Evolutionary Multiobjective Optimization for Dynamic Hospital Resource Management

Anke K. Hutzschenreuter[ab] Peter A.N. Bosman[b] Han La Poutré[ab]

[a] *Department Industrial Engineering & Innovation Sciences, Eindhoven University of Technology, The Netherlands,* A.K.Hutzschenreuter@tue.nl

[b] *Centre for Mathematics and Computer Science, P.O. Box 94079, 1090 GB Amsterdam, The Netherlands,* Peter.Bosman@cwi.nl, Han.La.Poutre@cwi.nl

The full version of this paper has been accepted for publication at the *2009 Evolutionary Multi-Criterion Optimization conference — EMO*.

1 Introduction

Today, many hospitals face great demands to reduce costs and improve quality of service, e.g. by reducing patient waiting times. In several European countries this is due to the introduction of a free market health care system, like in the Netherlands. To decrease costs, the occupancy rates of resources need to be increased. Increasing resource utilization, however, may lead to bottlenecks that cause the blocking of patient flows and thus increase patient waiting times. Therefore, the efficient allocation of resources is an important issue.

For the optimization of resource management three outcome measures are of interest to the hospital: patient throughput, i.e. the number of patients discharged from the hospital after treatment, resource costs and back-up capacity usage. In order to accommodate patients at the appropriate care level, a hospital unit may open an extra bed or transfer a patient temporarily to another unit until a bed becomes available. A well-designed hospital resource allocation features high patient throughput at low resource costs and back-up capacity usage. A trade-off is needed between these conflicting objectives.

To optimize hospital resource management, we apply strategy (or policy) optimization. Policies are parameterized functions that return an allocation decision. In cooperation with domain experts from the Catharina Hospital Eindhoven (CHE), the Netherlands, we designed policies that enable dynamic resource allocations. The policies can be easily understood by health care professionals which is important for implementation and understanding in practice.

Due to the stochastic patient processes and the actual patient flow being the result of resource availability, an analytical evaluation of a resource allocation is not feasible. Furthermore, changing the structure of the patient pathways or the underlying probability distributions is non-trivial in an analytical model. Therefore, we use a simulation tool to evaluate a resource allocation. This simulation tool has been validated recently and shown to provide an accurate representation of the real world [1]. The decision space comprises allocations for each unit in a hospital. Due to the need of a complex simulation tool for evaluation, the huge decision space and multiple conflicting objectives, evolutionary algorithms (EAs) were chosen as solution technique, as they are known to be very powerful for multi-objective (MO) optimization.

2 Policies

In our simulation model, we consider the number of allocated resources as free decision variables (i.e. control variables that impact the performance of the system). The policies that we design therefore return allocations. More detailed mathematical formulations are given in the full version of this paper.

Static resource allocation policy. The static allocation policy allocates a fixed number of resources to the different hospital units.

State-dependent allocation policy. A static allocation can do well in a relatively stable environment. This condition, however, does not hold in hospitals due to the stochastic patient treatment processes. Therefore, we consider dynamic policies that return an allocation for the units in the network, given the current state of the units. This allows the resources (i.e. decision variables) to switch and track changes in the environment (i.e. the optimization problem) dynamically.

Our state-dependent allocation policy is determined by five parameters: a base resource allocation, two adjustments, and two utilization thresholds. The current resource allocation is decreased by the first adjustment value if the resource utilization rate is below the lower utilization threshold. Similarly, an increase by the second adjustment value occurs if the resource utilization rate is above the upper utilization threshold.

The state-dependent strategy assumes a large supply and stock of beds, enabling the concurrent in- and decrease in resource capacity at the different units. In reality, however, bed availability is restricted by the available staff, in particular the number of personnel needed per bed at a specific unit. Staff schedules need to be fixed at least several weeks in advance. The use of stand-by personnel is not common in the hospital domain. Therefore, a direct implementation of the policy is often not practically feasible. For a more practically relevant setting, we additionally used a realistic exchange mechanism that is based on fixed personnel resources. The resources are exchanged among the hospital units to meet the current local need.

3 Optimization results

For optimization of the dynamic and complex multi-objective resource allocation problem, we apply the SDR-AVS-MID𝔼A algorithm that is known to be an efficient evolutionary algorithm for MO problems.

Results of optimizing the resource allocation strategies are shown in Figure 1. F_0 is the mean total throughput of patients, defined as the number of patients discharged from the hospital after treatment. F_1 is the mean total resource costs. F_2 is the mean total weighted back-up capacity usage.

The exact values for back-up capacity usage are of minor importance and a categorization is therefore sufficient for the representation of the optimization results, giving multiple 2D Pareto fronts.

The results show that the benchmarks obtained from hospital practice are dominated by all policies proposed in this paper. Moreover, the dynamic resource allocation policies show higher performance compared to the static allocation policies.

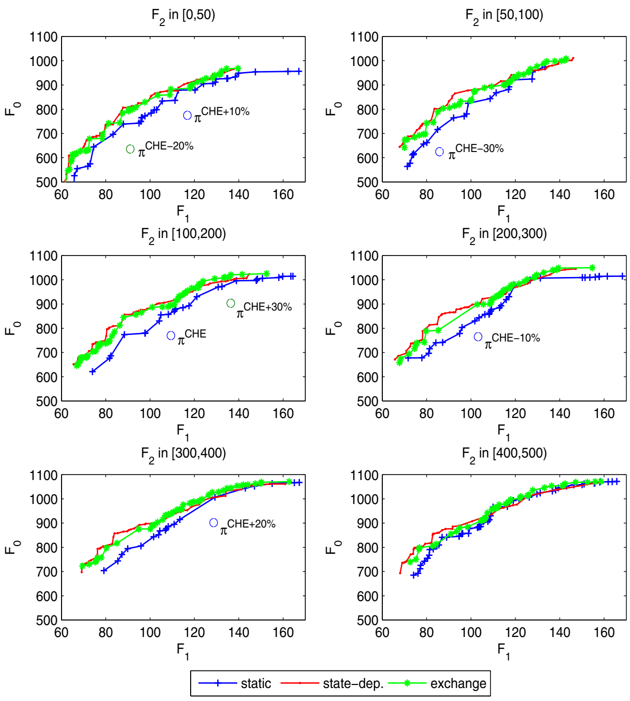

Fig. 1. Pareto-fronts for static and dynamic allocation policies. Also shown are benchmark points from CHE.

For F2-values of above 300 and F1-values higher than 120, the static and dynamic policies show similar performance. This can be explained by the small extent and frequency of allocation adjustments of the dynamic policies obtained for these F1 and F2 values. Since additional demand for care can be met by using back-up capacity, less allocation adjustments are necessary in these cases. The bed exchange policies show slightly lower performance compared to the state-dependent policies. The difference is due to the interaction between the hospital units due to which required allocation adjustments cannot always be fully undertaken.

4 Future work

In future work, we will develop allocation strategies that use more advanced anticipation models of the time-dependency effects. Furthermore, we will consider alternative orderings of care levels in the bed exchange mechanism. Also, we will further explore the settings of the EA in relation to the above extensions.

References

[1] A.K. Hutzschenreuter, P.A.N. Bosman, I. Blonk-Altena, J. van Aarle, and H. La Poutré. Agent-based patient admission scheduling in hospitals. In L. Padgham et al., editors, *Autonomous Agents and Multiagent Systems — AAMAS*, pages 45–54. IFAAMAS, 2008.