Enhanced Hospital Resource Management using Anticipatory Policies in Online Dynamic Multi-objective Optimization

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

[a] *Faculty of Industrial Engineering, Eindhoven University of Technology, The Netherlands*
[b] *Center for Mathematics and Computer Science, Amsterdam, The Netherlands, Peter.Bosman@cwi.nl, Han.La.Poutre@cwi.nl*
[c] *Department of Information and Computing Sciences, Utrecht University, The Netherlands*

The full version of this paper has been accepted for publication at the *2010 Genetic and Evolutionary Computation — GECCO* conference.

1 Introduction

Optimization problems in real-world settings are often dynamic and need to be solved online, i.e. as time goes by. Moreover, current decisions may have future consequences, i.e. time-dependence, requiring anticipation of future situations to make well-informed decisions [1]. In real-world optimization problems, often more than one objective needs to be optimized at the same time. The optimum then is no longer a single solution but a set of solutions called the Pareto front. Population-based methods such as evolutionary algorithms (EAs) are among the state-of-the-art in solving MO optimization problems. There is currently no literature on how to perform anticipation in the optimization of multiple dynamically-changing objectives.

Hospital resource management is concerned with the efficient allocation of resources, i.e. operating room (OR) time slots and hospital care beds. Here, uncertain patient arrivals and treatment processes cause resource usage to behave stochastically and dynamically. Moreover, multiple care units have to be considered as patient treatment processes often involve more than one care unit. Also, resources are often shared by different treatment processes. An extensive, validated simulation is available (see [3]) for four types of surgical and emergency patient pathways that involve the OR and six postoperative care units. Patient routing between units may deviate from pathways if no resource is available at the required unit.

We show how, for the dynamic MO problem of hospital resource management, anticipatory solutions can be obtained, providing better results than non-anticipatory solutions.

2 Definition and Approach

We denote the m objectives by $f_i(x)$, $i \in 0, 1, ..., m-1$ and without loss of generality aim to minimize all objectives. In dynamic optimization, all objectives are explicitly a function of a time parameter t in addition to the decision variables $\boldsymbol{x}$. Mathematically, we optimize the functional

$$\min_{\boldsymbol{x}(t)} \left\{ \left(\int_0^{t^{\text{end}}} f_0(\boldsymbol{x}(t), t)\, dt, \ldots, \int_0^{t^{\text{end}}} f_{m-1}(\boldsymbol{x}(t), t)\, dt \right) \right\} \tag{1}$$

This integral represents optimization over time. We determine the resource allocation using policies, π, i.e. parameterized functions that return an allocation decision given the current situation, and optimize the parameters of allocation policies using a MO evolutionary algorithm (MOEA).

Our policy, $\pi(t, u)$, for unit $u \in U$ is an iterative step-function that returns an allocation decision based on current problem variables. A (non-)anticipatory policy considers the (current) future interval and the outcome depends on the utilization observed (at t^{now}) in $[t^{now}; t^{now} + \Delta t]$.

We consider three objectives: maximize the mean total number of patients discharged after complete treatment ($f_0(\pi)$), minimize the mean resource costs ($f_1(\pi)$) and minimize the mean back-up capacity usage ($f_2(\pi)$). Detailed definitions can be found in [3]. Finally, increased availability of resource capacity is coupled with increased demand for care of related treatment types, modeling increased attraction of patients following capacity enlargement or reputation improvement of a hospital unit.

3 Experiments

The MOEA we used in our experiments is SDR-AVS-$\mathbb{MIDEA}$ [2]. The genes correspond to the policy parameters as described above. The parameters for SDR-AVS-$\mathbb{MIDEA}$ were based on [2]. The simulation is run with $\pi(t, u),\ u \in U$.

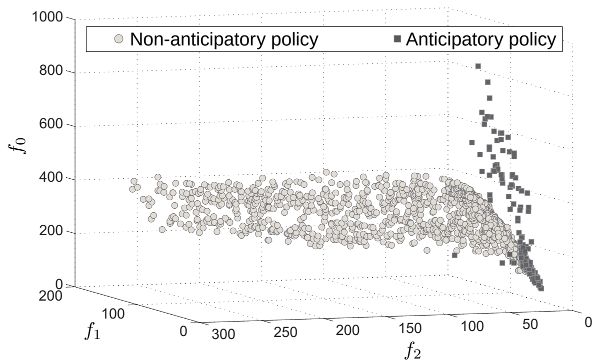

Figure 1: Three-dimensional Pareto fronts computed from all runs

Figure 1 shows Pareto fronts computed over all runs. We note that without anticipation, already better results than current real-world practice are obtained. With anticipation however, the anticipatory policy picks up on the fact that increased availability of resource capacity entails an increase in demand for care of this type of treatment. Performances are comparable for $f_1 \lesssim 70$. For larger f_1 values the use of predicted occupancy information in the anticipatory policy considerably improves throughput. The increased frequency in demand for care established by the optimized anticipatory policy also results in more efficient usage of beds.

4 Summary and Conclusions

We have focused on the design of EAs for solving dynamic MO optimization problems. Typically, optimization has to be performed online. Because for real-world problems running an optimizer online, i.e. in real time, is not always feasible, we have studied the feasibility of an offline policy-based approach. We note that our results may be slightly optimistic as we slightly strengthened the coupling of arrival processes to resource capacity to better illustrate the contribution of anticipation; in practice this coupling may be weaker and less instantaneous. However, our results illustrate, for the first time, that time-dependence can be detected and dealt with successfully for dynamic optimization problems with multiple objectives using anticipatory policies.

References

[1] P. A. N. Bosman and J.A. La Poutré. Learning and anticipation in online dynamic optimization with evolutionary algorithms: The stochastic case. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO-2007*, pages 1165–1172, New York, New York, 2007. ACM Press.

[2] P. A. N. Bosman and D. Thierens. Adaptive variance scaling in continuous multi-objective estimation-of-distribution algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference - GECCO-2007*, pages 500–507, New York, New York, 2007. ACM Press.

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