathan F., Zhichao Shu, and Luci Leykum. 2014. "A network-based approach for monthly scheduling of residents in primary care clinics." *Operations Research for Health Care* 3 (4): 200–214 (December).

Beaulieu, Huguette, Jacques A. Ferland, Bernard Gendron, and Philippe Michelon. 2000. "A mathematical programming approach for scheduling physicians in the emergency room." *Health care management science* 3 (3): 193–200.

Epstein, Kenneth, Esther Juarez, Anne Epstein, Kathy Loya, and Adam Singer. 2010. "The impact of fragmentation of hospitalist care on length of stay." *Journal of Hospital Medicine* 5 (6): 335–338 (July).

Ferrand, Y., M. Magazine, U. S. Rao, and T. F. Glass. 2011. "Building Cyclic Schedules for Emergency Department Physicians." *Interfaces* 41 (6): 521–533 (December).

Hjortdahl, P., and E. Laerum. 1992. "Continuity of care in general practice: effect on patient satisfaction." *BMJ* 304 (6837): 1287–1290.

Kleywegt, Anton J., Alexander Shapiro, and Tito Homem-de Mello. 2002. "The sample average approximation method for stochastic discrete optimization." *SIAM Journal on Optimization* 12 (2): 479–502.

Mak, Wai-Kei, David P. Morton, and R. Kevin Wood. 1999. "Monte Carlo bounding techniques for determining solution quality in stochastic programs." *Operations Research Letters* 24:47–56.

Norkin, Vladimir I., Georg Ch. Pflug, and Andrzej Ruszczyski. 1998. "A Branch and Bound Method for Stochastic Global Optimization." *Math. Program.* 83 (3): 425–450 (November).

Santos, Haroldo G., Tlio A. M. Toffolo, Rafael A. M. Gomes, and Sabir Ribas. 2014. "Integer programming techniques for the nurse rostering problem." *Annals of Operations Research*, May.